# Large State Spaces and Self-supervision in Reinforcement Learning. 

par

## Ahmed Touati

Département d'informatique et de recherche opérationnelle Faculté des arts et des sciences

Thèse présentée à la Faculté des arts et des sciences en vue de l'obtention du grade de Philosophiæ Doctor (Ph.D.) en informatique

August, 2021
© Ahmed Touati, 2021.

Université de Montréal
Faculté des arts et des sciences

Cette thèse intitulée:

# Large State Spaces and Self-supervision in Reinforcement Learning. 

présentée par:
Ahmed Touati
a été évaluée par un jury composé des personnes suivantes:
Yoshua Bengio, président-rapporteur
Pascal Vincent, directeur de recherche
Ioannis Mitliagkas, membre du jury
Matthieu Geist, examinateur externe

To my beloved parents Najoua and Mohammed, to whom I owe an eternal gratitude.

## Résumé

L'apprentissage par renforcement (RL) est un paradigme d'apprentissage orienté agent qui s'intéresse à l'apprentissage en interagissant avec un environnement incertain. Combiné à des réseaux de neurones profonds comme approximateur de fonction, l'apprentissage par renforcement profond (Deep RL) nous a permis récemment de nous attaquer à des tâches très complexes et de permettre à des agents artificiels de maîtriser des jeux classiques comme le Go, de jouer à des jeux vidéo à partir de pixels et de résoudre des tâches de contrôle robotique.

Toutefois, un examen plus approfondi de ces remarquables succès empiriques révèle certaines limites fondamentales. Tout d'abord, il a été difficile de combiner les caractéristiques souhaitables des algorithmes RL, telles que l'apprentissage hors politique et en plusieurs étapes, et l'approximation de fonctions, de manière à obtenir des algorithmes stables et efficaces dans de grands espaces d'états. De plus, les algorithmes RL profonds ont tendance à être très inefficaces en raison des stratégies d'exploration-exploitation rudimentaires que ces approches emploient. Enfin, ils nécessitent une énorme quantité de données supervisées et finissent par produire un agent étroit capable de résoudre uniquement la tâche sur laquelle il est entrainé. Dans cette thèse, nous proposons de nouvelles solutions aux problèmes de l'apprentissage hors politique et du dilemme exploration-exploitation dans les grands espaces d'états, ainsi que de l'auto-supervision dans la RL.

En ce qui concerne l'apprentissage hors politique, nous apportons deux contributions. Tout d'abord, pour le problème de l'évaluation des politiques, nous montrons que la combinaison des méthodes populaires d'apprentissage hors politique et à plusieurs étapes avec une paramétrisation linéaire de la fonction de valeur pourrait conduire à une instabilité indésirable, et nous dérivons une variante de ces méthodes dont la convergence est prouvée. Deuxièmement, pour l'optimisation des politiques, nous proposons de stabiliser l'étape d'amélioration des politiques par une régularisation de divergence hors politique qui contraint les distributions stationnaires d'états induites par des politiques consécutives à être proches les unes des autres.

Ensuite, nous étudions l'apprentissage en ligne dans de grands espaces d'états et nous nous concentrons sur deux hypothèses structurelles pour rendre le problème traitable : les environnements lisses et linéaires. Pour les environnements lisses, nous proposons un algorithme en ligne efficace qui apprend activement un partitionnement adaptatif de l'espace commun en zoomant sur les
régions les plus prometteuses et fréquemment visitées. Pour les environnements linéaires, nous étudions un cadre plus réaliste, où l'environnement peut maintenant évoluer dynamiquement et même de façon antagoniste au fil du temps, mais le changement total est toujours limité. Pour traiter ce cadre, nous proposons un algorithme en ligne efficace basé sur l'itération de valeur des moindres carrés pondérés. Il utilise des poids exponentiels pour oublier doucement les données qui sont loin dans le passé, ce qui pousse l'agent à continuer à explorer pour découvrir les changements.

Enfin, au-delà du cadre classique du RL, nous considérons un agent qui interagit avec son environnement sans signal de récompense. Nous proposons d'apprendre une paire de représentations qui mettent en correspondance les paires état-action avec un certain espace latent. Pendant la phase non supervisée, ces représentations sont entraînées en utilisant des interactions sans récompense pour encoder les relations à longue portée entre les états et les actions, via une carte d'occupation prédictive. Au moment du test, lorsqu'une fonction de récompense est révélée, nous montrons que la politique optimale pour cette récompense est directement obtenue à partir de ces représentations, sans aucune planification. Il s'agit d'une étape vers la construction d'agents entièrement contrôlables.

Un thème commun de la thèse est la conception d'algorithmes RL prouvables et généralisables. Dans la première et la deuxième partie, nous traitons de la généralisation dans les grands espaces d'états, soit par approximation de fonctions linéaires, soit par agrégation d'états. Dans la dernière partie, nous nous concentrons sur la généralisation sur les fonctions de récompense et nous proposons un cadre d'apprentissage non-supervisé de représentation qui est capable d'optimiser toutes les fonctions de récompense.

## Abstract

Reinforcement Learning (RL) is an agent-oriented learning paradigm concerned with learning by interacting with an uncertain environment. Combined with deep neural networks as function approximators, deep reinforcement learning (Deep RL) allowed recently to tackle highly complex tasks and enable artificial agents to master classic games like Go, play video games from pixels, and solve robotic control tasks. However, a closer look at these remarkable empirical successes reveals some fundamental limitations. First, it has been challenging to combine desirable features of RL algorithms, such as off-policy and multi-step learning with function approximation in a way that leads to both stable and efficient algorithms in large state spaces. Moreover, Deep RL algorithms tend to be very sample inefficient due to the rudimentary explorationexploitation strategies these approaches employ. Finally, they require an enormous amount of supervised data and end up producing a narrow agent able to solve only the task that it was trained on. In this thesis, we propose novel solutions to the problems of off-policy learning and exploration-exploitation dilemma in large state spaces, as well as self-supervision in RL.

On the topic of off-policy learning, we provide two contributions. First, for the problem of policy evaluation, we show that combining popular off-policy and multi-step learning methods with linear value function parameterization could lead to undesirable instability, and we derive a provably convergent variant of these methods. Second, for policy optimization, we propose to stabilize the policy improvement step through an off-policy divergence regularization that constrains the discounted state-action visitation induced by consecutive policies to be close to one another.

Next, we study online learning in large state spaces and we focus on two structural assumptions to make the problem tractable: smooth and linear environments. For smooth environments, we propose an efficient online algorithm that actively learns an adaptive partitioning of the joint space by zooming in on more promising and frequently visited regions. For linear environments, we study a more realistic setting, where the environment is now allowed to evolve dynamically and even adversarially over time, but the total change is still bounded. To address this setting, we propose an efficient online algorithm based on weighted least squares value iteration. It uses exponential weights to smoothly forget data that are far in the past, which drives the agent to keep exploring to discover changes.

Finally, beyond the classical RL setting, we consider an agent interacting with its environments without a reward signal. We propose to learn a pair of representations that map state-action pairs to some latent space. During the unsupervised phase, these representations are trained using reward-free interactions to encode long-range relationships between states and actions, via a predictive occupancy map. At test time, once a reward function is revealed, we show that the optimal policy for that reward is directly obtained from these representations, with no planning. This is a step towards building fully controllable agents.

A common theme in the thesis is the design of provable RL algorithms that generalize. In the first and the second part, we deal with generalization in large state spaces either by linear function approximation or state aggregation. In the last part, we focus on generalization over reward functions and we propose a task-agnostic representation learning framework that is provably able to solve all reward functions.

## Keywords-Mots-clés

reinforcement learning, markov decision process, artificial agent, offpolicy learning, function approximation, exploration-exploitation trade-off, self-supervision, generalization
apprentissage par renforcement, processus de décision markovien, agent artificiel, apprentissage hors-politique, approximation de fonction, compromis exploration-exploitation, auto-supervision, généralisation

## Contents

1 Introduction ..... 1
1.1 Research Contributions ..... 3
1.1.1 Off-Policy Learning ..... 3
1.1.2 Exploration v.s. Exploitation Dilemma ..... 3
1.1.3 Unsupervised Learning in RL ..... 4
1.2 List of Excluded Contributions ..... 5
2 Background ..... 7
2.1 Discounted Markov Decision Process ..... 7
2.1.1 The Model ..... 7
2.1.2 Policies and Value Functions ..... 8
2.1.3 Optimal Policies and Optimal Value Functions ..... 9
2.2 Episodic Markov Decision Processes ..... 9
2.3 Dynamic Programming ..... 10
2.3.1 Bellman Operators ..... 10
2.3.2 Value Iteration ..... 11
2.3.3 Policy Iteration ..... 12
2.4 Temporal Difference Learning ..... 13
2.4.1 Policy Evaluation ..... 14
2.4.2 Policy Learning ..... 15
2.5 Function Approximation ..... 16
2.5.1 Value-based Methods ..... 16
2.5.2 Policy Gradient ..... 17
2.6 Exploitation-Exploration Dilemma ..... 18
2.6.1 Online Performance ..... 20
3 Multi-step Off-policy Learning with Function Approximation ..... 22
3.1 Prologue to the Contribution ..... 22
3.1.1 Article Details ..... 22
3.1.2 Context ..... 22
3.1.3 Paper Abstract ..... 22
3.1.4 Recent Developments ..... 23
3.2 Introduction ..... 23
3.3 Tabular Off-policy Methods ..... 25
3.4 Off-policy instability with function approximation ..... 26
3.5 Convergent gradient off-policy algorithms ..... 29
3.6 Convergence Rate Analysis ..... 30
3.7 Related Work and Discussion ..... 34
3.8 Experimental Results ..... 35
3.8.1 Evidence of instability in practice ..... 35
3.8.2 Comparison with existing methods ..... 36
3.9 Conclusion ..... 38
4 Stable Policy Optimization via Off-Policy Divergence Regularization ..... 40
4.1 Prologue to the Contribution ..... 40
4.1.1 Article Details ..... 40
4.1.2 Context ..... 40
4.1.3 Paper Abstract ..... 41
4.1.4 Recent Developments ..... 41
4.2 Introduction ..... 41
4.3 Conservative Update Approaches ..... 43
4.4 Theoretical Insights ..... 44
4.5 Off-policy Formulation Of Divergences ..... 46
4.6 A Practical Algorithm Using Adversarial Divergence ..... 48
4.7 Related Work ..... 52
4.8 Experiments And Results ..... 53
4.8.1 Important Aspects Of PPO-DICE ..... 54
4.8.2 Results On Atari ..... 56
4.8.3 Results On OpenAI Gym MuJoCo ..... 57
4.9 Conclusion ..... 57
5 Online Learning in Smooth Markov Decision Processes ..... 58
5.1 Protologue To The Contribution ..... 58
5.1.1 Article Details ..... 58
5.1.2 Context ..... 58
5.1.3 Paper Abstract ..... 59
5.1.4 Recent Developments ..... 59
5.2 Introduction ..... 60
5.3 Related Work ..... 61
5.4 Problem Statement ..... 62
5.4.1 Episodic Reinforcement Learning and Regret ..... 62
5.4.2 Metric Space ..... 63
5.5 The ZoomRL algorithm ..... 64
5.6 Main results ..... 68
5.6.1 Result For The Misspecified Case ..... 69
5.7 Proof Outline ..... 70
5.7.1 Regret Analysis ..... 72
5.8 Conclusion ..... 75
6 Online Learning in Non-stationary Linear Markov Decision Processes ..... 76
6.1 Prologue to the Contribution ..... 76
6.1.1 Article Details ..... 76
6.1.2 Context . ..... 76
6.1.3 Paper Abstract ..... 77
6.1.4 Recent Developments ..... 77
6.2 Introduction ..... 77
6.3 Problem Statement ..... 78
6.3.1 Notation ..... 78
6.3.2 Non-Stationary Reinforcement Learning and Dynamic Re- gret ..... 78
6.3.3 Linear Markov Decision Processes ..... 79
6.4 The Proposed Algorithm ..... 80
6.5 Non-stationary Linear Bandits ..... 81
6.6 Theoretical guarantee of OPT-WLSVI ..... 85
6.6.1 Unknown variation budget ..... 86
6.7 Technical Highlights ..... 88
6.8 Related Work ..... 90
6.9 Conclusion ..... 92
7 Learning One Representation to Optimize All Rewards ..... 93
7.1 Prologue to the Contribution ..... 93
7.1.1 Article Details ..... 93
7.1.2 Context ..... 93
7.1.3 Paper Abstract ..... 93
7.2 Introduction ..... 94
7.3 Problem and Notation ..... 96
7.4 Encoding All Optimal Policies via the Forward-Backward Repre- sentation ..... 96
7.5 Learning and Using Forward-Backward Representations ..... 98
7.6 Experiments ..... 103
7.6.1 Environments and Experimental Setup ..... 103
7.6.2 Goal-Oriented Setting: Quantitative Comparisons ..... 104
7.6.3 More Complex Rewards: Qualitative Results ..... 105
7.6.4 Embedding Visualizations ..... 106
7.7 Related work ..... 107
7.8 Extended Results: Approximate Solutions and General Goals ..... 109
7.8.1 The Forward-Backward Representation With a Goal or Feature Space ..... 110
7.8.2 Existence of Exact $F B$ Solutions, Influence of Dimension $d$, Uniqueness ..... 112
7.8.3 Approximate Solutions Provide Approximately Optimal Policies ..... 114
7.8.4 $F$ and $B$ as Successor and Predecessor Features of Each Other, Optimality for Rewards in the Span of $B$ ..... 119
7.8.5 Estimating $z_{R}$ from a Different State Distribution at Test Time ..... 122
7.8.6 A Note on the Measure $M^{\pi}$ and its Density $m^{\pi}$ ..... 123
7.9 Conclusion ..... 124
8 Conclusion ..... 125
8.1 Summary of Contributions ..... 125
8.2 Future Research ..... 126
8.2.1 Towards Fully Controllable Agents ..... 126
8.2.2 Optimization for RL ..... 126
8.2.3 Statistical RL with General Function Approximation ..... 127
A Appendix from Chapter 3 ..... 128
A. 1 Proof of Proposition 1 ..... 128
A. 2 Proof of Proposition 2 ..... 130
A. 3 Proof of Proposition 3 ..... 130
A. 4 Convergence Rate Analysis ..... 131
A. 5 True on-line equivalence ..... 135
B Appendix For Chapter 4 ..... 138
B. 1 Proof of Lemma 2 ..... 138
B. 2 Score Function Estimator of the gradient with respect to the policy ..... 138
B. 3 Comparison with AlgaeDICE ..... 139
B. 4 Additional Empirical Results on MuJoCo ..... 140
B. 5 Hyperparameters ..... 141
C Appendix For Chapter 5 ..... 142
C. 1 Omitted proofs for the Lipschitz setting ..... 142
C.1.1 Proof of Lemma 3 ..... 142
C.1.2 Proof of Lemma 4 ..... 143
C.1.3 Proof of Lemma 5 ..... 143
C.1.4 Proof of lemma 7 ..... 147
C.1.5 Bounding $\sum_{h=1}^{H} \sum_{k=1}^{K} \xi_{h+1}^{k}$ ..... 149
C. 2 Misspecified Setting: Approximately Lipschtiz Case ..... 150
C.2.1 Recursive Formula of $\widehat{Q}_{h}^{k}(B)-Q_{h}^{\star}(s, a)$ ..... 150
C.2.2 Bounding of $Q_{h}^{k}(B)-Q_{h}^{*}(s, a)$ ..... 151
C.2.3 High Probability Bound On The Sampling Noise ..... 151
C.2.4 Approximate Optimism Of $Q$-values ..... 152
C.2.5 Upper Bound of $\widehat{Q}_{h}^{k}(B)-Q_{h}^{\star}(s, a)$ ..... 153
C.2.6 Regret Analysis ..... 154
C. 3 Technical Lemmas ..... 155
C.3.1 Few Reminders on Probability Theory ..... 155
D Appendix for Chapter 6 ..... 156
D. 1 Technical Gaps in Published Bandit Papers ..... 156
D. 2 Regret Reanalysis of D-LINUCB ..... 157
D. 3 Regret Analysis of OPT-WLSVI and Proof Outline ..... 160
D.3.1 Single Step Error Decomposition ..... 160
D.3.2 High Probability Bound on the Transition Variance ..... 161
D.3.3 Optimism ..... 162
D.3.4 Final Regret Analysis ..... 163
D. 4 Missing Proofs of Regret Analysis of OPT-WLSVI ..... 165
D.4.1 Linearity of $Q$-values: Lemma 8 ..... 165
D.4.2 Non-Stationarity Bias ..... 165
D.4.3 Single Step Error Decomposition ..... 168
D.4.4 Boundness of iterates ..... 169
D.4.5 Transition Concentration ..... 170
D.4.6 Single-Step High Probability Upper Bound ..... 171
D.4.7 Optimism ..... 172
D. 5 Technical Lemmas ..... 173
E Appendix for Chapter 7 ..... 177
E. 1 Proofs ..... 177
E. 2 Experimental Setup ..... 189
E.2.1 Environments ..... 189
E.2.2 Architectures ..... 190
E.2.3 Implementation Details ..... 191
E.2.4 Experimental results ..... 192

## List of Tables

3.1 Properties of different off-policy algorithms for policy evaluation. ..... 26
3.2 Convergence results for gradient-based TD algorithms shown in previous work [Sutton et al., 2009b,c, Liu et al., 2015, Wang et al., 2017, Lakshminarayanan and Szepesvári, 2017, Dalal et al., 2017]. $\bar{\theta}_{k}$ stand for the Polyak-average of iterates: $\bar{\theta}_{k} \triangleq \frac{\sum_{k} \alpha_{k} \theta_{k}}{\sum_{k} \alpha_{k}}$. Our algo- rithms achieve $O(1 / k)$ without the need for projections or Polyak averaging. ..... 33
4.1 Mean final reward and 1 standard error intervals across 10 seeds for Atari games evaluated at 10 M steps. ..... 55
6.1 Comparison of our regret bound with state-of-the-art bounds for both linear bandits and linear MDPs. $d$ is the dimension of the features space, $H$ is the planning horizon of the MDP, $K$ is the number of episodes and $\Delta$ is the variation budget. When we go from a bandit setting to MDPs, the work of Jin et al. [2020b] in the stationary case and our work in the non-stationary case incur an extra $d^{1 / 2}$ factor and $d^{3 / 8}$ respectively. Zanette et al. [2020b] achieve a linear dependence on $d$ in the stationary case but their proposed algorithm is computationally intractable. ..... 86
B. 1 A complete overview of used hyper parameters for all methods. ..... 141
E. 1 Hyperparameters of the FB algorithm ..... 193
E. 2 Hyperparameters of the goal-oriented DQN algorithm ..... 193

## List of Figures

1.1 Standard Reinforcement Learning setting ..... 1
2.1 Restaurant selection example. Exploitation: Go to favorite restau- rant. Exploration: Try a new restaurant ..... 18
2.2 n-chain environment ..... 19
3.1 Two-state counterexample. We assign the features $\left\{(1,0)^{\top},(2,0)^{\top},(0,1)^{\top},(0,2)^{\top}\right\}$ to the state-action pairs $\{(1$, right $),(2$, right $),(1$, left $),(2$, left $)\}$. The target policy is given by $\pi($ right $\mid \cdot)=1$ and the behavior policy is $\mu($ right $\mid \cdot)=0.5$. ..... 28
3.2 Baird's counterexample. The combination of linear function approximation with TB and RETRACE leads to divergence (left panel) while the proposed gradient extensions GTB and GRETRACE converge (right panel). ..... 35
3.3 In the 2-states counterexample of section 3.4 showing that the gradient-based TB and RETRACE converge while TB and RE- TRACE diverge. ..... 36
3.4 Each curves shows the 5th percentile of NMSE (over all possible combination of step-size values) achieved by each algorithm for different values of $\lambda$. ..... 37
3.5 Comparison between the best empirical MSPBE obtained by each algorithm for different values of $\lambda$. Only $\operatorname{GRETRACE}(\lambda)$ and AB- $\operatorname{TRACE}(\lambda)$ are showed here because the other algorithms do not have the same operators and hence not the same MSPBE. Note that MSPBEs depend on $\lambda$. Thus, MSPBEs are not directly com- parable across different values of $\lambda$. Both $\operatorname{GRETRACE}(\lambda)$ and AB- $\operatorname{TRACE}(\lambda)$ have very similar performances. AB-TRACE $(\lambda)$ per- forms slightly better. ..... 38
3.6 Comparison of empirical performance of $G Q(\lambda)$, $\operatorname{AB}-\operatorname{TRACE}(\lambda)$, $\operatorname{GRETRACE}(\lambda)$ and $\operatorname{GTB}(\lambda)$ on an off-policy evaluation task in Mountain Car domain. Each box plot shows the distribution of the NMSE achieved by each algorithm after 2000 episodes for different values of $\lambda$. NMSE distributions are computed over all the possible combinations of step-size values $\left(\alpha_{k}, \eta_{k}\right) \in$ [0.001, $0.005,0.01,0.05,0.1]^{2}$. ..... 39
4.1 Comparison of $\chi^{2}$ and KL divergences for PPO-DICE for two ran-
domly selected environments in OpenAI Gym MuJoCo and Atari,
respectively. We see that KL performs better than $\chi^{2}$ in both set-
tings. Performance plotted across 10 seeds with 1 standard error
shaded. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 53
4.2 Varying $\lambda$ in Hopper_v2, 10 seeds, 1 standard error shaded. PPODICE is somewhat sensitive to $\lambda$ value, but the theoreticallymotivated adaptive version works well.55

4.3 Comparison of PPO-DICE with clipped loss $L^{\text {clip }}$ and without $L$.
We see that clipping the action loss is crucial for good performance. 55
4.4 Results from OpenAI Gym MuJoCo suite in more complex domains, with 10 seeds and 1 standard error shaded. Results on the full suite of environments can be found in B.4.56

5.1 Left: a possible partition for 2-dimensional state-action space. For a
given state $s$, we show in orange the three relevant balls $A, B$ and
C. Right: For the three relevant balls, we show how the index is
constructed based on the interpolation between $Q$-value estimates of
each ball. The gray piecewise linear curve corresponds to the function
$: a \rightarrow \min _{B^{\prime}}\left\{\widehat{\mathrm{Q}}_{h}\left(B^{\prime}\right)+L \cdot \operatorname{dist}\left((s, a), x_{B^{\prime}}\right)\right\}$ for a given state. ..... 66

7.1 Comparative performance of FB for different dimensions and DQN in
the FetchReach. Left: success rate averaged over 20 randomly selected
goals as function of the first 100 training epochs. Right: success rate
averaged over 20 random goals after 800 training epochs.
7.2 Comparative performance of FB for different dimensions and DQN in Ms. Pacman. Left: success rate averaged over 20 randomly selected goals as function of the first 200 training epochs. Right: success rate averaged over the goal space after 800 training epochs. ..... 105
7.3 Heatmap of $\max _{a} F\left(s, a, z_{R}\right)^{\top} z_{R}$ for $z_{R}=B(\xi)$ Left: $d=25$. Right: $d=75.106$
7.4 Trajectories generated by the $F^{\top} B$ policies for the task of reaching a target position (star shape $\underset{\sim}{r}$ ) while avoiding forbidden positions (red circle). ..... 106
7.5 Trajectories generated by the $F^{\top} B$ policies for the task of reaching the closest among two equally rewarding positions (star shapes $\begin{gathered}\text { rus) . (Opti- }\end{gathered}$ mal $Q$-values are not linear over such mixtures.) ..... 107
7.6 Contour plot of $\max _{a \in A} F\left(s, a, z_{R}\right)^{\top} z_{R}$ in Continuous Maze. Left: for the task of reaching a target while avoiding a forbidden region, Right: for two equally rewarding targets. ..... 107
7.7 Visualization of FB embedding vectors on Continuous Maze after projecting them in two-dimensional space with t-SNE. Left: the states to be mapped. Middle: the $F$ embedding. Right: the $B$ embedding. The walls appear as large dents; the smaller dents correspond to the number of steps needed to get past a wall. . . . ..... 108
B. 1 Our method with KL divergences in comparison to PPO and TRPO on MuJoCo, with 10 seeds. Standard error shaded
E. 1 Discrete maze: Comparative performance of FB for different dimensions and DQN. Left: the policy quality averaged over 20 randomly selected goals as function of the training epochs. Right: the policy quality averaged over the goal space after 800 training epochs.
E. 2 Continuous maze: Comparative performance of FB for different dimensions and DQN. Left: the success rate averaged over 20 randomly selected goals as function of the training epochs. Right: the success rate averaged over 1000 randomly sampled goals after 800 training epochs.194
E. 3 FetchReach: Comparative performance of FB for different dimensions and DQN. Left: the success rate averaged over 20 randomly selected goals as function of the training epochs. Right: the success rate averaged over 1000 randomly sampled goals after 800 training epochs.
E. 4 Ms. Pacman: Comparative performance of FB for different dimensions and DQN. Left: the success rate averaged over 20 randomly selected goals as function of the training epochs. Right: the success rate averaged over the 184 handcrafted goals after training epochs. Note that FB-50 and F-100 have been trained only for 200 epochs.194
E. 5 Distance to goal of FB for different dimensions and DQN as function of training epochs. Left: Continuous maze. Right: FetchReach. 195
E. 6 Discrete Maze: Heatmap plots of $\max _{a \in A} F\left(s, a, z_{R}\right)^{\top} z_{R}$ (left) and trajectories of the Boltzmann policy with respect to $F\left(s, a, z_{R}\right)^{\top} z_{R}$ with temperature $\tau=1$ (right). Top row: for the task of reaching a target while avoiding a forbidden region, Middle row: for the task of reaching the closest goal among two equally rewarding positions, Bottom row: choosing between a small, close reward and a large, distant one.196
E. 7 Continuous Maze: Contour plots plot of $\max _{a \in A} F\left(s, a, z_{R}\right)^{\top} z_{R}$ (left) and trajectories of the $\varepsilon$ greedy policy with respect to $F\left(s, a, z_{R}\right)^{\top} z_{R}$ with $\varepsilon=0.1$ (right). Left: for the task of reaching a target while avoiding a forbidden region, Middle: for the task of reaching the closest goal among two equally rewarding positions, Right: choosing between a small, close reward and a large, distant one ..... 197
E. 8 Ms. Pacman: Trajectories of the $\varepsilon$ greedy policy with respect to $F\left(s, a, z_{R}\right)^{\top} z_{R}$ with $\varepsilon=0.1$ (right). Top row: for the task of reaching a target while avoiding a forbidden region, Middle row: for the task of reaching the closest goal among two equally rewarding positions, Bottom row: choosing between a small, close reward and a large, distant one.198
E. 9 Full series of frames in Ms. Pacman along the trajectory generated by the $F^{\top} B$ policy for the task of reaching a target position (star shape $s$ ) while avoiding forbidden positions (red circle).199
E. 10 Discrete maze: Visualization of FB embedding vectors after projecting them in two-dimensional space with t -SNE. Left: the $F$ embedding for $z=0$. Right: the $B$ embedding. Note how both embeddings recover the foor-room and door structure of the original environment. The spread of $B$ embedding is due to the regularization that makes B closer to orthonormal.
E. 11 Continuous maze: Visualization of FB embedding vectors after projecting them in two-dimensional space with t-SNE. Left: the states to be mapped. Middle: the $F$ embedding. Right: the $B$ embedding.
E. 12 Ms. Pacman: Visualization of FB embedding vectors after projecting them in two-dimensional space with t-SNE. Left: the agent's position corresponding to the state to be mapped. Middle: the $F$ embedding for $z=0$. Right: the $B$ embedding. Note how both embeddings recover the cycle structure of the environment. F acts on visual inputs and B acts on the agent's position.
E. 13 Continuous maze: visualization of $F$ embedding vectors for different $z$ vectors, after projecting them in two-dimensional space with t-SNE.

# List of acronyms and abbreviations 

RL reinforcement learning
e.g. exempli gratia [for instance]

MDP Markov decision process
TD temporal difference
GTD gradient temporal difference
MSPBE mean squared projected Bellman error
OFU optimism in face to the uncertainty
DQN deep Q-network
TRPO trust region policy optimization
PPO proximal policy optimization
CPI conservative Policy Iteration
KL Kullback-Leibler
FB forward-backward representations

## Acknowledgements

My PhD was both a humanly and scientifically rich journey. I am incredibly fortunate to have done my PhD at Mila, certainly one of the most thriving environments in the field. I feel greatly grateful to be surrounded by an uncountable number of beautiful humans and brilliant minds. The abovementioned acknowledgments are surely brief and incomplete and won't reflect perfectly my deep gratitude.

To my mother and my father, Najoua and Mohammed, to whom this thesis is dedicated, for their unconditional love, prayers and for raising me to value education. To my sister and brother, Yosra and Khalil, for their support and eternal companionship despite the geographical distance between us.

To Pascal Vincent, for believing in me and for giving me the opportunity to spend these last 5 years at Mila. Being advised by him is a privilege. His scientific rigor, honesty and openness helped me take shape as a researcher. Besides being a remarkable academic mentor, he provided me a big emotional support. He always managed to lift my spirit and helped me to go through my tough times. None of this thesis would have existed without him.

To Doina Precup, for being a source of inspiration. Her introductory course of RL is the reason why I decided to do my PhD in this field. I felt honored when she invited me to attend my first conference RLDM 2017 to discover the active research topics in RL. She is a wonderful woman.

To Pierre Luc Bacon, for all his advice and helpful technical discussions, especially at the early stage of my PhD. He also helped me to know better the RL community by introducing me to many prominent researchers during conferences.

To Joelle Pineau, for hosting me as part-time researcher within her Facebook Research team in Montreal, for all the interesting discussions and thoughts she shared with me. Her role as charismatic leader and great researcher makes her a source of inspiration for many of us.

To Marc G. Bellemare, for offering me the opportunity to do an internship at Google Brain in Montreal under his supervision, for his clarity of thoughts and for always challenging my views.

To Yann Ollivier, for offering me an internship at Facebook research in Paris under his supervision and whom I learnt a disproportionate amount from in the short time. He is remarkable researcher and was fully committed to our project. It was fantastic to collaborate with him.

To Michal Valko, for being a good friend, for all his advice and guidance, for our technical discussions and for helping me to overcome my insecurities.

To Alessandro Lazaric, for giving me a two-hour crash course on online learning in RL, for being my mentor during my Facebook internship and for his willingness to answer my technical questions.

To Matteo Pirotta, for the stimulating discussions and valuable feedback.
To all my collaborators at Mila, for the exciting work we have done together, Chin-Wei Huang, Amy Zhang, Joshua Romoff and Harsh Satija,... to name few of them. Special thanks to Harsh for providing me feedback on almost all my papers and for always showing up at my posters during online events so I don't feel alone in the zoom room.

To my big Montreal family, who made my Montreal life more colorful and warmer. To Max, Pier, Ben and Brett, for our beach volley ball sessions, for all the the craziness and laughter we had together. To Léo, who has been my roommate for almost 4 years, for her support and for making work from home enjoyable. To my other roommates Sylvain and Anna, for all the moments we shared together in our home. To Rim, for being like a sister from the first sight. To Gabriel, Rémi, Gauthier, Hugo, Tom, Adrien, Valentin, Salem, David, AnneMarie, Victor, Angèle, Lucile, Benoît, Tess, Elodie, for all the experiences and memories we created together. Finally, to Clement for always helping me to go through the toughest moments of my PhD with his enthusiasm and creativity.

To my Friends from Tunisia and France, Omar, Achref, Nouha, Frikha, Tarak, Dali, Lamloum, Qays, Kaabar, Amal, the two Nicolas, Florent, Fadi, Simon, Matthieu, Solomane, Maxime, Rémi, Benjamin whom I keep a strong tie with despite the geographical distance, for their constant support. I'm thankful to be surrounded by such wonderful friends.

## Notation

The set of real numbers ..... $\mathbb{R}$
The set $\{1,2, \ldots, N\}$ for an integer $N>0$ ..... [ $N$ ]
The set of states ..... $\mathcal{S}$
The set of actions ..... $\mathcal{A}$
The discount factor ..... $\gamma$
The reward upon taking the action $a$ in state $s$ ..... $r(s, a)$
Transitioning probability into $s^{\prime}$ conditioned on $(s, a) \ldots$ ..... $P\left(d s^{\prime} \mid s, a\right)$
Planning horizon ..... H
The reward upon taking the action $a$ in state $s$ at stage $h$ ..... $r_{h}(s, a)$
Transitioning probability into $s^{\prime}$ conditioned on $(s, a)$ at stage $h$ ..... $P_{h}\left(d s^{\prime} \mid s, a\right)$
Policy$\pi$
The value function and action-value function of policy $\pi$ ..... $V^{\pi}, Q^{\pi}$
The advantage function for policy $\pi$ ..... $A^{\pi}$
Optimal value function and action-value function ..... $V^{\star}, Q^{\star}$
One-step transition operator for policy $\pi$ ..... $P^{\pi}$
Bellman evaluation operator for policy $\pi$ ..... $\mathcal{T}^{\pi}$
Bellman optimality operator ..... $\mathcal{T}$
The discounted state visitation for $\pi$ and initial distribution $\rho$$d_{\rho}^{\pi}$
Successor states measureThe Euclidien norm of vector $x$$\|x\|$
The supremum norm of vector $x$$\|x\|_{\infty}$
The Euclidien norm of vector $x$ weighted by the matrix $V$. ..... $\|x\|_{V}$
The spectral norm of matrix $V$The Kullback-Leibler divergence$\|V\|$
The Frobenius norm of the matrix $V$ ..... $\|V\|_{F}$
The total variation distance ..... $D_{\text {TV }}$$D_{\mathrm{KL}}$
Indicator function that returns 1 if $a=b$ and 0 otherwise

$$
\mathbb{1}_{\{a=b\}}
$$

## Introduction

## «Intelligence is the computational part of the ability to achieve goals in the world» -John McCarthy

Reinforcement Learning (RL) is a computational approach to solving sequential decision-making problems. Through trial and error an agent must learn to act optimally in an unknown environment in order to maximize its expected utility. A standard setting where a RL agent evolves is illustrated in Figure 1.1: the agent observes the environment's state and a reward associated with the last state transition. It chooses an action. The environment makes a transition to a new state and sends it back to the agent, and the process is repeated. The aim of the agent is to learn to control the environment so as to maximize the longterm total reward. What makes RL different from supervised learning is that the agent receives only partial feedback about his executed actions. Moreover, actions may have a long term effects. For example, a RL agent may take many decisions over multiple steps where immediate rewards are (almost) zero and where more relevant events are rather distant in the future. Such delayed signal makes it challenging to assign credit to temporally distant states and actions that have an effect on the rewards. Finally, the agent's actions affect the subsequent data it receives. The agent needs sometimes to sacrifice its current utility and seek actions that reveal new information about its environment. This will prevent it from prematurely exploiting early limited knowledge and falling into local optima.


Figure 1.1: Standard Reinforcement Learning setting

While RL remained on the margins of the broader artificial intelligence (AI) community until relatively recently, today, RL flourishes as a vital learning paradigm that captures many important AI applications, including self-driving cars, robots, and adaptive medical treatments. This recent success is partly due to the deployment of flexible function approximation schemes (such as the use of deep neural networks in Deep RL) to deal with large state spaces. However, a closer look at these impressive empirical results reveals some fundamental limitations:

- It has been challenging to combine desirable features of RL algorithms, such as off-policy and multi-step learning with function approximation in a way that leads to both stable and efficient algorithms.
- Deep RL algorithms tend to be very sample inefficient due to the rudimentary exploration-exploitation strategies these approaches employ.
- Furthermore, They require an enormous amount of supervised data and end up producing a narrow agent able to solve only the task that it is trained on.

Addressing the aforementioned challenges and bridging the gap between practise and theory in RL are crucial steps to bring RL to real life and fulfill its promise, and this is what inspired my research.

This thesis is structured as follows: The next section provides a synthetic overview of the research contributions that will be detailed in subsequent chapters. Readers less familiar with RL and its terminology and open challenges will benefit from first reading the background presented in Chapter 2, which attempts to be a compact formal introduction to RL and to the notions this thesis builds on. Each of the following chapters 3 to 7 corresponds to a published or submitted research paper. Chapter 8 concludes the thesis, discussing future research directions.

As this is a thesis by article, the review and discussion of the literature most relevant to each contribution is to be found, in its context, within that contribution's chapter, as in the corresponding article, rather than in a separate dedicated literature review chapter.

### 1.1 Research Contributions

### 1.1.1 Off-Policy Learning

Off-policy learning is concerned with evaluating or optimizing a target policy based on data from a different behavior policy. This can provide many benefits: efficient parallel exploration and, reuse of past experience with experience replay and, in many practical contexts, learning form data produced by policies that are currently deployed, but which we want to improve.

However, we show in «Convergent Tree Backup and Retrace with Function Approximation» [Touati et al., 2018] (Chapter 3) that combining popular off-policy and multi-step learning methods, for policy evaluation, with linear value function parameterization could lead to undesirable instability. Then, we derive a provably convergent variant of these methods by converting the problem into a primal-dual saddle point problem. Finally, we provide their finite sample analysis.

In another work «Stable Policy Optimization via Off-Policy Divergence Regularization» Touati et al. [2020c] (Chapter 4), we ask a different question:

Given a stream of experience generated by a given policy, how could we improve it in a safe and stable way?

For that, we revisit the foundations of conservative approaches to policy optimization that optimize a surrogate objective that can provide local improvements to the current policy at each iteration. We highlight that these methods rely solely on constraining immediate action probabilities and argue that the latter might not be enough and that we should rather reason about the longterm effect of the policies on the distribution of the future states. In particular, we directly consider the divergence between state-action occupancy measures induced by successive policies and use it as a regularization term added to the surrogate objective. This regularization term is itself optimized in an adversarial and off-policy manner. Finally, we empirically show that our proposed method can have a beneficial effect on stability and improve final performance in benchmark high-dimensional control tasks.

### 1.1.2 Exploration v.s. Exploitation Dilemma

An RL agent, through trial and error, must learn to act optimally in an unknown environment to maximize its expected utility. Exploration refers to the execution of diversified actions to collect data that provide a well-rounded characterization of the environment. Thus, an efficient learning requires balancing exploration (to gain more knowledge) and exploitation (acting optimally according to the available knowledge).

While a wealth of research has been developed to design systematic exploration strategies in small state-space problems (known as tabular problems), their theoretical measure guarantee typically scales with the number of discrete states and the number of discrete actions. This precludes applying them to arbitrarily large state-action spaces.

An appealing challenge is to combine exploration strategies with generalization methods in a way that leads to both provable sample efficient and computationally efficient RL algorithms for large-scale problems. This would require some structural assumptions.

A simple way to ensure generalization over states is to aggregate them into a finite set of meta-states and run a tabular exploration mechanism on the latter. In this direction, we proposed in «Zooming for Efficient Model-free Reinforcement Learning in Metric Spaces» [Touati et al., 2020b] (Chapter 5) to actively explore the state-action space by learning on the-fly an adaptive partitioning that takes into account the shape of the optimal value function. When the state-action space is assumed to be a compact metric space, such adaptive discretization based algorithms yield sublinear regret.

Another structural assumption, that received attention in the recent literature, is when both reward and transition dynamics are linear functions with respect to a given feature mapping. This assumption enables the design of efficient algorithms with a linear representation of the action value function. In our work «Efficient learning in non-stationary linear Markov decision processes» [Touati and Vincent, 2020a] (Chapter 6), we add exogenous non-stationarity where the environment is now allowed to evolve dynamically and even adversarially over time, but the total change is still bounded. This is a more challenging setting, since what has been learned in the past may be obsolete in the present. To address this setting, our proposed algorithm is based on weighted least squares value iteration that uses exponential weights to smoothly forget data that are far in the past, which drives the agent to keep exploring to discover changes. We prove that this algorithm achieves state-of-the-art online performance.

### 1.1.3 Unsupervised Learning in RL

In standard RL, the agent is given a training signal from the environment in the form of a scalar reward and aims to maximize the expected cumulative rewards along its trajectories. This strategy ends up producing a narrow agent that is only able to solve the single task that it is trained on. Furthermore, it is arguably not how humans learn in real life. For example, a child might get supervisions from his parents in various forms such as verbal instructions or applause. But most of the time, he is exploring the world through curiosity and play and gaining knowledge about his surroundings.

In our work «Learning One Representation to Optimize All Rewards» [Touati and

Ollivier, 2021] (Chapter 7), we ask this question:
Given an environment without reward information, is it possible to learn and store a compact object that, for any reward function specified later, provides the optimal policy for that reward, with a minimal amount of additional computation?

To address the question above, we propose to learn a pair of representations $F$ (for forward) and $B$ (for backward) that map state-action pairs to some latent space. During the unsupervised learning phase, these representations are trained using reward-free interactions such that $F^{\top} B$ encodes long-range relationships between states and actions, via a predictive occupancy map. At test time, once a reward function is revealed, we show that the optimal policy for that reward is directly obtained from these representations, with no planning. Empirically, our approach compares well to goal-oriented RL algorithms on discrete and continuous mazes, pixel-based MsPacman (an Atari game), and a virtual robot arm. We also illustrate how the agent can immediately adapt to new tasks beyond goal-oriented RL.

### 1.2 LIST OF ExCLUDED CONTRIBUTIONS

During my PhD, I produced other contributions on various machine learning topics, that are not included in this thesis.

- Reinforcement Learning:
- Joshua Romoff, Peter Henderson, David Kanaa, Emmanuel Bengio, Ahmed Touati, Pierre-Luc Bacon, Joelle Pineau. TDprop: Does Jacobi Preconditioning Help Temporal Difference Learning?. AAMAS 2021. [Romoff et al., 2021]
- Ahmed Touati, Pascal Vincent. Sharp Analysis of Smoothed Bellman Error Embedding. ICML 2020 Workshop on Theoretical Foundations of Reinforcement Learning. [Touati and Vincent, 2020b]
- Zilun Peng*, Ahmed Touati*, Pascal Vincent and Doina Precup. SVRG for Policy Evaluation with Fewer Gradient Evaluations. IJCAI 2020. [Peng et al., 2019]
- Joshua Romoff, Peter Henderson, Ahmed Touati, Yann Ollivier, Emma Brunskill, Joelle Pineau. Separating value functions across time-scales. ICML 2019. [Romoff et al., 2019]
- Nan Rosemary Ke, Amanpreet Singh, Ahmed Touati, Anirudh Goyal, Yoshua Bengio, Devi Parikh and Dhruv Batra. Learning Dynamics Model in Reinforcement Learning by Incorporating the Long Term Future. ICLR 2019. [Ke et al., 2019]
- Ahmed Touati, Harsh Satija, Joshua Romoff, Joelle Pineau, Pascal Vincent. Randomized Value Functions via Multiplicative Normalizing Flows. UAI 2019. [Touati et al., 2020a]
- Bayesian Deep Learning:
- Chin-Wei Huang, Ahmed Touati, Pascal Vincent, Gintare Karolina Dziugaite, Alexandre Lacoste, Aaron Courville. Stochastic Neural Network with Kronecker Flow. AISTATS 2020. [Huang et al., 2020]
- Chin-Wei Huang, Ahmed Touati, Laurent Dinh, Michal Drozdzal, Mohammad Havaei, Laurent Charlin, Aaron Courville. Learnable Explicit Density for Continuous Latent Space and Variational Inference. ICML 2017 workshop. [Huang et al., 2017a]


## - Generative Modeling:

- Gabriel Huang, Hugo Berard, Ahmed Touati, Gauthier Gidel, Simon Julien-Lacoste. Adversarial Divergences are Good Task Losses for Generative Modeling. Under review at Journal of Machine Learning Research. [Huang et al., 2017b]


## - Convex Optimization:

- Rémi Le Priol, Ahmed Touati and Simon Lacoste-Julien. Adaptive Stochastic Dual Coordinate Ascent for Conditional Random Fields. NIPS 2017 Workshop on Optimization for Machine Learning.


## 2 <br> Background

### 2.1 Discounted Markov Decision Process

### 2.1.1 The Model

Markov Decision Processes [Puterman, 1994] are a mathematical tool for modeling sequential decision-making problems where a decision maker or an agent interacts with a system in a sequential fashion. A discounted Markov Decision Process (MDP) is a tuple $(\mathcal{S}, \mathcal{A}, P, r, \gamma)$ with state space $\mathcal{S}$, action space $\mathcal{A}$, transition probabilities $P$ mapping state-action pairs to distributions over next states, immediate reward function $r$, and a discount fator $\gamma \in[0,1)$. The reward could be stochastic. If $\mathcal{S}$ is discrete, for each $(s, a), P\left(s^{\prime} \mid s, a\right)$ is a probability mass function on $s^{\prime} \in \mathcal{S}$. For general $\mathcal{S}, P\left(d s^{\prime} \mid s, a\right)$ is a probability measure on $s^{\prime} \in \mathcal{S}$.

For the ease of exposition in this chapter, we restrict our attention to finite MDPs ( $\mathcal{S}$ and $\mathcal{A}$ are finite). We also assume that the reward function is bounded, $r: \mathcal{S} \times \mathcal{A} \rightarrow\left[0, r_{\text {max }}\right]$ for some positive scalar $r_{\text {max }}$.

An agent can select its actions at any stage based on the observed history. A rule describing the way the actions are selected is called a policy. Starting from an initial state $s_{0}$ and executing the policy, the agent generates a random sequence of state-action-reward $\left\{s_{0}, a_{0}, r_{0}, s_{k+1}, \ldots s_{t}, a_{t}, r_{t}, s_{t+1} \ldots\right\}$ where $r_{t}=$ $r\left(s_{t}, a_{t}\right)$ and $s_{t+1} \sim p\left(\cdot \mid s_{t}, a_{t}\right)$. The return underlying a behavior is defined as the total discounted sum of the rewards incurred:

$$
\begin{equation*}
\mathcal{R}=\sum_{t=0}^{\infty} \gamma^{t} r_{t} \tag{2.1.1}
\end{equation*}
$$

The goal of the agent is to choose a behavior that maximizes the expected return. Such a maximizing policy is said to be optimal. It is known that an optimal policy can be memoryless in the sense that it does not need the whole observed history and may only depend on the current state. Moreover, it is known that to be optimal over an infinite horizon setting, a policy can be stationary in the sense that it does not depend on the time that a state was encountered.

### 2.1.2 Policies and Value Functions

As stated above, we will restrict our attention to memoryless and stationary policies. Formally, we define a policy as a mapping from states to probabilities of selecting each possible action. We denote by $\pi(a \mid s)$ the probability of choosing action $a$ in state $s$ under the policy $\pi: \mathcal{S} \rightarrow \operatorname{Prob}(\mathcal{A})$, where $\operatorname{Prob}(\mathcal{A})$ denotes the space of probability distribution over actions.

An action-value function, know also as $Q$-function, is a function of stateaction pairs that estimate how good it is for the agent to be in a given state-action pair. Thus, it is a way to evaluate policy and compare different policies. The value of a state-action pair $(s, a)$ under a policy $\pi$, denoted $Q^{\pi}(s)$, is the expected return when starting in state $s$, taking action $a$ and following $\pi$ thereafter:

$$
\begin{equation*}
Q^{\pi}(s, a) \triangleq \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^{t} r_{t} \mid\left(s_{0}, a_{0}\right)=(s, a), \pi\right] \tag{2.1.2}
\end{equation*}
$$

Where the expectation is over trajectories drawn by executing $\pi$ : $\left\{\left(s_{t}, a_{t}, r_{t}, s_{t+1}\right)\right\}_{t}$ where $a_{t} \sim \pi\left(. \mid s_{t}\right), \quad r_{t}=r\left(s_{t}, a_{t}\right), \quad$ and $\quad s_{t+1} \sim p\left(\cdot \mid s_{t}, a_{t}\right)$. Similarly, the value of a state $s$ under a policy $\pi$, denoted $V^{\pi}(s)$, is the expected return when starting in state $s$ and following $\pi$ thereafter:

$$
\begin{equation*}
V^{\pi}(s) \triangleq \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^{t} r_{t} \mid s_{0}=s, \pi\right] \tag{2.1.3}
\end{equation*}
$$

Value function and action-value function are related as follows:

$$
\begin{equation*}
V^{\pi}(s)=\mathbb{E}_{a \in \pi(\cdot \mid s)}\left[Q^{\pi}(s, a)\right]=\sum_{a \in \mathcal{A}} \pi(a \mid s) Q^{\pi}(s, a) \tag{2.1.4}
\end{equation*}
$$

A fundamental property of action-value functions used throughout reinforcement learning is that they satisfy the following recursive relationship:

$$
\begin{align*}
Q^{\pi}(s, a) & =\mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^{t} r_{t} \mid\left(s_{0}, a_{0}\right)=(s, a), \pi\right] \\
& =r(s, a)+\mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^{t+1} r_{t+1} \mid\left(s_{0}, a_{0}\right)=(s, a), \pi\right] \\
& =r(s, a)+\gamma \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right) \mathbb{E} \underbrace{\left[\sum_{t=0}^{\infty} \gamma^{t} r_{t} \mid s_{0}=s^{\prime}, \pi\right]}_{V^{\pi}\left(s^{\prime}\right)} \\
& =r(s, a)+\gamma \sum_{s^{\prime}, a^{\prime}} P\left(s^{\prime} \mid s, a\right) \pi\left(a^{\prime} \mid s^{\prime}\right) Q^{\pi}\left(s^{\prime}, a^{\prime}\right) \tag{2.1.5}
\end{align*}
$$

This is the Bellman equation for $Q^{\pi}$. It expresses a relationship between the value of a state-action pair and the values of its successor states and actions. Similarly, the Bellman equation for $V^{\pi}$ is defined as follows:

$$
\begin{equation*}
V^{\pi}(s)=\sum_{a} \pi(a \mid s) r(s, a)+\gamma \sum_{s^{\prime}, a} \pi(a \mid s) P\left(s^{\prime} \mid s, a\right) V^{\pi}\left(s^{\prime}\right) \tag{2.1.6}
\end{equation*}
$$

### 2.1.3 Optimal Policies and Optimal Value Functions

The optimal value functions, denoted by $V^{\star}$ and $Q^{\star}$ respectively, are defined by:

$$
\begin{align*}
V^{\star}(s) & \triangleq \max _{\pi} V^{\pi}(s)  \tag{2.1.7}\\
Q^{\star}(s, a) & \triangleq \max _{\pi} Q^{\pi}(s, a) \tag{2.1.8}
\end{align*}
$$

A remarkable property of MDPs is that there exists a stationary and deterministic policy $\pi$ that simultaneously maximizes $V^{\pi}(s)$ for all $s \in \mathcal{S}$ [Puterman, 1994]. We refer to such a $\pi$ as an optimal policy. The optimal value and action-value functions are connected by the following equations:

$$
\begin{align*}
V^{\star}(s) & =\max _{a \in \mathcal{A}} Q^{\star}(s, a)  \tag{2.1.9}\\
Q^{\star}(s, a) & =r(s, a)+\gamma \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right) V^{\star}\left(s^{\prime}\right) \tag{2.1.10}
\end{align*}
$$

The latter equation is known as the optimal Bellman equation. We know that for all policies $\pi, V^{\pi}(s)=\sum_{a} \pi(a \mid s) Q^{\pi}(s, a) \quad \forall s \in \mathcal{S}$. In order latter equation holds for optimal value function $V^{\star}, \pi^{\star}(a \mid s)>0$ only if $Q^{\star}(s, a)=\max _{a^{\prime} \in \mathcal{A}} Q^{\star}\left(s, a^{\prime}\right)$, which simply corresponds to choosing the action of maximum value in each state. Therefore, the greedy policy with respect to $Q^{\star}$ is an optimal policy.

### 2.2 Episodic Markov Decision Processes

We presented in the last section the infinite-horizon discounted setting of MDPs where the return of a trajectory is given by $\sum_{t=0}^{\infty} \gamma^{t} r_{t}$. Popular alternative choices include the finite-horizon undiscounted setting (return of a trajectory is $\sum_{t=1}^{T} r_{t}$ with some finite horizon $H<\infty$ ) and the infinite-horizon average reward setting (return is $\lim _{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^{T} r_{t}$ ). The latter case often requires additional conditions on the transition dynamics (such as ergodicity) so that values can be well-defined.

A finite horizon MDP is a tuple $(\mathcal{S}, \mathcal{A}, P, r, H)$ where $\mathcal{S}$ and $\mathcal{A}$ are the state and action space, $H$ is the planning horizon i.e number of steps in each episode,
$P=\left\{P_{h}\right\}$ is a collection of the transition probabilities such that $P_{h}(\cdot \mid s, a)$ gives the distribution over next states if action $a$ is taken at state $s$ at step $h \in[H]$, $r=\left\{r_{h}\right\}$ is a collection of reward functions such that $r_{h}(s, a)$ is the reward of taking action $a$ at state $s$ at time step $h$.

For any step $h \in[H]$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$, the action-value function of a non-stationary policy $\pi=\left(\pi_{1}, \ldots, \pi_{H}\right)$ is defined as $Q_{h}^{\pi}(s, a)=r_{h}(s, a)+$ $\left.\mathbb{E}\left[\sum_{i=h+1}^{H} r_{i}\left(s_{i}, a_{i}\right)\right) \mid\left(s_{h}, a_{h}\right)=(s, a), \pi\right]$, and the value function is $V_{h}^{\pi}(s)=$ $\sum_{a} \pi_{h}(a \mid s) Q_{h}^{\pi}(s, a)$. As the horizon is finite, there always exists an optimal policy $\pi^{\star}$ whose value and action-value functions are defined as $V_{h}^{\star}(x)=V_{h}^{\pi^{\star}}(s)=$ $\max _{\pi} V_{h}^{\pi}(s)$ and $Q_{h}^{\star}(s, a)=Q_{h}^{\pi^{\star}}(s, a)=\max _{\pi} Q_{h}^{\pi}(s, a)$. Both $Q^{\pi}$ and $Q^{\star}$ can be conveniently written as the result of the following Bellman equations

$$
\begin{align*}
& Q_{h}^{\pi}(s, a)=r_{h}(s, a)+\sum_{s^{\prime}} P_{h}\left(s^{\prime} \mid s, a\right) V_{h+1}^{\pi}\left(s^{\prime}\right),  \tag{2.2.1}\\
& Q_{h}^{\star}(s, a)=r_{h}(s, a)+\sum_{s^{\prime}} P_{h}\left(s^{\prime} \mid s, a\right) V_{h+1}^{\star}\left(s^{\prime}\right), \tag{2.2.2}
\end{align*}
$$

where $V_{H+1}^{\pi}(s)=V_{H+1}^{\star}(s)=0$ and $V_{h}^{\star}(s)=\max _{a \in \mathcal{A}} Q_{h}^{\star}(s, a)$, for all $s \in \mathcal{S}$.

### 2.3 Dynamic Programming

«An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision» -Bellman's Principle of Optimality [Bellman, 1957]

The term dynamic programming (DP) refers to a collection of algorithms that can be used to compute optimal policies given a perfect model of the environment as a Markov decision process (MDP). DP leverages the structure of MDP to simplify the problem solving. While DP has limited practical importance as it assumes access to a perfect model, it remains of great theoretical interest and gives valuable insights for the design of RL algorithms.

### 2.3.1 Bellman Operators

Many DP algorithms operate in the value function space by repeatedly applying transformations. Therefore, it is convenient to think of the latter as operators. If we consider action-value function as mapping : $\mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. An operator $\mathcal{B}$ over action-value functions then takes a function $Q: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ and maps it to a function over the same domain, $\mathcal{B Q}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$.

Let us first define the transition operator $\mathcal{P}^{\pi}$ for a policy $\pi$ :

$$
\begin{equation*}
\mathcal{P}^{\pi} Q(s, a) \triangleq \sum_{s^{\prime}, a^{\prime}} P\left(s^{\prime} \mid s, a\right) \pi\left(a^{\prime} \mid s^{\prime}\right) Q\left(s^{\prime}, a^{\prime}\right) \tag{2.3.1}
\end{equation*}
$$

The spectral radius of $\mathcal{P}^{\pi}$, as stochastic matrix, is equal to one. This implies that the Neuman expansion $\sum_{t=0}^{\infty} \gamma^{t}\left(\mathcal{P}^{\pi}\right)^{t}$ is convergent and is equal to $\left(I-\gamma \mathcal{P}^{\pi}\right)^{-1}$. Therefore, we can write $Q^{\pi}$ as

$$
\begin{equation*}
Q^{\pi}=\sum_{t=0}^{\infty} \gamma^{t}\left(\mathcal{P}^{\pi}\right)^{t} r=\left(I-\gamma P^{\pi}\right)^{-1} r \tag{2.3.2}
\end{equation*}
$$

Now, we can write The Bellman evaluation operator over Q-functions, which corresponds to applying (2.1.5) to all state-action pairs, as for any $Q: \mathcal{S} \times \mathcal{A} \rightarrow$ $\mathbb{R}$,

$$
\begin{equation*}
\mathcal{T}^{\pi} Q \triangleq r+\gamma \mathcal{P}^{\pi} Q \tag{2.3.3}
\end{equation*}
$$

Similarly to this evaluation operator, we can define the Bellman optimality operator as, for any $Q: \mathcal{S} \times \mathcal{A} \times \mathbb{R}$,

$$
\begin{equation*}
\mathcal{T} Q \triangleq r+\max _{\pi} \mathcal{P}^{\pi} Q \tag{2.3.4}
\end{equation*}
$$

Both operators share nice properties that provide the basis for DP algorithms such as Value Iteration and Policy Iteration. First, the $Q$-function of $Q^{\pi}$ of the policy $\pi$ is a fixed point of $\mathcal{P}^{\pi}$ while the optimal $Q$-function $Q^{\star}$ is a fixed point of $\mathcal{T}$ i.e $\mathcal{T}^{\pi} Q^{\pi}=Q^{\pi}$ and $\mathcal{T} Q^{\star}=Q^{\star}$. Moreover, Both operators are $\gamma$-contraction in supremum norm $l_{\infty}$ as for any pair of functions $Q_{1}, Q_{2}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, we have,

$$
\begin{align*}
\left\|\mathcal{T}^{\pi} Q_{1}-\mathcal{T}^{\pi} Q_{2}\right\|_{\infty} & \leq \gamma\left\|Q_{1}-Q_{2}\right\|_{\infty}  \tag{2.3.5}\\
\left\|\mathcal{T} Q_{1}-\mathcal{T} Q_{2}\right\|_{\infty} & \leq \gamma\left\|Q_{1}-Q_{2}\right\|_{\infty} \tag{2.3.6}
\end{align*}
$$

Finally, both operators are monotone in the sense that for any pair of functions $Q_{1}$ and $Q_{2}$ such that $Q_{1} \leq Q_{2}$ component-wise, we have $\mathcal{T}^{\pi} Q_{1} \leq \mathcal{T}^{\pi} Q_{2}$ and $\mathcal{T} Q_{1} \leq \mathcal{T} Q_{2}$.

### 2.3.2 Value Iteration

The Value Iteration (VI) [Bellman, 1957] algorithm operates on the Q-function space. Starting with an arbitrary, initial $Q$-function estimate, it repeatedly applies the Bellman optimality operator $\mathcal{T}$ so that the $Q$-function estimate approaches the optimal $Q$-function in the limit. A basic form of VI is given in algorithm 1. The behaviour of VI can be analyzed by the Banach fixed point

```
Algorithm 1 Value Iteration Algorithm
Input: \(\operatorname{MDP}(\mathcal{S}, \mathcal{A}, P, r, \gamma)\)
    Initialize a value \(Q_{0}\)
    for \(\mathrm{k}=1 \ldots \mathrm{~K}\) do
        \(Q_{k+1}=\mathcal{T} Q_{k}\)
    end for
    Return the greedy policy \(\pi_{K}(s)=\arg \max _{a} Q_{K}(s, a)\)
```

theorem that states, if $\mathcal{X}$ is a complete metric space ${ }^{1}$ with a contraction mapping $\mathcal{B}: \mathcal{X} \rightarrow \mathcal{X}$, then $\mathcal{B}$ admits a unique fixed point and the sequence $x, \mathcal{B} x, \mathcal{B}^{2} x, \ldots$ converge to that fixed point for any $x \in \mathcal{X}$. In particular, it shows that the optimal $Q$-function $Q^{\star}$ is indeed the unique fixed point of $\mathcal{T}$ and repeated applications of $\mathcal{T}$ are guaranteed to yield it. Furthermore, it can be shown that VI converges at a geometric rate. In fact, we have for any initial function $Q$,

$$
\begin{aligned}
\left\|\mathcal{T}^{k} Q-Q^{\star}\right\|_{\infty} & =\left\|\mathcal{T}\left(\mathcal{T}^{k-1} Q\right)-\mathcal{T} Q^{\star}\right\|_{\infty} \\
& \leq \gamma\left\|\mathcal{T}^{k-1} Q-Q^{\star}\right\|_{\infty} \\
& \leq \ldots \leq \gamma^{k}\left\|Q-Q^{\star}\right\|_{\infty}
\end{aligned}
$$

### 2.3.3 Policy Iteration

Another classic DP approach to solve MDPs is Policy Iteration (PI). This algorithm 2 directly searches the optimal policy by interleaving between a policy evaluation step and a policy improvement step. Starting with an arbitrary policy, the algorithm proceeds in an iterative way. At iteration $k>0$, it computes the action-value function $Q^{\pi_{k}}$ of $\pi_{k}$. Next, it improves the policy by updating the policy to be the greedy policy with respect to $Q^{\pi_{k}}$ i.e $\pi_{k+1}(s)=\arg \max _{a} Q^{\pi_{k}}(s, a)$ for all $s \in \mathcal{S}$. In terms of the Bellman operators, this can be equivalently expressed as $\mathcal{T}^{\pi_{k+1}} Q^{\pi_{k}}=\mathcal{T} Q^{\pi_{k}}$. We have that

$$
\begin{equation*}
Q^{\pi_{k}}=\mathcal{T}^{\pi_{k}} Q^{\pi_{k}} \leq \mathcal{T} Q^{\pi_{k}}=\mathcal{T}^{\pi_{k+1}} Q^{\pi_{k}} \text { (component-wise). } \tag{2.3.7}
\end{equation*}
$$

By applying the operator $\mathcal{T}^{\pi_{k+1}}$ on both sides and using the monotonicity property, we obtain for any integer $n>0$

$$
\begin{equation*}
Q^{\pi_{k}} \leq \mathcal{T}^{\pi_{k+1}} Q^{\pi_{k}} \leq \ldots \leq\left(\mathcal{T}^{\pi_{k+1}}\right)^{n} Q^{\pi_{k}}(\text { component-wise }) \tag{2.3.8}
\end{equation*}
$$

[^0]```
Algorithm 2 Policy Iteration Algorithm
Input: \(\operatorname{MDP}(\mathcal{S}, \mathcal{A}, P, r, \gamma)\)
    Initialize a policy \(\pi_{1}\)
    for \(k=1 \ldots K\) do
        given \(\pi_{k}\), compute \(Q^{\pi_{k}} / *\) Policy Evaluation
        \(\pi_{k+1}(s)=\arg \max _{a} Q^{\pi_{k}}(s, a) / *\) Policy Improvement
    end for
    Return the last policy \(\pi_{K}\)
```

By applying the Banach fixed-point theorem on the contraction $\mathcal{T}^{\pi_{k+1}}$, we obtain $\lim _{n \rightarrow}\left(\mathcal{T}^{\pi_{k+1}}\right)^{n} Q^{\pi_{k}}=Q^{\pi_{k+1}}$. Which implies that $Q^{\pi_{k}} \leq Q^{\pi_{k+1}}$. Therefore, PI generates a sequences of policies with non-decreasing performance. Since we have now $\mathcal{T} Q^{\pi_{k}}=\mathcal{T}^{\pi_{k+1}} Q^{\pi_{k}} \leq Q^{\pi_{k+1}} \leq Q^{\star}$, we obtain

$$
\begin{aligned}
\left\|Q^{\pi_{k+1}}-Q^{\star}\right\|_{\infty} & \leq\left\|\mathcal{T}^{\pi_{k+1}} Q^{\pi_{k}}-Q^{\star}\right\|_{\infty} \\
& =\left\|\mathcal{T} Q^{\pi_{k}}-\mathcal{T} Q^{\star}\right\|_{\infty} \\
& \leq \gamma\left\|Q^{\pi_{k}}-Q^{\star}\right\|_{\infty} \\
& \leq \ldots \leq \gamma^{k+1}\left\|Q^{\pi_{1}}-Q^{\star}\right\|_{\infty}
\end{aligned}
$$

Thus, PI converge at a linear rate.

### 2.4 TEMPORAL DIFFERENCE LEARNING

«TD learning is learning a prediction from another, later, learned prediction, i.e, learning a guess from a guess» -Richard Sutton at Montreal summer school, 2017.

Dynamic programming algorithms assume that both the transition probabilities $P$ and the reward function $r$ are explicitly known. On the other hand, in the general reinforcement learning setting, this information is not available to the agent which needs a direct interaction with the environment in order to solve it.

For value function estimation at a given state, a straightforward alternative is to approximate the expectation of the random return by an average over independent trajectories started from the given state. This is an instance to the so-called Monte-Carlo method. However, this technique faces several issues: the estimator tends to have high variance, it is computed only when all trajectories are terminated, and it assumes that we are able to reset the environment to some particular state.

To address these shortcomings, temporal difference (TD) learning [Sutton, 1988] has been proposed and has become the main building block of many RL algorithms. TD learning approaches combine ideas from both Monte Carlo and DP. Similarly to Monte Carlo, they learn from raw experiences without access to an explicit model of the environment. Similarly to DP, they update estimates of the values of states based on estimates of the values of successor states. This latter is referred as bootsrapping.

### 2.4.1 Policy Evaluation

For the problem of policy evaluation, TD learning, in its simplest form TD(0), updates the value of a state-action pair $\left(s_{t}, a_{t}\right)$ towards an estimate of the return $r_{t}+\gamma Q\left(s_{t+1}, a_{t+1}\right)$, called the TD target, where $r_{t}$ is the reward given by the environment at time $t, s_{t+1}$ is the state reached after the transition, and $a_{t+1}$ is the action selected by the policy being evaluated. More precisely, TD(0) update rule is as follows :

$$
\begin{equation*}
Q\left(s_{t}, a_{t}\right) \leftarrow Q\left(s_{t}, a_{t}\right)+\alpha_{t}\left(s_{t}, a_{t}\right)\left(r_{t}+\gamma Q\left(s_{t+1}, a_{t+1}\right)-Q\left(s_{t}, a_{t}\right)\right) \tag{2.4.1}
\end{equation*}
$$

where $\alpha_{t}(s, a) \in(0,1)$ is the learning rate associated to the state-action pair $(s, a)$ at time $t$. The term $r_{t}+\gamma Q\left(s_{t+1}, a_{t+1}\right)-Q\left(s_{t}, a_{t}\right)$ is called the TD error, and is usually denoted by $\delta_{t}$. We have in expectation:

$$
\begin{equation*}
\mathbb{E}\left[\delta_{t} \mid\left(s_{t}, a_{t}\right)=(s, a)\right]=\left(\mathcal{T}^{\pi} Q-Q\right)(s, a), \forall(s, a) \tag{2.4.2}
\end{equation*}
$$

Therefore, $\mathrm{TD}(0)$ can be interpreted as a stochastic approximation scheme for solving a fixed point of the Bellman equation. Contraction property together with general results from stochastic approximation theory can then be used to show asymptotic convergence under the minimal assumption that for every state-action pair $(s, a)$, the sequence of step-size $\left\{\alpha_{t}(s, a)\right\}_{t}$ satisfies RobbinsMonro conditions i.e $\sum_{t} \alpha_{t}(s, a)=\infty$ and $\sum_{t} \alpha_{t}(s, a)^{2}<\infty$ [Jaakkola et al., 1994]. Since the learning rate $\alpha_{t}$ is less than 1 , the first condition $\sum_{t} \alpha_{t}(s, a)=\infty$ implies that every state and every action should be visited infinitely often.

While the Monte Carlo method has no bias but could suffer from large variance, $\mathrm{TD}(0)$, a one-step TD, has a small variance but may introduce bias. We can extend both methods by n-step TD which can shift from one to the other smoothly as needed.

The target for an arbitrary n-step update is the n-step return:

$$
\begin{equation*}
R_{t: t+n}=r_{t}+\gamma r_{t+1}+\cdots+\gamma^{n-1} r_{t+n-1}+\gamma^{n} Q\left(s_{t+n}, a_{t+n}\right) \tag{2.4.3}
\end{equation*}
$$

All n-step returns can be considered approximations to the full return, truncated after n steps and then corrected for the remaining missing terms by $Q\left(s_{t+n}, a_{t+n}\right)$.

Therefore, n -step methods enable bootstrapping to occur over multiple steps. Now we note that a valid update can be done not only toward any n-step return, but also toward any average of n-step returns. A particular and popular way of averaging n -step updates is to weight each n -step updates by $\lambda^{n-1}$, where $\lambda \in[0,1]$ and then normalize it by a factor of $1-\lambda$ to ensure that the weights sum to 1 . This particular average is called the $\lambda$-return and is defined as

$$
\begin{equation*}
R_{t}^{\lambda} \triangleq(1-\lambda) \sum_{n=0}^{\infty} \lambda^{n-1} R_{t: t+n} \tag{2.4.4}
\end{equation*}
$$

$\operatorname{TD}(\lambda)$ algorithm updates the value function estimates towards this $\lambda$-return as follows:

$$
\begin{equation*}
Q\left(s_{t}, a_{t}\right) \leftarrow Q\left(s_{t}, a_{t}\right)+\alpha_{t}\left(s_{t}, a_{t}\right)\left(R_{t}^{\lambda}-Q\left(s_{t}, a_{t}\right)\right) \tag{2.4.5}
\end{equation*}
$$

$\lambda=0$ corresponds to the $\mathrm{TD}(0)$ (one-step TD) update, and $\lambda \rightarrow 1$ removes the recursion on Q-function estimate, and recover the Monte Carlo estimate. $\lambda$ controls clearly the amount of sampling versus bootstrapping. Therefore, it is a trade-off parameter between bias and variance. We have that $\mathbb{E}\left[R_{t}^{\lambda}\right]=\mathcal{T}_{\lambda}^{\pi} V$ where $\mathcal{T}_{\lambda}^{\pi}$ is the Bellman $\lambda$-operator defined as

$$
\begin{equation*}
\mathcal{T}_{\lambda}^{\pi} \triangleq(1-\lambda) \sum_{n=0}^{\infty} \lambda^{n}\left(\mathcal{T}^{\pi}\right)^{n+1} \tag{2.4.6}
\end{equation*}
$$

$Q^{\pi}$ remains the fixed point of $\mathcal{T}_{\lambda}^{\pi}$. Therefore, $\operatorname{TD}(\lambda)$ could be interpreted as a stochastic approximation scheme for finding the fixed point of the operator $\mathcal{T}_{\lambda}^{\pi}$.

### 2.4.2 Policy Learning

So far, we dealt with TD learning for the problem of approximating the value function of a fixed policy. Now, we present the Q-learning algorithm [Watkins and Dayan, 1992], one natural extension of TD learning to control problems, where the goal is to learn an effective policy from data. Given a transition $\left(s_{t}, a_{t}, r_{t}, s_{t+1}\right)$ at time $t$, Q-learning's update rule is as follows:

$$
\begin{equation*}
Q\left(s_{t}, a_{t}\right) \leftarrow Q\left(s_{t}, a_{t}\right)+\alpha_{t}\left(s_{t}, a_{t}\right)\left(r_{t}+\max _{a} Q\left(s_{t+1}, a\right)-Q\left(s_{t}, a_{t}\right)\right) \tag{2.4.7}
\end{equation*}
$$

If we consider the TD error $\delta_{t}=r_{t}+\max _{a} Q\left(s_{t+1}, a\right)-Q\left(s_{t}, a_{t}\right)$, we have in expectation:

$$
\begin{equation*}
\mathbb{E}\left[\delta_{t} \mid s_{t}=s, a_{t}=a\right]=(\mathcal{T} Q-Q)(s, a), \forall(s, a) \tag{2.4.8}
\end{equation*}
$$

Q-learning is then a stochastic approximation scheme for finding the fixed point of the optimality Bellman operator $\mathcal{T}$. The sequence of estimates is showed to converge to $Q^{\star}$ with probability one under the assumption that for every
state-action pair $(s, a)$ the sequence of step-sizes $\left\{\alpha_{t}(s, a)\right\}_{t}$ satisfies RobbinsMonro conditions [Jaakkola et al., 1994]. This condition requires that every stateaction pair is visited infinitely.

While $\operatorname{TD}(0)$, presented earlier, requires that actions are executed by the fixed policy that we would like to evaluate, Q -learning allows us to learn from any stream of experiences. Q-learning is thus an off-policy method. The commonly used strategies in practice is to sample the actions following the $\epsilon$-greedy or the Botzmann policy with respect to the current estimate of the Q-value. While such sampling strategies ensure the full coverage of the state-action space in the limit of infinite time and eventually lead to asymptotic convergence, a more directed exploration strategy would be critical to achieve a good performance in a reasonable amount of time.

### 2.5 FUNCTION APPROXIMATION

When the state space is large, which is the case of most domains of interest, it is not feasible to keep a separate value for each state in the memory. In order to obtain generalization between different state-action pairs, value functions or policies should be represented in a compact functional form.

### 2.5.1 Value-based Methods

Value-based method learn parametrized value-functions. A traditional choice is the linear function approximation of the form:

$$
Q(s, a)=\theta^{\top} \phi(s, a)
$$

where $\theta \in \Theta \subset \mathbb{R}^{d}$ is a weight vector and $\phi: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}$ is a feature map from a state-action pairs to a given $d$-dimensional feature space.

The natural extension of Q-learning to function approximation with parametric form $\left\{Q_{\theta}, \theta \in \mathbb{R}^{d}\right\}$ is:

$$
\begin{equation*}
\theta_{t+1}=\theta_{t}+\alpha_{t} \delta_{t} \nabla_{\theta} Q_{\theta_{t}}\left(s_{t}, a_{t}\right) \tag{2.5.1}
\end{equation*}
$$

where $\delta_{t}=r_{r}+\gamma \max _{a} Q_{\theta_{t}}\left(s_{t+1}, a\right)-Q_{\theta_{t}}\left(s_{t}, a_{t}\right)$. Note that with linear function, $\nabla_{\theta} Q_{\theta_{t}}\left(s_{t}, a_{t}\right)=\phi\left(s_{t}, a_{t}\right)$ and when features are defined by one-hot encoding vectors, we recover the Q-learning algorithm in the tabular case. Moreover, if we ignore the effect of $\theta$ on the target $y_{t}=r_{r}+\gamma \max _{a} Q_{\theta_{t}}\left(s_{t+1}, a\right)$, the update rule looks like a stochastic gradient descent on:

$$
\begin{equation*}
\mathbb{E}\left[\left(Q_{\theta}(s, a)-y\right)^{2}\right] \tag{2.5.2}
\end{equation*}
$$

However, we include only a part of the gradient by taking into account the effect of changing the weight vector $\theta_{t}$ on the estimate $Q_{\theta_{t}}$, but ignore its effect, via the bootstrapping, on the target $y_{t}$. Such methods are called semi-gradient methods.

Although update rule 2.5.1 is commonly used in practice, not much can be said regarding its convergence properties. In fact, we know that Q-learning with linear function approximation may diverge. Some properties of combining temporal difference with function approximation will be extensively studied in our first contribution.

The limitation of linear functions is that we need to handcraft feature extractors for each domain. This could be painful and time-consuming especially in high-dimensional domains or when a good prior for feature design is not available.

This partially motivates the recent trend of combining reinforcement learning with deep leaning, which is known as deep reinforcement learning (deep RL). Deep RL opens up new applications as it allows us solve a wide range of complex decision-making tasks that were previously out of reach for a machine.

Deep neural networks are made up of multiple processing layers. Each layer consists in an affine transformation followed by non-linear activations. The sequence of non-linear transformations allows to learn different levels of abstractions. This enables representations learning and obviates the need to handcraft features to describe the inputs.

Deep Q-Networks (DQN) [Mnih et al., 2015] incorporates a deep neural network, parameterized by $\theta$, as a function approximator for the action-value function. The neural network parameters are estimated by minimizing the squared temporal difference residual:

$$
\begin{equation*}
L(\theta)=\mathbb{E}_{\left(s, a, r, s^{\prime}\right) \sim \mathcal{D}}\left[\left(Q_{\theta}(s, a)-y\right)^{2}\right], \quad y=r+\gamma \max _{a \in \mathcal{A}} Q_{\theta^{-}}\left(s^{\prime}, a\right) \tag{2.5.3}
\end{equation*}
$$

where transitions $\left(s, a, r=r(s, a), s^{\prime} \sim P\left(s^{\prime} \mid s, a\right)\right)$ are sampled from a replay buffer of recent observed transitions $\mathcal{D}$. Here $\theta^{-}$denotes the parameters of a target network which is updated $\left(\theta^{-} \leftarrow \theta\right)$ regularly and held fixed between individual updates of $\theta$.

### 2.5.2 Policy Gradient

Policy gradients methods search over a parameterized class of polices by performing stochastic gradient ascent on an objective function capturing the cumulative expected rewards. Specifically, they consider the scalar objective function

$$
\begin{equation*}
J^{\pi} \triangleq \mathbb{E}_{s \sim \rho}\left[V^{\pi}(s)\right] \tag{2.5.4}
\end{equation*}
$$

which averages the total value function $V^{\pi}$ over a random initial state distribution $\rho$.


Figure 2.1: Restaurant selection example. Exploitation: Go to favorite restaurant. Exploration: Try a new restaurant

Let's first define the discounted state visitation distribution $d_{\rho}^{\pi_{\theta}}$ of policy $\pi$ :

$$
\begin{equation*}
d_{\rho}^{\pi}(s)=(1-\gamma) \sum_{t=0}^{\infty} \gamma^{t} \operatorname{Pr}\left(s_{t}=s \mid s_{0} \sim \rho, \pi\right) \tag{2.5.5}
\end{equation*}
$$

where $\operatorname{Pr}\left(s_{t}=s \mid s_{0} \sim \rho, \pi\right)$ is the state visitation probability that $s_{t}=s$, after we execute $\pi$ starting at a state $s_{0}$ sampled from $\rho$.

For the class of stochastic differentiable policies $\left\{\pi_{\theta}, \theta \in \Theta \subset \mathbb{R}^{d}\right\}$, the fundamental result underlying these algorithms is the policy gradient theorem [Sutton et al., 1999a]:

$$
\begin{equation*}
\nabla_{\theta} J^{\pi_{\theta}}=\frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{\rho}^{\pi_{\theta}}, a \sim \pi_{\theta}}\left[\nabla_{\theta}\left(\log \pi_{\theta}(a \mid s)\right) Q^{\pi_{\theta}}(s, a)\right] \tag{2.5.6}
\end{equation*}
$$

The latter expression suggests that we can estimate the gradient by approximating the action-value function. It can be done by $\mathrm{TD}(0)$ with function approximation. This gave rise to the actor-critic architecture [Konda and Tsitsiklis, 2000].

Another instance of policy-gradient is the REINFORCE algorithm [Williams, 1992] which, instead of a learned value function, uses the cumulative return averaged over rollouts generated by the policy $\pi_{\theta}$.

### 2.6 EXPLOITATION-EXPLORATION DILEMMA

A RL agent aims at maximizing its long-term utility by exploiting its knowledge about the environment and allocating resources to the choices that are identified to be highly rewarding. However, this knowledge has to be acquired by the agent itself through exploring different choices that may reduce short-term utility. An illustrative situation is depicted in figure 2.1. An agent has the choice between going to its favorite restaurant so far or trying a new restaurant. A conservative agent would stick to its usual place that may be suboptimal choice. An active learner would experiment the new restaurant to figure out if it is worth


Figure 2.2: n-chain environment
going back to it and gain more knowledge. An efficient learning should balance carefully between exploration and exploitation.

The simplest approach to exploration is the $\epsilon$-greedy strategy, which selects the greedy action with probability $1-\epsilon$ and otherwise selects uniformly at random from all currently available actions. It is commonly used because it is easy to implement and disentangles explicitly between exploration and exploitation. However, such an approach does not take into account the level of information that actions could bring, and thus wastes exploratory resources.

To illustrate the inefficiency of random exploration, let us consider the nchain environment in Figure 2.2. The environment consists of n states. The agent always starts at the second state $s_{2}$ and has two possible actions: move right or move left. A small reward $r=0.001$ is received in the first state $s_{1}$, a large reward $r=1$ in the final state $s_{n}$, otherwise the reward is zero. Q-learning with $\epsilon$ greedy does poorly in this example. It takes $O\left(2^{n}\right)$ samples in expectation to visit the right end state for one time, resulting in an exponential sample complexity. What makes exploration particularly challenging in this toy example is that the agent only observes the informative signal after committing to the action right over an extended period of time. Therefore a directed exploration, involving a planning over time, is critical for effective learning.

Optimism in the face of uncertainty (OFU) is one of the traditional guiding principles that offers provably efficient learning algorithms. We can distinguish two classes of approaches: confidence-intervals based methods [Kearns and Singh, 2002, Strehl and Littman, 2005, Jaksch et al., 2010] and exploration-bonus based methods [Azar et al., 2017, Jin et al., 2018, Jian et al., 2019]. In the former, the agent builds a set of statistically plausible Markov Decision Processes (MDPs) that contains the true MDP with high probability. Then, the agent selects the most optimistic version of its model and acts optimally with respect to it. In the latter, discoveries of poorly understood states and actions are rewarded by an exploration bonus. Such bonus is designed to bound estimation errors on the value function.

An entire body of algorithms for efficient exploration is inspired by Thompson sampling [Thompson, 1933]. Bayesian dynamic programming was first introduced
in Strens [2000] and is more recently known as posterior sampling for reinforcement learning (PSRL) [Osband et al., 2013]. In PSRL, the agent starts with a prior belief over world models and then proceeds to update its full posterior distribution over models with the newly observed samples. A model hypothesis is then sampled from this distribution, and a greedy policy with respect to the sampled model is followed thereafter.

Why, intuitively, does PSRL works? let $H_{k}$ the history of observations up to episode $k$. The posterior distribution $f\left(. \mid H_{k}\right)$ generates only plausible models and it concentrates around the True MDP.

- In early iterations, the posterior distribution is flat and its spread capture the uncertainty which enables exploration.
- As we collect more data, the mode of the posterior distribution captures the maximum likelihood estimates, which enable exploitation.


### 2.6.1 Online Performance

In order to quantitatively assess the exploration-exploitation trade-off, we use some online performance measures. the most common performance measures widely used in the literature are the sample complexity of exploration [Kakade, 2003] (defined for discounted and episodic MDPs) and the regret (defined traditionally for the average-reward and episodic MDPs and extended recently to discounted MDPs [Liu and Su, 2020])

In episodic MDPs, if an algorithm follows policy $\pi_{k}$ in episode $k$, then the optimality gap in episode $k$ is $\Delta_{k} \triangleq V_{1}^{*}\left(s_{1}\right)-V_{1}^{\pi_{k}}\left(s_{1}\right)$, where $V_{1}^{*}$ and $V_{1}^{\pi_{k}}$ is respectively the optimal value function and the value function of $\pi_{k}$ at stage $h=1$, and $s_{1}$ is the initial state.

The sample complexity of exploration is defined as $N_{(\epsilon, K)} \triangleq \sum_{k=1}^{K} \mathbb{1}_{\left\{\Delta_{k}>\epsilon\right\}}$ i.e the total number of errors ( $\epsilon$-suboptimal policies). We say that a learning algorithm is $(\epsilon, \delta)$-PAC (Probably Approximately Correct) if there exists a polynomial function of $1 / \epsilon, \delta$ and MDP's parameters such that $\left.N_{( } \epsilon, K\right)$ is smaller than this function with probability at least $1-\delta$. In other term, the number of errors can be bounded by a number polynomial in the relevant quantities with high probability.

The regret is defined as the total sum of optimality gap over episodes: $\operatorname{REGRET}(K)=\sum_{k=1}^{K} \Delta_{k}$. A learning algorithm is said "efficient" if it achieves a sublinear regret i.e REGRET $(K)=o(K)$ when $K \leftarrow \infty$ in expectation or with high probability. A high probability bound is a stronger results as it can always be converted to a bound in expectation. Another quantity is the expected Bayesian regret where we consider the expectation of the regret over a prior distribution on MDPs. This is a weaker guarantee as the measure holds only in expectation over MDPs and not for a specific MDP.

The tow aforementioned performance measures are not directly comparable. An algorithm with small regret can make small mistakes infinitely often and thus it cannot be $(\epsilon, \delta)$-PAC for $\epsilon$ small enough. Dann et al. [2017] provide an extensive study of the connection between the two quantities.

## Multi-step Off-policy Learning with Function Approximation

> «Before a decision is made, a set of options constitute candidates for future reality. After the decision has been made and its consequences are known, those unchosen options can never become reality, but they may nevertheless haunt us, amuse us, or influence our perception of the decision process itself.» - [Roese, 1999]

### 3.1 PROLOGUE TO THE CONTRIBUTION

### 3.1.1 Article Details

This chapter is based on the article «Convergent Tree-backup and Retrace with Function Approximation» [Touati et al., 2018]. This is joint work with PierreLuc Bacon, Doina Precup and Pascal Vincent. This paper was accepted as a long oral at ICML 2018. I am the main author. I performed the theoretical analysis, the derivation of the new algorithm and the experiments.

### 3.1.2 Context

This work started originally as a course project in the context of the RL class taught by Doina Precup at Mcgill university. I was interested in optimizationbased RL algorithms. Pierre-luc Bacon, who was the teaching assistant at the time, suggested to me to study the behavior of RETRACE algorithm with function approximation. RETRACE was recently introduced as a safe and efficient off-policy method in Munos et al. [2016] and became quickly at the heart of several successful deep RL architectures such as REACTOR (Retrace-Actor) [Gruslys et al., 2017] and ACER [Wang et al., 2016].

### 3.1.3 Paper Abstract

Off-policy learning is key to scaling up reinforcement learning as it allows to learn about a target policy from the experience generated by a different behavior policy. Unfortunately, it has been challenging to combine off-policy learning
with function approximation and multi-step bootstrapping in a way that leads to both stable and efficient algorithms.

In this work, we show that popular multi-step off-policy methods such as Tree Backup [Precup, 2000] and Retrace [Munos et al., 2016] are unstable with linear function approximation, both in theory and in practice with specific examples. Based on our analysis, we then derive stable and efficient gradientbased algorithms using a quadratic convex-concave saddle-point formulation. By exploiting the problem structure proper to these algorithms, we are able to provide convergence guarantees and finite-sample bounds. The applicability of our new analysis also goes beyond TREE BACKUP and RETRACE and allows us to provide new state-of-the-art convergence rates for the GTD and GTD2 algorithms without having recourse to projections or Polyak averaging.

### 3.1.4 Recent Developments

While our work focuses on linear function approximation, Wai et al. [2019] extend the primal-dual formulation of policy evaluation to nonlinear smooth function approximation, show that it converges to a stationary point and provide its finite sample analysis.

I worked on a followup in [Peng et al., 2019], jointly led by Zilun Peng and in collaboration with Pascal Vincent and Doina Precup. We proposed a reduced variance methods for the problem of off-policy policy evaluation that achieve a linear convergence rate while addressing the computational bottleneck that previously proposed reduced variance methods in Du et al. [2017] suffer from.

### 3.2 InTRODUCTION

Rather than being confined to their own stream of experience, off-policy learning algorithms are capable of leveraging data from a different behavior than the one being followed, which can provide many benefits: efficient parallel exploration as in Mnih et al. [2016] and Wang et al. [2016], reuse of past experience with experience replay [Lin, 1992] and, in many practical contexts, learning form data produced by policies that are currently deployed, but which we want to improve (as in many scenarios of working with an industrial or health care partner). Moreover, a single stream of experience can be used to learn about a variety of different targets which may take the form of value functions corresponding to different policies and time scales [Sutton et al., 1999c] or to predicting different reward functions as in Sutton and Tanner [2004] and Sutton et al. [2011]. Therefore, the design and analysis of off-policy algorithms using all the features of reinforcement learning, e.g. bootstrapping, multi-step updates (eligibility traces), and function approximation has been explored extensively over
three decades. While off-policy learning and function approximation have been understood in isolation, their combination with multi-steps bootstrapping produces a so-called deadly triad [Sutton, 2015, Sutton and Barto, 2018], i.e., many algorithms in this category are unstable.

A typical off-policy approach to is provided by importance sampling, which bends the behavior policy distribution onto the target one [Precup, 2000, Precup et al., 2001]. However, as the length of the trajectories increases, the variance of importance sampling corrections tends to become very large. The Tree Backup algorithm [Precup, 2000] is an alternative approach which remarkably does not rely on importance sampling ratios directly. More recently, Munos et al. [2016] introduced the RETRACE algorithm which also builds on TREE BACKUP to perform off-policy learning without importance sampling.

Until now, Tree Backup and Retrace $(\lambda)$ had only been shown to converge in the tabular case, and their behavior with linear function approximation was not known. In this paper, we show that this combination with linear function approximation is in fact divergent. We obtain this result by analyzing the mean behavior of TREE BACKUP and Retrace using the ordinary differential equation (ODE) [Borkar and Meyn, 2000] associated with them. We also demonstrate this instability with a concrete counterexample.

Insights gained from this analysis allow us to derive a new gradient-based algorithm with provable convergence guarantees. Instead of adapting the derivation of Gradient Temporal Difference (GTD) learning from [Sutton et al., 2009c], we use a primal-dual saddle point formulation [Liu et al., 2015, Macua et al., 2015] which facilitates the derivation of sample complexity bounds. The underlying saddle-point problem combines the primal variables, function approximation parameters, and dual variables through a bilinear term.

In general, stochastic primal-dual gradient algorithms like the ones derived in this paper can be shown to achieve $O(1 / k)$ convergence rate (where $k$ is the number of iterations). For example, this has been established for the class of forward-backward algorithms with added noise [Rosasco et al., 2016]. Furthermore, this work assumes that the objective function is composed of a convexconcave term and a strongly convex-concave regularization term that admits a tractable proximal mapping. In this paper, we are able to achieve the same $O(1 / k)$ convergence rate without having to assume strong convexity with respect to the primal variables and in the absence of proximal mappings. As corollary, our convergence rate result extends to the well-known gradient-based temporal difference algorithms GTD [Sutton et al., 2009c] and GTD2 [Sutton et al., 2009b] and hence improves the previously published results.

The algorithms resulting from our analysis are simple to implement, and perform well in practice compared to other existing multi-steps off-policy learning algorithms such as $G Q(\lambda)$ [Maei and Sutton, 2010] and AB-TrACE $(\lambda)$ [Mahmood et al., 2017].

### 3.3 Tabular Off-policy Methods

In this work, we are concerned with the policy evaluation problem [Sutton and Barto, 1998] under model-free off-policy learning. That is, we will evaluate a target policy $\pi$ using trajectories (i.e. sequences of states, actions and rewards) obtained from a different behavior policy $\mu$.

Munos et al. [2016] provided a unified perspective on several off-policy learning algorithms, namely: those using explicit importance sampling corrections [Precup, 2000] as well as Tree Backup (TB $(\lambda)$ ) [Precup, 2000] and $Q^{\pi}(\lambda)$ [Harutyunyan et al., 2016] which do not involve importance ratios. As a matter of fact, all these methods share a general form based on the $\lambda$-return [Sutton and Barto, 2018] but involve different coefficients $\kappa_{i}$ in :

$$
\begin{aligned}
G_{k}^{\lambda} & \triangleq Q\left(s_{k}, a_{k}\right)+\sum_{t=k}^{\infty}(\lambda \gamma)^{t-k}\left(\prod_{i=k+1}^{t} \kappa_{i}\right)\left(r_{t}+\gamma \mathbb{E}_{\pi} Q\left(s_{t+1}, \cdot\right)-Q\left(s_{t}, a_{t}\right)\right) \\
& =Q\left(s_{k}, a_{k}\right)+\sum_{t=k}^{\infty}(\lambda \gamma)^{t-k}\left(\prod_{i=k+1}^{t} \kappa_{i}\right) \delta_{t}
\end{aligned}
$$

where $\mathbb{E}_{\pi} Q\left(s_{t+1},.\right) \triangleq \sum_{a \in \mathcal{A}} \pi\left(a \mid s_{t+1}\right) Q\left(s_{t+1}, a\right)$ and $\delta_{t} \triangleq r_{t}+\gamma \mathbb{E}_{\pi} Q\left(s_{t+1},.\right)-$ $Q\left(s_{t}, a_{t}\right)$ is the temporal-difference (TD) error. The coefficients $\kappa_{i}$ determine how the TD errors would be scaled in order to correct for the discrepancy between target and behavior policies. From this unified representation, Munos et al. [2016] derived the $\operatorname{Retrace}(\lambda)$ algorithm. Both $\operatorname{TB}(\lambda)$ and $\operatorname{Retrace}(\lambda)$ consider this form of return, but set $\kappa_{i}$ differently. The $\operatorname{TB}(\lambda)$ updates correspond to the choice $\kappa_{i}=\pi\left(a_{i} \mid s_{i}\right)$ while $\operatorname{RETRACE}(\lambda)$ sets $\kappa_{i}=\min \left(1, \frac{\pi\left(a_{i} \mid s_{i}\right)}{\left.\mu\left(a_{i}\right) s_{i}\right)}\right)$, which is intended to allow learning from full returns when the target and behavior policies are very close. The importance sampling approach [Precup, 2000] converges in the tabular case by correcting the behavior data distribution to the distribution that would be induced by the target policy $\pi\left(\kappa_{i}=\frac{\pi\left(a_{i} \mid s_{i}\right)}{\mu\left(a_{i} \mid s_{i}\right)}\right)$.However, these correction terms lead to high variance in practice. Since $Q(\lambda)$ does not involve importance ratios ( $\kappa_{i}=1$ ), this variance problem is avoided but at the cost of restricted convergence guarantees satisfied only when the behavior and target policies are sufficiently close.

The analysis provided in this paper concerns $\operatorname{TB}(\lambda)$ and $\operatorname{RETRACE}(\lambda)$, which are convergent in the tabular case, but have not been analyzed in the function approximation case. We start by noting that the Bellman operator ${ }^{1} \mathcal{R}$ underly-

[^1]|  | Value of $\kappa_{i}$ | $l_{\infty}$ contraction <br> guarantee | Estimation <br> variance |
| :---: | :---: | :---: | :---: |
| Importance | $\frac{\pi\left(a_{i} \mid s_{i}\right)}{\mu\left(a_{i} \mid s_{i}\right)}$ | for any $\pi, \mu$ | High |
| sampling | 1 | only for $\pi$ close to $\mu$ | Low |
| $Q^{\pi}(\lambda)$ | $\pi\left(a_{i} \mid s_{i}\right)$ | for any $\pi, \mu$ | Low |
| $\operatorname{TB}(\lambda)$ | $\min \left(1, \frac{\pi\left(a_{i} \mid s_{i}\right)}{\mu\left(a_{i} \mid s_{i}\right)}\right)$ | for any $\pi, \mu$ | Low |
| RETRACE $(\lambda)$ | $\operatorname{men}$ |  |  |

Table 3.1: Properties of different off-policy algorithms for policy evaluation.
ing these algorithms can be written in the following form:

$$
\begin{aligned}
(\mathcal{R} Q)(s, a) & \triangleq Q(s, a)+\mathbb{E}_{\mu}\left[\sum_{t=0}^{\infty}(\lambda \gamma)^{t}\left(\prod_{i=1}^{t} \kappa_{i}\right)\left(r_{t}+\gamma \mathbb{E}_{\pi} Q\left(s_{t+1}, \cdot\right)-Q\left(s_{t}, a_{t}\right)\right)\right] \\
& =Q(s, a)+\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1}\left(\mathcal{T}^{\pi} Q-Q\right)(s, a)
\end{aligned}
$$

where $\mathbb{E}_{\mu}$ is the expectation over the behavior policy and MDP transition probabilities and $P^{\kappa \mu}$ is the operator defined by:

$$
\left(P^{\kappa \mu} Q\right)(s, a) \triangleq \sum_{\substack{s^{\prime} \in \mathcal{S} \\ a^{\prime} \in \mathcal{A}}} P\left(s^{\prime} \mid s, a\right) \mu\left(a^{\prime} \mid s^{\prime}\right) \kappa\left(s^{\prime}, a^{\prime}\right) Q\left(s^{\prime}, a^{\prime}\right)
$$

In the tabular case, these operators were shown to be contraction mappings with respect to the supremum norm [Precup, 2000, Munos et al., 2016]. In this paper, we focus on what happens to these operators when combined with linear function approximation.

### 3.4 OFF-POLICY INSTABILITY WITH FUNCTION APPROXIMATION

In order to obtain generalization between different state-action pairs, $Q^{\pi}$ should be represented in a functional form. In this paper, we focus on linear function approximation of the form:

$$
Q(s, a) \triangleq \theta^{\top} \phi(s, a)
$$

where $\theta \in \Theta \subset \mathbb{R}^{d}$ is a weight vector and $\phi: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}$ is a feature map from a state-action pairs to a given $d$-dimensional feature space.

When combined with function approximation, the temporal difference updates corresponding to the $\lambda$-return $G_{k}^{\lambda}$ are given by

$$
\begin{align*}
\theta_{k+1} & =\theta_{k}+\alpha_{k}\left(G_{k}^{\lambda}-Q\left(s_{k}, a_{k}\right)\right) \nabla_{\theta} Q\left(s_{k}, a_{k}\right) \\
& =\theta_{k}+\alpha_{k}\left(\sum_{t=k}^{\infty}(\lambda \gamma)^{t-k}\left(\prod_{i=k+1}^{t} \kappa_{i}\right) \delta_{t}^{k}\right) \phi\left(s_{k}, a_{k}\right) \tag{3.4.1}
\end{align*}
$$

where $\delta_{t}^{k}=r_{t}+\gamma \theta_{k}^{\top} \mathbb{E}_{\pi} \phi\left(s_{t+1}, \cdot\right)-\theta_{k}^{\top} \phi\left(s_{t}, a_{t}\right)$ and $\alpha_{k}$ are positive non-increasing step sizes. The updates (3.4.1) implies off-line updating as $G_{k}^{\lambda}$ is a quantity which depends on future rewards. This will be addressed later using eligibility traces: a mechanism to transform the off-line updates into efficient on-line ones. Since (3.4.1) describes stochastic updates, the following standard assumption is necessary:

Assumption 1. The Markov chain induced by the behavior policy $\mu$ is ergodic and admits a unique stationary distribution, denoted by $\xi$, over state-action pairs. We write $\Xi$ for the diagonal matrix whose diagonal entries are $(\xi(s, a))_{s \in \mathcal{S}, a \in \mathcal{A}}$.

Our first proposition establishes the expected behavior of the parameters in the limit.

Proposition 1. If the behavior policy satisfies Assumption 1 and $\left(\theta_{k}\right)_{k \leq 0}$ is the Markov process defined by (3.4.1) then:

$$
\mathbb{E}\left[\theta_{k+1} \mid \theta_{0}\right]=\left(I+\alpha_{k} A\right) \mathbb{E}\left[\theta_{k} \mid \theta_{0}\right]+\alpha_{k} b
$$

where matrix $A$ and vector $b$ are defined as follows:

$$
\begin{aligned}
& A \triangleq \Phi^{\top} \Xi\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1}\left(\gamma P^{\pi}-I\right) \Phi \\
& b \triangleq \Phi^{\top} \Xi\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1} r
\end{aligned}
$$

Sketch of Proof (The full proof is in the appendix).

$$
\begin{aligned}
\theta_{k+1} & =\theta_{k}+\alpha_{k}\left(\sum_{t=k}^{\infty}(\lambda \gamma)^{t-k}\left(\prod_{i=k+1}^{t} \kappa_{i}\right) \phi\left(s_{k}, a_{k}\right)\left(\left[\gamma \mathbb{E}_{\pi} \phi\left(x_{t+1}, \cdot\right)-\phi\left(x_{t}, a_{t}\right)\right]^{\top} \theta_{k}+r_{t}\right)\right) \\
& =\theta_{k}+\alpha_{k}\left(A_{k} \theta_{k}+b_{k}\right)
\end{aligned}
$$

So, $\mathbb{E}\left[\theta_{k+1} \mid \theta_{k}\right]=\left(I+\alpha_{k} A\right) \theta_{k}+\alpha_{k} b$ where $A=\mathbb{E}\left[A_{k}\right]$ and $b=\mathbb{E}\left[b_{k}\right]$
The ODE (Ordinary Differential Equations) approach [Borkar and Meyn, 2000] is the main tool to establish convergence in the function approximation case [Bertsekas and Tsitsiklis, 1995, Tsitsiklis et al., 1997]. In particular, we use Proposition 4.8 in Bertsekas and Tsitsiklis [1995], which states that under some
conditions, $\theta_{k}$ converges to the unique solution $\theta^{*}$ of the system $A \theta^{*}+b=0$. This crucially relies on the matrix $A$ being negative definite i.e $y^{\top} A y<0, \forall y \neq 0$. In the on-policy case, when $\mu=\pi$, we rely on the fact that the stationary distribution is invariant under the transition matrix $P^{\pi}$ i.e $d^{\top} P^{\pi}=d^{\top}$ [Tsitsiklis et al., 1997, Sutton et al., 2015]. However, this is no longer true for off-policy learning with arbitrary target/behavior policies and the matrix $A$ may not be negative definite: the series $\theta_{k}$ may then diverge. We will now see that the same phenomenon may occur with $\operatorname{TB}(\lambda)$ and $\operatorname{RETRACE}(\lambda)$.

Counterexample: We extend the two-states MDP of Tsitsiklis et al. [1997], originally proposed to show the divergence of off-policy $\mathrm{TD}(0)$, to the case of function approximation over state-action pairs. This environment has only two states, as shown in Figure 3.1, and two actions: left or right.


Figure 3.1: Two-state counterexample. We assign the features $\left\{(1,0)^{\top},(2,0)^{\top},(0,1)^{\top},(0,2)^{\top}\right\}$ to the state-action pairs $\{(1$, right $),(2$, right $),(1$, left $),(2$, left $)\}$. The target policy is given by $\pi($ right $\mid \cdot)=1$ and the behavior policy is $\mu(\operatorname{right} \mid \cdot)=0.5$

In this particular case, both $\operatorname{TB}(\lambda)$ and $\operatorname{RETRACE}(\lambda)$ share the same matrix $P^{\kappa \mu}$ and $P^{\kappa \mu}=0.5 P^{\pi}$ :

$$
P^{\pi}=\left(\begin{array}{llll}
0 & 1 & 0 & 0 \\
0 & 1 & 0 & 0 \\
1 & 0 & 0 & 0 \\
1 & 0 & 0 & 0
\end{array}\right),\left(P^{\pi}\right)^{n}=\left(\begin{array}{llll}
0 & 1 & 0 & 0 \\
0 & 1 & 0 & 0 \\
0 & 1 & 0 & 0 \\
0 & 1 & 0 & 0
\end{array}\right) \forall n \geq 2
$$

If we set $\beta:=0.5 \gamma \lambda$, we then have:

$$
\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1}=\left(\begin{array}{cccc}
1 & \frac{\beta}{1-\beta} & 0 & 0 \\
0 & \frac{1}{1-\beta} & 0 & 0 \\
\beta & \frac{\beta^{2}}{1-\beta} & 1 & 0 \\
\beta & \frac{\beta^{2}}{1-\beta} & 0 & 1
\end{array}\right), A=\left(\begin{array}{cc}
\frac{6 \gamma-\beta-5}{1-\beta} & 0 \\
\frac{3\left(\gamma \beta-\beta^{2}-\beta-\gamma\right)}{1-\beta} & -5
\end{array}\right)
$$

Therefore, $\forall \gamma \in\left(\frac{5}{6}, 1\right)$ and $\forall \lambda \in\left[0, \min \left(1, \frac{12 \gamma-10}{\gamma}\right)\right)$, the first eigenvalue $e_{1}=\frac{6 \gamma-\beta-5}{1-\beta}$ of $A$ is positive. The basis vectors $(1,0)^{\top}$ and $(0,1)^{\top}$ are eigenvectors of A associated with $e_{1}$ and -5 , then if $\theta_{0}=\left(\eta_{1}, \eta_{2}\right)^{\top}$, we obtain
$\mathbb{E}\left[\theta_{k} \mid \theta_{0}\right]=\left(\eta_{1} \prod_{i=0}^{k-1}\left(1+\alpha_{i} e_{1}\right), \eta_{2} \prod_{i=0}^{k-1}\left(1-5 \alpha_{i}\right)\right)^{\top}$ implying that $\left\|\mathbb{E}\left[\theta_{k} \mid \theta_{0}\right]\right\| \geq$ $\left|\eta_{1}\right| \prod_{i=0}^{k-1}\left(1+\alpha_{i} e_{1}\right)$. Hence, as $\sum_{k} \alpha_{k} \rightarrow \infty,\left\|\mathbb{E}\left[\theta_{k} \mid \theta_{0}\right]\right\| \rightarrow \infty$ if $\eta_{1} \neq 0$.

### 3.5 CONVERGENT GRADIENT OFF-POLICY ALGORITHMS

If $A$ were to be negative definite, $\operatorname{Retrace}(\lambda)$ or $\operatorname{TB}(\lambda)$ with function approximation would converge to $\theta^{*}=-A^{-1} b$. It is known [Bertsekas, 2011] that $\Phi \theta^{*}$ is the fixed point of the projected Bellman operator :

$$
\Phi \theta^{*}=\Pi^{\mu} \mathcal{R}\left(\Phi \theta^{*}\right),
$$

where $\Pi^{\mu}=\Phi\left(\Phi^{\top} \Xi \Phi\right)^{-1} \Phi^{\top} \Xi$ is the orthogonal projection onto the space $S=\left\{\Phi \theta \mid \theta \in \mathbb{R}^{d}\right\}$ with respect to the weighted Euclidean norm $\||\cdot|\|_{\Xi}$. Rather than computing the sequence of iterates given by the projected Bellman operator, another approach for finding $\theta^{*}$ is to directly minimize [Sutton et al., 2009a, Liu et al., 2015] the Mean Squared Projected Bellman Error (MSPBE):

$$
\operatorname{MSPBE}(\theta)=\frac{1}{2}\left\|\Pi^{\mu} \mathcal{R}(\Phi \theta)-\Phi \theta\right\|_{\mathbb{E}}^{2} .
$$

This is the route that we take in this paper to derive convergent forms of $\operatorname{TB}(\lambda)$ and $\operatorname{Retrace}(\lambda)$. To do so, we first define our objective function in terms of $A$ and $b$ which we introduced in Proposition 1.
Proposition 2. Let $M \triangleq \Phi^{\top} \Xi \Phi=\mathbb{E}\left[\Phi \Phi^{\top}\right]$ be the covariance matrix of features. We have:

$$
\operatorname{MSPBE}(\theta)=\frac{1}{2}\|A \theta+b\|_{M^{-1}}^{2}
$$

(The proof is provided in the appendix.)
In order to derive parameter updates, we could compute gradients of the above expression explicitly as in Sutton et al. [2009c], but we would then obtain a gradient that is a product of expectations. The implied double sampling makes it difficult to obtain an unbiased estimator of the gradient. Sutton et al. [2009c] addressed this problem with a two-timescale stochastic approximations. However, the algorithm obtained in this way is no longer a true stochastic gradient method with respect to the original objective. Liu et al. [2015] suggested an alternative which converts the original minimization problem into a primal-dual saddle-point problem. This is the approach that we chose in this paper.

The convex conjugate of a real-valued function $f$ is defined as:

$$
\begin{equation*}
f^{*}(y)=\sup _{x \in X}(\langle y, x\rangle-f(x)), \tag{3.5.1}
\end{equation*}
$$

and $f$ is convex, we have $f^{* *}=f$. Also, if $f(x)=\frac{1}{2}\|x\|_{M^{-1}}$, then $f^{*}(x)=$ $\frac{1}{2}\|x\|_{M}$. Note that by going to the convex conjugate, we do not need to invert matrix $M$. We now go back to the original minimization problem:

$$
\begin{aligned}
\min _{\theta} \operatorname{MSPBE}(\theta) & \Leftrightarrow \min _{\theta} \frac{1}{2}\|A \theta+b\|_{M^{-1}}^{2} \\
& \Leftrightarrow \min _{\theta} \max _{\omega}\left(\langle A \theta+b, \omega\rangle-\frac{1}{2}\|\omega\|_{M}^{2}\right)
\end{aligned}
$$

The gradient updates resulting from the saddle-point problem (ascent in $\omega$ and descent in $\theta$ ) are then:

$$
\begin{align*}
\omega_{k+1} & =\omega_{k}+\eta_{k}\left(A \theta_{k}+b-M \omega_{k}\right) \\
\theta_{k+1} & =\theta_{k}-\alpha_{k} A^{\top} \omega_{k} \tag{3.5.2}
\end{align*}
$$

where $\left\{\eta_{k}\right\}$ and $\left\{\alpha_{k}\right\}$ are non-negative step-size sequences. As the $A, b$ and $M$ are all expectations, we can derive stochastic updates by drawing samples, which would yield unbiased estimates of the gradient.

On-line updates: We now derive on-line updates by exploiting equivalences in expectation between forward views and backward views outlined in Maei [2011].

Proposition 3. Let $e_{k}$ be the eligibility traces vector, defined as $e_{-1}=0$ and :

$$
e_{k}=\lambda \gamma \kappa\left(s_{k}, a_{k}\right) e_{k-1}+\phi\left(s_{k}, a_{k}\right) \quad \forall k \geq 0
$$

Furthermore, let $\left.\hat{A}_{k}=e_{k}\left(\gamma \mathbb{E}_{\pi}\left[\phi\left(s_{k+1},.\right)\right]-\phi\left(s_{k}, a_{k}\right)\right]\right)^{\top}, \quad \hat{b}_{k}=r\left(s_{k}, a_{k}\right) e_{k}, \quad \hat{M}_{k}=$ $\phi\left(s_{k}, a_{k}\right) \phi\left(s_{k}, a_{k}\right)^{\top}$. Then, we have $\mathbb{E}\left[\hat{A}_{k}\right]=A, \mathbb{E}\left[\hat{b}_{k}\right]=b$ and $\mathbb{E}\left[\hat{M}_{k}\right]=M$.
(The proof is provided in the appendix.)
This proposition allows us to replace the expectations in Eq. (3.5.2) by corresponding unbiased estimates. The resulting detailed procedure is provided in Algorithm 3.

### 3.6 Convergence Rate Analysis

In order to characterize the convergence rate of the algorithm 3, we need to introduce some new notations and state new assumptions.

We denote by $\|A\| \triangleq \sup _{\|x\|=1}\|A x\|$ the spectral norm of the matrix A and by $c(A)=\|A\|\left\|A^{-1}\right\|$ its condition number. If the eigenvalues of a matrix $A$

```
Algorithm 3 Gradient Off-policy with eligibility traces
Given: target policy \(\pi\), behavior policy \(\mu\)
    Initialize \(\theta_{0}\) and \(\omega_{0}\)
    for \(\mathrm{n}=0 \ldots\) do
        set \(e_{0}=0\)
        for \(\mathrm{k}=0 \ldots\) end of episode do
            Observe \(s_{k}, a_{k}, r_{k}, s_{k+1}\) according to \(\mu\)
            /* Update traces
            \(e_{k}=\lambda \gamma \kappa\left(s_{k}, a_{k}\right) e_{k-1}+\phi\left(s_{k}, a_{k}\right)\)
            \({ }^{*}\) Update parameters
            \(\delta_{k}=r_{k}+\gamma \theta_{k}^{\top} \mathbb{E}_{\pi} \phi\left(s_{k+1},.\right)-\theta_{k}^{\top} \phi\left(s_{k}, a_{k}\right)\)
            \(\omega_{k+1}=\omega_{k}+\eta_{k}\left(\delta_{k} e_{k}-\omega_{k}^{\top} \phi\left(s_{k}, a_{k}\right) \phi\left(s_{k}, a_{k}\right)\right)\)
            \(\theta_{k+1}=\theta_{k}-\alpha_{k} \omega_{k}^{\top} e_{k}\left(\gamma \mathbb{E}_{\pi} \phi\left(s_{k+1},.\right)-\phi\left(s_{k}, a_{k}\right)\right)\)
        end for
    end for
```

are real, we use $\lambda_{\max }(A)$ and $\lambda_{\min }(A)$ to denote respectively the largest and the smallest eigenvalue.

If we set $\eta_{k}=\beta \alpha_{k}$ for a positive constant $\beta$, it is possible to combine the two iterations present in our algorithm as a single iteration using a parameter vector $z_{k} \triangleq\binom{\theta_{k}}{\frac{1}{\sqrt{\beta}} \omega_{k}}$ where :

$$
z_{k+1}=z_{k}-\alpha_{k}\left(\hat{G}_{k} z_{k}-\hat{g}_{k}\right)
$$

where:

$$
\hat{G}_{k} \triangleq\left(\begin{array}{cc}
0 & \sqrt{\beta} \hat{A}_{k}^{\top} \\
-\sqrt{\beta} \hat{A}_{k} & \beta \hat{M}_{k}
\end{array}\right) \quad \hat{g}_{k} \triangleq\binom{0}{\sqrt{\beta} \hat{b}_{k}}
$$

Let $G \triangleq \mathbb{E}\left[\hat{G}_{k}\right]$ and $g=\mathbb{E}\left[\hat{g}_{k}\right]$. It follows from the proposition 3 that $G$ and $g$ are well defined and more specifically:

$$
G=\left(\begin{array}{cc}
0 & \sqrt{\beta} A^{\top} \\
-\sqrt{\beta} A & \beta M
\end{array}\right) \quad g=\binom{0}{\sqrt{\beta} b}
$$

Furthermore, let $\mathcal{F}_{k}=\sigma\left(z_{0}, \hat{G}_{0}, \hat{g}_{0} \ldots, z_{k}, \hat{G}_{k}, \hat{g}_{k}, z_{k+1}\right)$ be the sigma-algebra generated by the variables up to time $k$. With these definitions, we can now state our assumptions.

Assumption 2. The matrices $A$ and $M$ are nonsingular. This implies that the saddlepoint problem admits a unique solution $\left(\theta^{\star}, \omega^{\star}\right)=\left(-A^{-1} b, 0\right)$ and we define $z^{\star} \triangleq$ $\left(\theta^{\star}, \frac{1}{\sqrt{\beta}} \omega^{\star}\right)$.

Assumption 3. The features and reward functions are uniformly bounded. This implies that the features and rewards have uniformly bounded second moments. It follows that there exists a constant $\sigma$ such that:

$$
\mathbb{E}\left[\left\|\hat{G}_{k} z_{k}-\hat{g}_{k}\right\|^{2} \mid \mathcal{F}_{k-1}\right] \leq \sigma^{2}\left(1+\left\|z_{k}\right\|^{2}\right)
$$

Before stating our main result, the following key quantities needs to be defined:

$$
\begin{gathered}
\rho \triangleq \lambda_{\max }\left(A^{\top} M^{-1} A\right), \quad \delta \triangleq \lambda_{\min }\left(A^{\top} M^{-1} A\right), \\
L_{G} \triangleq\left\|\mathbb{E}\left[\hat{G}_{k}^{\top} \hat{G}_{k} \mid \mathcal{F}_{k-1}\right]\right\|
\end{gathered}
$$

The following proposition characterize the convergence in expectation of $\| z_{k}-$ $z^{\star}\left\|^{2}=\right\| \theta_{k}-\theta^{\star}\left\|^{2}+\frac{1}{\beta}\right\| w_{k} \|^{2}$

Proposition 4. Suppose assumptions 2 and 3 holds and if we choose $\beta=\frac{8 \rho}{\lambda_{\min }(M)}$ and $\alpha_{k}=\frac{9^{2} \times 2 \delta}{8 \delta^{2}(k+2)+9^{2} \zeta}$ where $\zeta=2 \times 9^{2} c(M)^{2} \rho^{2}+32 c(M) L_{G}$. Then the mean square error $\mathbb{E}\left[\left\|z_{k}-z^{\star}\right\|^{2}\right]$ is upper bounded by:

$$
9^{2} \times 8 c(M)\left\{\frac{(8 \delta+9 \zeta)^{2} \mathbb{E}\left[\left\|z_{0}-z^{\star}\right\|^{2}\right]}{\left(8^{2} \delta^{2} k+9^{2} \zeta\right)^{2}}+\frac{8 \sigma^{2}\left(1+\left\|z^{\star}\right\|^{2}\right)}{\left(8^{2} \delta^{2} k+9^{2} \zeta\right)}\right\}
$$

Sketch of Proof (The full proof is in the appendix). The beginning of our proof relies on Du et al. [2017] which shows the linear convergence rate of deterministic primal-dual gradient method for policy evaluation. More precisely, we make use of the spectral properties of matrix $G$ shown in the appendix of this paper. The rest of the proof follows a different route exploiting the structure of our problem.

The above proposition 4 shows that the mean square error $\mathbb{E}\left[\left\|z_{k}-z^{\star}\right\|^{2}\right]$ at iteration $k$ is upper bounded by tow terms. The first bias term tells that the initial error $\mathbb{E}\left[\left\|z_{0}-z^{\star}\right\|^{2}\right]$ is forgotten at a rate $O\left(1 / k^{2}\right)$ and the constant depends on the condition number of the covariance matrix $c(M)$. The second variance term shows that noise is rejected at a rate $O(1 / k)$ and the constant depends on the variance of estimates $\sigma^{2}$ and $c(M)$. The overall convergence rate is $O(1 / k)$.

Existing stochastic saddle-point problem results: Chen et al. [2014] provides a comprehensive review of stochastic saddle-point problem. When the objective function is convex-concave, the overall convergence rate is $O(1 / \sqrt{k})$. Although several accelerated techniques could improve the dependencies on the smoothness constants of the problem in their convergence rate, the dominant term that depends on the gradient variance still decays only as $O(1 / \sqrt{k})$.
$\left.\begin{array}{|ccccc|}\hline \text { Paper } & \text { step-sizes } & \text { Projection } & \begin{array}{c}\text { Polyak } \\ \text { averaging }\end{array} & \begin{array}{c}\text { Convergence } \\ \text { rate }\end{array} \\ \hline \begin{array}{c}\text { Sutton et al. [2009c], } \\ \text { Sutton et al. [2009b] }\end{array} & \begin{array}{c}\eta_{k}=\beta \alpha_{k}, \beta>0, \\ \sum_{k} \alpha_{k}=\infty, \sum_{k} \alpha_{k}^{2}<\infty\end{array} & \text { No } & \text { No } & \begin{array}{c}\theta_{k} \rightarrow \theta^{\star} \\ \text { with probability one }\end{array} \\ \hline \text { Liu et al. [2015] } & \begin{array}{c}\text { constant step-size, } \\ \alpha_{k}=\eta_{k}\end{array} & \text { Yes } & \text { Yes } & \begin{array}{c}\text { MSPBE }\left(\bar{\theta}_{k}\right) \in O(1 / \sqrt{k}) \\ \text { with high probability }\end{array} \\ \hline \text { Wang et al. [2017] } & \begin{array}{c}\alpha_{k}=\eta_{k}, \sum_{k} \alpha_{k}=\infty, \\ \sum_{k} \alpha_{k}^{2}<\infty\end{array} & \text { Yes } & \text { Yes } & \begin{array}{c}\text { MSPBE }\left(\bar{\theta}_{k}\right) \in O\left(\frac{\sum_{k} \alpha_{k}^{2}}{\sum_{k} \alpha_{k}}\right) \\ \text { with high probability }\end{array} \\ \hline \begin{array}{c}\text { Lakshminarayanan } \\ \text { and Szepesvári [2017] }\end{array} & \begin{array}{c}\text { constant step-size, } \\ \alpha_{k}=\eta_{k}\end{array} & \text { No } & \text { Yes } & \mathbb{E}\left[\left\|\bar{\theta}_{k}-\theta^{\star}\right\|^{2}\right] \in O(1 / k) \\ \hline \text { Dalal et al. [2017] } & \begin{array}{c}\alpha_{k}=\frac{1}{k^{1-c}}, \\ \eta_{k}=\frac{1}{k(2 / 3)(1-c)} \\ \text { wher } c \in(0,1)\end{array} & \text { Yes } & \text { No } & \left\|\theta_{k}-\theta^{\star}\right\| \in O\left(k^{\left.-\frac{1}{3}+\frac{c}{3}\right)}\right. \\ \text { with high probability }\end{array}\right]$

Table 3.2: Convergence results for gradient-based TD algorithms shown in previous work [Sutton et al., 2009b,c, Liu et al., 2015, Wang et al., 2017, Lakshminarayanan and Szepesvári, 2017, Dalal et al., 2017]. $\bar{\theta}_{k}$ stand for the Polyak-average of iterates: $\bar{\theta}_{k} \triangleq \frac{\sum_{k} \alpha_{k} \theta_{k}}{\sum_{k} \alpha_{k}}$. Our algorithms achieve $O(1 / k)$ without the need for projections or Polyak averaging.

When the objective function is strongly convex-concave, Rosasco et al. [2016] and Palaniappan and Bach [2016] showed that stochastic forward-backward algorithms can achieve $O(1 / k)$ convergence rate. Algorithms in this class are feasible in practice only if their proximal mappings can be computed efficiently. In our case, our objective function is strongly concave because of the positivedefiniteness of $M$ but is otherwise not strongly convex. Because our algorithms are vanilla stochastic gradient methods, they do not rely on proximal mappings.

Singularity: If assumption 2 does not hold, the matrix $G$ is singular and either $G z+g=0$ has infinitely many solutions or it has no solution. In the case of many solutions, we could still get asymptotic convergence. In Wang and Bertsekas [2013], it was shown that under some assumptions on the null space of matrix $G$ and using a simple stabilization scheme, the iterates converge to the Drazin [Drazin, 1958] inverse solution of $G z+g=0$. However, it is not clear how extend our finite-sample analysis because the spectral analysis of the matrix $G$ [Benzi and Simoncini, 2006] in our proof assumes that the matrices $A$ and $M$ are nonsingular.

### 3.7 Related Work and Discussion

Convergent Retrace: Mahmood et al. [2017] have recently introduced the $\mathrm{ABQ}(\zeta)$ algorithm which uses an action-dependent bootstrapping parameter that leads to off-policy multi-step learning without importance sampling ratios. They also derived a gradient-based algorithm called AB-TrACE $(\lambda)$ which is related to $\operatorname{Retrace}(\lambda)$. However, the resulting updates are different from ours, as they use the two-timescale approach of Sutton et al. [2009a] as basis for their derivation. In contrast, our approach uses the saddle-point formulation, avoiding the need for double sampling. Another benefit of this formulation is that it allows us to provide a bound of the convergence rate (proposition 4) whereas Mahmood et al. [2017] is restricted to a more general two-timescale asymptotic result from Borkar and Meyn [2000]. The saddle-point formulation also provides a rich literature on acceleration methods which could be incorporated in our algorithms. Particularly in the batch setting, Du et al. [2017] recently introduced Stochastic Variance Reduction methods for state-value estimation combining GTD with SVRG Johnson and Zhang [2013] or SAGA Defazio et al. [2014]. This work could be extended easily to ours algorithms in the batch setting.

Existing Convergence Rates: Our convergence rate result 4 can apply to GTD/GTD2 algorithms. Recall that GTD/GTD2 are off-policy algorithms designed to estimate the state-value function using temporal difference $\mathrm{TD}(0)$ return while our algorithms compute the action-value function using Retrace and Tree BACKUP returns. In both GTD and GTD2, the quantities $\hat{A}_{k}$ and $\hat{b}_{k}$ involved in their updates are the same and equal to $\hat{A}_{k}=\phi\left(s_{k}\right)\left(\gamma \phi\left(s_{k+1}\right)-\phi\left(s_{k}\right)\right)^{\top}$, $\hat{b}_{k}=r\left(s_{k}, a_{k}\right) \phi\left(s_{k}\right)$ while the matrix $\hat{M}_{k}$ is equal to $\phi\left(s_{k}\right) \phi\left(s_{k}\right)^{\top}$ for GTD2 and to identity matrix for GTD.
The table 3.2 show in chronological order the convergence rates established in the literature of Reinforcement learning. GTD was first introduced in Sutton et al. [2009c] and its variant GTD2 was introduced later in Sutton et al. [2009b]. Both papers established the asymptotic convergence with RobbinsMonro step-sizes. Later, Liu et al. [2015] provided the first sample complexity by reformulating GTD/GTD2 as an instance of mirror stochastic approximation [Nemirovski et al., 2009]. Liu et al. [2015] showed that with high probability, $\operatorname{MSPBE}\left(\bar{\theta}_{k}\right) \in O(1 / \sqrt{k})$ where $\bar{\theta}_{k} \triangleq \frac{\sum_{k} \alpha_{k} \theta_{k}}{\sum_{k} \alpha_{k}}$. However, they studied an alternated version of GTD/GTD2 as they added a projection step into bounded convex set and Polyak-averaging of iterates. Wang et al. [2017] studied also the same version as Liu et al. [2015] but for the case of Markov noise case instead of the i.i.d assumptions. They prove that with high probability $\operatorname{MSPBE}\left(\bar{\theta}_{k}\right) \in O\left(\frac{\sum_{k} \alpha_{k}^{2}}{\sum_{k} \alpha_{k}}\right)$ when the step-size sequence satisfies $\sum_{k} \alpha_{k}=\infty, \frac{\sum_{k} \alpha_{k}^{2}}{\sum_{k} \alpha_{k}}<\infty$. The optimal
rate achieved in this setup is then $O(1 / \sqrt{k})$. Recently, Lakshminarayanan and Szepesvári [2017] improved on the existing results by showing for the first time that $\mathbb{E}\left[\left\|\bar{\theta}_{k}-\theta^{\star}\right\|^{2}\right] \in O(1 / k)$ without projection step. However, the result still consider the Polyak-average of iterates. Moreover, the constants in their bound depend on the data distribution that are difficult to relate to the problem-specific constants, such as those present in our bound 4. Finally, Dalal et al. [2017] studied sparsily projected version of GTD/GTD2 and they showed that for step-sizes $\alpha_{k}=\frac{1}{k^{1-c}}, \eta_{k}=\frac{1}{k^{(2 / 3)(1-c)}}$ where $c \in(0,1),\left\|\theta_{k}-\theta^{\star}\right\| \in O\left(k^{-\frac{1}{3}+\frac{c}{3}}\right)$ with high probability. The projection is called sparse as they project only on iterations which are powers of 2 .
Our work is the first to provide a finite-sample complexity analysis of GTD/GTD2 in its original setting, i.e without assumption a projection step or Polyak-averaging and with diminishing step-sizes.

### 3.8 EXPERIMENTAL RESULTS

### 3.8.1 Evidence of instability in practice

To validate our theoretical results about instability, we implemented $\operatorname{TB}(\lambda)$, $\operatorname{RETRACE}(\lambda)$ and compared them against their gradient-based counterparts $\operatorname{GTB}(\lambda)$ and $\operatorname{GRETRACE}(\lambda)$ derived in this paper. The first one is the 2-states counterexample that we detailed in the third section and the second is the 7 -states versions of Baird's counterexample [Baird et al., 1995]. Figures 3.2 and 3.3 show the MSBPE (averaged over 20 runs) as a function of the number of iterations. We can see that our gradient algorithms converge in these two counterexamples whereas $\operatorname{TB}(\lambda)$ and $\operatorname{RETRACE}(\lambda)$ diverge.


Figure 3.2: Baird's counterexample. The combination of linear function approximation with TB and RETRACE leads to divergence (left panel) while the proposed gradient extensions GTB and GRETRACE converge (right panel).


Figure 3.3: In the 2-states counterexample of section 3.4 showing that the gradient-based TB and RETRACE converge while TB and Retrace diverge.

### 3.8.2 Comparison with existing methods

We also compared $\operatorname{GTB}(\lambda)$ and $\operatorname{GRETRACE}(\lambda)$ with two recent state-of-theart convergent off-policy algorithms for action-value estimation and function approximation: $\mathrm{GQ}(\lambda)$ [Maei, 2011] and $\operatorname{AB}-\operatorname{TrACE}(\lambda)$ [Mahmood et al., 2017]. As in Mahmood et al. [2017], we also consider a policy evaluation task in the Mountain Car domain. In order to better understand the variance inherent to each method, we designed the target policy and behavior policy in such a way that the importance sampling ratios can be as large as 30 . We chose to describe state-action pairs by a 96 -dimensional vector of features derived by tile coding [Sutton and Barto, 1998]. We ran each algorithm over all possible combinations of step-size values $\left(\alpha_{k}, \eta_{k}\right) \in[0.001,0.005,0.01,0.05,0.1]^{2}$ for 2000 episodes and reported their normalized mean squared errors (NMSE):

$$
\operatorname{NMSE}(\theta)=\frac{\left\|\Phi \theta-Q^{\pi}\right\|_{\Xi}^{2}}{\left\|Q^{\pi}\right\|_{\Xi}^{2}}
$$

where $Q^{\pi}$ is estimated by simulating the target policy and averaging the discounted cumulative rewards overs trajectories. As AB-TRACE $(\lambda)$ and $\operatorname{GRETRACE}(\lambda)$ share both the same operator, we can evaluate them using the empirical MSPBE $=\frac{1}{2}\|\hat{A} \theta+\hat{b}\|_{\hat{M}^{-1}}^{2}$ where $\hat{A}, \hat{b}$ and $\hat{M}$ are Monte-Carlo estimates obtained by averaging $\hat{A}_{k}, \hat{b}_{k}$ and $\hat{M}_{k}$ defined in proposition 3 over 10000 episodes.
Figure 3.5 shows that the best empirical MSPBE achieved by $\operatorname{AB-TrACE}(\lambda)$ and $\operatorname{GRETRACE}(\lambda)$ are almost identical across value of $\lambda$. This result is consistent with the fact that they both minimize the MSPBE objective function. However, significant differences can be observed when computing the 5th percentiles of NMSE (over all possible combination of step-size values) for different values of


Figure 3.4: Each curves shows the 5th percentile of NMSE (over all possible combination of step-size values) achieved by each algorithm for different values of $\lambda$.
$\lambda$ in Figure 3.4. When $\lambda$ increases, the NMSE of $\operatorname{GQ}(\lambda)$ increases sharply due to increased influence of importance sampling ratios. This clearly demonstrate the variance issues of $\mathrm{GQ}(\lambda)$ in contrast with the other methods based on the Tree Backup and Retrace returns (that are not using importance ratios). For intermediate values of $\lambda, \operatorname{AB}-\operatorname{TrACE}(\lambda)$ performs better but its performance is matched by $\operatorname{GRETRACE}(\lambda)$ and $\operatorname{TB}(\lambda)$ for small and very large values of $\lambda$. In fact, $\operatorname{AB-TRACE}(\lambda)$ updates the function parameters $\theta$ as follows:

$$
\theta_{k+1}=\theta_{k}-\alpha_{k}\left(\delta_{k} e_{k}-\Delta_{k}\right)
$$

where $\Delta_{k} \triangleq \gamma w_{k}^{\top} e_{k}\left(\mathbb{E}_{\pi} \phi\left(s_{k+1},.\right)-\lambda \sum_{a} \kappa\left(s_{k}, a\right) \mu\left(a \mid s_{k}\right) \phi\left(s_{k}, a\right)\right)$ is a gradient correction term. When the instability is not an issue, the correction term could be very small and the update of $\theta$ would be essentially $\theta_{k+1} \sim \theta_{k}-\alpha_{k} \delta_{k} e_{k}$ so that $\theta_{k+1}$ follows the semi-gradient of the mean squared error $\left\|\Phi \theta-G_{k}^{\lambda}\right\|_{\square}^{2}$.
To better understand the errors of each algorithm and their robustness to stepsize values, we propose the box plots shown in Figure 3.6. Each box plot shows the distribution of NMSE obtained by each algorithm for different values of $\lambda$. NMSE distributions are computed over all possible combinations of step-size values. $\operatorname{GTB}(\lambda)$ has the smallest variance as it scaled its return by the target probabilities which makes it conservative in its update even with large step-size values. $\operatorname{GRETRACE}(\lambda)$ tends to more more efficient than $\operatorname{GTB}(\lambda)$ since it could benefit from full returns. The latter observation agrees with the tabular case of Tree Backup and Retrace [Munos et al., 2016]. Finally, we observe that AB$\operatorname{Trace}(\lambda)$ has lower error, but at the cost of increased variance with respect to


Figure 3.5: Comparison between the best empirical MSPBE obtained by each algorithm for different values of $\lambda$. Only $\operatorname{GREtRACE}(\lambda)$ and $\operatorname{AB}-\operatorname{TRACE}(\lambda)$ are showed here because the other algorithms do not have the same operators and hence not the same MSPBE. Note that MSPBEs depend on $\lambda$. Thus, MSPBEs are not directly comparable across different values of $\lambda$. Both $\operatorname{GRETRACE}(\lambda)$ and $\operatorname{AB}-\operatorname{TRACE}(\lambda)$ have very similar performances. AB-TRACE $(\lambda)$ performs slightly better.
step-size values.

### 3.9 CONCLUSION

Our analysis highlighted for the first time the difficulties of combining the Tree Backup and Retrace algorithms with function approximation. We addressed these issues by formulating gradient-based algorithm versions of these algorithms which minimize the mean-square projected Bellman error. Using a saddle-point formulation, we were also able to provide convergence guarantees and characterize the convergence rate of our algorithms GTB and GRETRACE. We also developed a novel analysis method which allowed us to establish a $O(1 / k)$ convergence rate without having to use Polyak averaging or projections. Furthermore, our proof technique is general enough that we were able to apply it to the existing GTD and GTD2 algorithms. Our experiments finally suggest that the proposed $\operatorname{GTB}(\lambda)$ and $\operatorname{GRETRACE}(\lambda)$ are robust to step-size selection and have less variance than both $\mathrm{GQ}(\lambda)$ [Maei, 2011] and AB-Trace $(\lambda)$ [Mahmood et al., 2017].


Figure 3.6: Comparison of empirical performance of $\operatorname{GQ}(\lambda), \operatorname{AB}-\operatorname{TrACE}(\lambda), \operatorname{GRETRACE}(\lambda)$ and GTB $(\lambda)$ on an off-policy evaluation task in Mountain Car domain. Each box plot shows the distribution of the NMSE achieved by each algorithm after 2000 episodes for different values of $\lambda$. NMSE distributions are computed over all the possible combinations of step-size values $\left(\alpha_{k}, \eta_{k}\right) \in[0.001,0.005,0.01,0.05,0.1]^{2}$.

## Stable Policy Optimization via Off-Policy Divergence Regularization

### 4.1 Prologue to the Contribution

### 4.1.1 Article Details

This chapter is based on the article «Stable Policy Optimization via Off-Policy Divergence Regularization» [Touati et al., 2020c], jointly led with Amy Zhang and in collaboration with Joelle Pineau and Pascal Vincent. This paper was accepted as oral at UAI 2020. I share the first authorship with Amy Zhang. I was at the origin of the project's idea and the derivation of the new algorithm. I helped with the implementation of the algorithm and Amy performed the large-scale experiments of the paper.

### 4.1.2 Context

In the previous chapter, we saw how the discrepancy between the policy that generated the data to be learnt from (the behavior policy) and the policy being evaluated (the target policy) can lead to undesirable behavior. This behavior can be arbitrarily more complex in the control setting where we aim at learning a near-optimal policy. Some policy optimization algorithms control this distribution shift by ensuring that the improved policy stays in the vicinity of the current policy in term of their action probabilities.

In this work, we want to improve how these algorithms leverage off-policy data and stabilize the policy learning. Especially, we believe that the regularization based on immediate actions is not enough and that we should reason about long-term future state-action distribution. This is well supported when we look at formal analysis of offline RL algorithms such as approximate Value Iteration and Fitted Q-iteration [Munos, 2005, Munos and Szepesvári, 2008, Chen and Jiang, 2019, Touati and Vincent, 2020b]. These algorithms are concerned with finding a near-optimal policy given a fixed dataset of transitions and their analyses rely on strong assumptions about the distribution of batch data. Especially, they assume that the ratio between the induced state-action distribution of any non-stationary policy and the state-action distribution of the batch data is upper bounded by a constant, referred to as the concentrability coefficient. Unfortunately, this coefficient may be arbitrarily large or infinite, due the large policy class, which makes the error bounds of these analysis vacuous.

Even though our setting is easier than the fully batch setting as we allow ourselves to collect online data, we believe that, to use better off-policy data, we need to bias algorithmically the policy class in the improvement step to the one that is well-supported by the current off-policy data distribution.

### 4.1.3 Paper Abstract

Trust Region Policy Optimization (TRPO) and Proximal Policy Optimization (PPO) are among the most successful policy gradient approaches in deep reinforcement learning (RL). While these methods achieve state-of-the-art performance across a wide range of challenging tasks, there is room for improvement in the stabilization of the policy learning and how the off-policy data are used. In this paper we revisit the theoretical foundations of these algorithms and propose a new algorithm which stabilizes the policy improvement through a proximity term that constrains the discounted state-action visitation distribution induced by consecutive policies to be close to one another. This proximity term, expressed in terms of the divergence between the visitation distributions, is learned in an off-policy and adversarial manner. We empirically show that our proposed method can have a beneficial effect on stability and improve final performance in benchmark high-dimensional control tasks. Our code is available at https://github.com/facebookresearch/ppo-dice.

### 4.1.4 Recent Developments

In the fully batch setting, the work of [Liu et al., 2020] has the same underlying motivation as our work i.e constraining the learned policy not only in term of action distribution, like in previous batch RL methods, but rather in term of its coverage of the provided state-action distribution. They implement this idea as follows: given a batch of data, they learn a state-action density function, then, in both policy evaluation and improvement steps, they filter out state-action pairs whose estimated densities are below than some threshold. Theoretically, they show that they can find the approximately best policy with sufficient support in the batch data, without requiring a priori strong concentrability assumptions.

### 4.2 INTRODUCTION

Combined with deep neural networks as function approximators, policy gradient methods have enjoyed many empirical successes on RL problems such as video games [Mnih et al., 2016] and robotics [Levine et al., 2016]. Their recent success can be attributed to their ability to scale gracefully to high dimensional state-action spaces and complex dynamics.

The main idea behind policy gradient methods is to parametrize the policy and perform stochastic gradient ascent on the discounted cumulative reward directly [Sutton et al., 2000]. To estimate the gradient, we sample trajectories from the distribution induced by the policy. Due to the stochasticity of both policy and environment, variance of the gradient estimation can be very large, and lead to significant policy degradation.

Instead of directly optimizing the cumulative rewards, which can be challenging due to large variance, some approaches [Kakade and Langford, 2002, Azar et al., 2012, Pirotta et al., 2013, Schulman et al., 2015] propose to optimize a surrogate objective that can provide local improvements to the current policy at each iteration. The idea is that the advantage function of a policy $\pi$ can produce a good estimate of the performance of another policy $\pi^{\prime}$ when the two policies give rise to similar state visitation distributions. Therefore, these approaches explicitly control the state visitation distribution shift between successive policies.

However, controlling the state visitation distribution shift requires measuring it, which is non-trivial. Direct methods are prohibitively expensive. Therefore, in order to make the optimization tractable, the aforementioned methods rely on constraining action probabilities by mixing policies [Kakade and Langford, 2002, Pirotta et al., 2013], introducing trust regions [Schulman et al., 2015, Achiam et al., 2017] or clipping the surrogate objective [Schulman et al., 2017, Wang et al., 2019b].

Our key motivation in this work is that constraining the probabilities of the immediate future actions might not be enough to ensure that the surrogate objective is still a valid estimate of the performance of the next policy and consequently might lead to instability and premature convergence. Instead, we argue that we should reason about the long-term effect of the policies on the distribution of the future states.

In particular, we directly consider the divergence between state-action visitation distributions induced by successive policies and use it as a regularization term added to the surrogate objective. This regularization term is itself optimized in an adversarial and off-policy manner by leveraging recent advances in off-policy policy evaluation [Nachum et al., 2019a] and off-policy imitation learning [Kostrikov et al., 2019]. We incorporate these ideas in the PPO algorithm in order to ensure safer policy learning and better reuse of off-policy data. We call our proposed method PPO-DICE.

The present paper is organized as follows: after reviewing conservative approaches for policy learning, we provide theoretical insights motivating our method. We explain how off-policy adversarial formulation can be derived to optimize the regularization term. We then present the algorithmic details of our proposed method. Finally, we show empirical evidences of the benefits of PPODICE as well as ablation studies.

### 4.3 Conservative Update Approaches

We consider a discounted $\operatorname{MDP}(\mathcal{S}, \mathcal{A}, \gamma, P, r, \rho)$ with continuous state and action spaces $\mathcal{S}$ and $\mathcal{A}$. We assume in this chapter that all probability distributions are absolutely continuous with respect to the Lebesgue measure and we confound probability distributions with their probability density functions, by abuse of notation.

The goal of the agent is to find a policy $\pi$ that maximizes the expected value from under the initial state distribution $\rho$ :

$$
\max _{\pi} J(\pi) \triangleq(1-\gamma) \mathbb{E}_{s \sim \rho}\left[V^{\pi}(s)\right]
$$

We define the normalized discounted state visitation density $d_{\rho}^{\pi}$ induced by a policy $\pi$ :

$$
d_{\rho}^{\pi}(s) \triangleq(1-\gamma) \sum_{t=0}^{\infty} \gamma^{t} \operatorname{Pr}\left(s_{t}=s \mid s_{0} \sim \rho, \pi\right)
$$

where $\operatorname{Pr}\left(s_{t}=s \mid s_{0} \sim \rho, \pi\right)$ is the probability density that $s_{t}=s$, after we execute $\pi$ for $t$ steps, starting from initial state $s_{0}$ distributed according to $\rho$. Similarly, we define the discounted state-action visitation density $\mu_{\rho}^{\pi}(s, a)$ of policy $\pi$

$$
\mu_{\rho}^{\pi}(s, a) \triangleq(1-\gamma) \sum_{t=0}^{\infty} \gamma^{t} \operatorname{Pr}\left(s_{t}=s, a_{t}=a \mid s_{0} \sim \rho, \pi\right)
$$

It is known [Puterman, 1990] that $\mu_{\rho}^{\pi}(s, a)=d_{\rho}^{\pi}(s) \cdot \pi(a \mid s)$ and that $\mu^{\pi}$ is characterized via: $\forall\left(s^{\prime}, a^{\prime}\right) \in \mathcal{S} \times \mathcal{A}$

$$
\mu_{\rho}^{\pi}\left(s^{\prime}, a^{\prime}\right)=(1-\gamma) \rho\left(s^{\prime}\right) \pi\left(a^{\prime} \mid s^{\prime}\right)+\gamma \int \pi\left(a^{\prime} \mid s^{\prime}\right) P\left(s^{\prime} \mid s, a\right) \mu_{\rho}^{\pi}(s, a) d s d a
$$

Most policy training approaches in RL can be understood as updating a current policy $\pi$ to a new improved policy $\pi^{\prime}$ based on the advantage function $A^{\pi} \triangleq Q^{\pi}(s, a)-V^{\pi}(s)$ or an estimate $\hat{A}$ of it. We review here some popular approaches that implement conservative updates in order to stabilize policy training.

First, let us state a key lemma from the seminal work of Kakade and Langford [2002] that relates the performance difference between two policies to the advantage function.

Lemma 1 (The performance difference lemma [Kakade and Langford, 2002]). For all policies $\pi$ and $\pi^{\prime}$,

$$
\begin{equation*}
J\left(\pi^{\prime}\right)=J(\pi)+\mathbb{E}_{s \sim d_{\rho}^{\pi^{\prime}}} \mathbb{E}_{a \sim \pi^{\prime}(. \mid s)}\left[A^{\pi}(s, a)\right] \tag{4.3.1}
\end{equation*}
$$

This lemma implies that maximizing Equation (4.3.1) will yield a new policy $\pi^{\prime}$ with guaranteed performance improvement over a given policy $\pi$. Unfortunately, a naive direct application of this procedure would be prohibitively expensive since it requires estimating $d_{\rho}^{\pi^{\prime}}$ for all $\pi^{\prime}$ candidates. To address this issue, Conservative Policy Iteration (CPI) [Kakade and Langford, 2002] optimizes a surrogate objective defined based on current policy $\pi^{i}$ at each iteration $i$,

$$
\begin{equation*}
L_{\pi_{i}}\left(\pi^{\prime}\right)=J\left(\pi_{i}\right)+\mathbb{E}_{s \sim d_{\rho}^{\pi_{i}}} \mathbb{E}_{a \sim \pi^{\prime}(. \mid s)}\left[A^{\pi_{i}}(s, a)\right] \tag{4.3.2}
\end{equation*}
$$

by ignoring changes in state visitation distribution due to changes in the policy. Then, CPI returns the stochastic mixture $\pi_{i+1}=\alpha_{i} \pi_{i}^{+}+\left(1-\alpha_{i}\right) \pi_{i}$ where $\pi_{i}^{+}=$ $\arg \max _{\pi^{\prime}} L_{\pi_{i}}\left(\pi^{\prime}\right)$ is the greedy policy and $\alpha_{i} \in[0,1]$ is tuned to guarantee a monotonically increasing sequence of policies.

Inspired by CPI, the Trust Region Policy Optimization algorithm (TRPO) [Schulman et al., 2015] extends the policy improvement step to any general stochastic policy rather than just mixture policies. TRPO maximizes the same surrogate objective as CPI subject to a Kullback-Leibler (KL) divergence constraint that ensures the next policy $\pi_{i+1}$ stays within $\delta$-neighborhood of the current policy $\pi_{i}$ :

$$
\begin{align*}
& \pi_{i+1}=\arg \max _{\pi^{\prime}} L_{\pi_{i}}\left(\pi^{\prime}\right)  \tag{4.3.3}\\
& \text { s.t } \quad \mathbb{E}_{s \sim d_{\rho}}^{\pi_{i}}\left[D_{\mathrm{KL}}\left(\pi^{\prime}(\cdot \mid s) \| \pi_{i}(\cdot \mid s)\right)\right] \leq \delta,
\end{align*}
$$

where $D_{\text {KL }}$ is the Kullback-Leibler divergence. In practise, TRPO considers a differentiable parameterized policy $\left\{\pi_{\theta}, \theta \in \Theta\right\}$ and solves the constrained problem (4.3.3) in parameter space $\Theta$. In particular, the step direction is estimated with conjugate gradients, which requires the computation of multiple Hessian-vector products. Therefore, this step can be computationally heavy.

To address this computational bottleneck, Proximal Policy Optimization (PPO) [Schulman et al., 2017] proposes replacing the KL divergence constrained objective (4.3.3) of TRPO by clipping the objective function directly as:

$$
\begin{gather*}
L_{\pi_{i}}^{\operatorname{clip}}\left(\pi^{\prime}\right)=\mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi_{i}}}\left[\operatorname { m i n } \left\{A^{\pi_{i}}(s, a) \cdot \kappa_{\pi^{\prime} / \pi_{i}}(s, a),\right.\right. \\
\left.\left.A^{\pi_{i}}(s, a) \cdot \operatorname{clip}\left(\kappa_{\pi^{\prime} / \pi_{i}}(s, a), 1-\epsilon, 1+\epsilon\right)\right\}\right] \tag{4.3.4}
\end{gather*}
$$

where $\epsilon>0$ and $\kappa_{\pi^{\prime} / \pi_{i}}(s, a)=\frac{\pi^{\prime}(s, a)}{\pi_{i}(s, a)}$ is the importance sampling ratio.

### 4.4 Theoretical Insights

In this section, we present the theoretical motivation of our proposed method.

At a high level, algorithms CPI, TRPO, and PPO follow similar policy update schemes. They optimize some surrogate performance objective $\left(L_{\pi_{i}}\left(\pi^{\prime}\right)\right.$ for CPI and TRPO and $L_{\pi}^{\text {clip }}\left(\pi^{\prime}\right)$ for PPO) while ensuring that the new policy $\pi_{i+1}$ stays in the vicinity of the current policy $\pi_{i}$. The vicinity requirement is implemented in different ways:

- CPI computes a sequence of stochastic policies that are mixtures between consecutive greedy policies.
- TRPO imposes a constraint on the KL divergence between old policy and new one $\left(\mathbb{E}_{s \sim d_{\rho}^{\pi_{i}}}\left[D_{\mathrm{KL}}\left(\pi^{\prime}(\cdot \mid s) \| \pi_{i}(\cdot \mid s)\right)\right] \leq \delta\right)$.
- PPO directly clips the objective function based on the value of the importance sampling ratio $\kappa_{\pi^{\prime} / \pi_{i}}$ between the old policy and new one.

Such conservative updates are critical for the stability of the policy optimization. In fact, the surrogate objective $L_{\pi_{i}}\left(\pi^{\prime}\right)$ (or its clipped version) is valid only in the neighbourhood of the current policy $\pi_{i}$, i.e, when $\pi^{\prime}$ and $\pi_{i}$ visit all the states with similar probabilities. The following lemma more precisely formalizes this ${ }^{1}$ :
Lemma 2. For all policies $\pi$ and $\pi^{\prime}$,

$$
\begin{align*}
J\left(\pi^{\prime}\right) & \geq L_{\pi}\left(\pi^{\prime}\right)-\epsilon^{\pi} D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi}\right)  \tag{4.4.1}\\
& \geq L_{\pi}^{\text {clip }}\left(\pi^{\prime}\right)-\epsilon^{\pi} D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi}\right)
\end{align*}
$$

where $\epsilon^{\pi}=\max _{s \in \mathcal{S}}\left|\mathbb{E}_{a \sim \pi^{\prime}(\cdot \mid s)}\left[A^{\pi}(s, a)\right]\right|$ and $D_{\mathrm{TV}}$ is the total variation distance.
The proof is provided in appendix for completeness. Lemma 2 states that $L_{\pi}\left(\pi^{\prime}\right)$ (or $L_{\pi}^{\text {clip }}\left(\pi^{\prime}\right)$ ) is a sensible lower bound to $J\left(\pi^{\prime}\right)$ as long as $\pi$ and $\pi^{\prime}$ are close in terms of total variation distance between their corresponding state visitation distributions $d_{\rho}^{\pi^{\prime}}$ and $d_{\rho}^{\pi}$. However, the aforementioned approaches enforce closeness of $\pi^{\prime}$ and $\pi$ in terms of their action probabilities rather than their state visitation distributions. This can be justified by the following inequality [Achiam et al., 2017]:

$$
\begin{equation*}
D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi}\right) \leq \frac{2 \gamma}{1-\gamma} \mathbb{E}_{s \sim d_{\rho}^{\pi}}\left[D_{\mathrm{TV}}\left(\pi^{\prime}(. \mid s) \| \pi(. \mid s)\right)\right] \tag{4.4.2}
\end{equation*}
$$

Plugging the last inequality (4.4.2) into (4.4.1) leads to the following lower bound:

$$
\begin{equation*}
J\left(\pi^{\prime}\right) \geq L_{\pi}\left(\pi^{\prime}\right)-\frac{2 \gamma \epsilon^{\pi}}{1-\gamma} \mathbb{E}_{s \sim d_{\rho}^{\pi}}\left[D_{\mathrm{TV}}\left(\pi^{\prime}(. \mid s) \| \pi(. \mid s)\right)\right] \tag{4.4.3}
\end{equation*}
$$

[^2]The obtained lower bound (4.4.3) is, however, clearly looser than the one in inequality (4.4.2). Lower bound (4.4.3) suffers from an additional multiplicative factor $\frac{1}{1-\gamma}$, which is the effective planning horizon. It is essentially due to the fact that we are characterizing a long-horizon quantity, such as the state visitation distribution $d_{\rho}^{\pi}(s)$, by a one-step quantity, such as the action probabilities $\pi(\cdot \mid s)$. Therefore, algorithms that rely solely on action probabilities to define closeness between policies should be expected to suffer from instability and premature convergence in long-horizon problems.

Furthermore, in the exact case if we take at iteration $i$, $\pi_{i+1} \leftarrow$ $\arg \max _{\pi^{\prime}} L_{\pi_{i}}\left(\pi^{\prime}\right)-\epsilon^{\pi_{i}} D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi_{i}}\right)$, then

$$
\begin{array}{rlr}
J\left(\pi_{i+1}\right) & \geq L_{\pi_{i}}\left(\pi_{i+1}\right)-\epsilon^{\pi_{i}} D_{\mathrm{TV}}\left(d_{\rho}^{\pi_{i+1}} \| d_{\rho}^{\pi_{i}}\right) \\
& \geq L_{\pi_{i}}\left(\pi_{i}\right) & \left.\quad \text { (by optimality of } \pi_{i+1}\right) \\
& =J\left(\pi_{i}\right) &
\end{array}
$$

Therefore, this provides a monotonic policy improvement, while TRPO suffers from a performance degradation that depends on the level of the trust region $\delta$ (see Proposition 1 in Achiam et al. [2017]).

It follows from our discussion that $D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi}\right)$ is a more natural proximity term to ensure safer and more stable policy updates. Previous approaches excluded using this term because we don't have access to $d_{\rho}^{\pi^{\prime}}$, which would require executing $\pi^{\prime}$ in the environment. In the next section, we show how we can leverage recent advances in off-policy policy evaluation to address this issue.

### 4.5 Off-policy Formulation Of Divergences

In this section, we explain how divergences between state-visitation distributions can be approximated. This is done by leveraging ideas from recent works on off-policy learning [Nachum et al., 2019a, Kostrikov et al., 2019].

Consider two different policies $\pi$ and $\pi^{\prime}$. Suppose that we have access to state-action samples generated by executing the policy $\pi$ in the environment, i.e, $(s, a) \sim \mu_{\rho}^{\pi}$. As motivated by the last section, we aim to estimate $D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi}\right)$ without requiring on-policy data from $\pi^{\prime}$. Note that in order to avoid using importance sampling ratios, it is more convenient to estimate $D_{\mathrm{TV}}\left(\mu_{\rho}^{\pi^{\prime}} \| \mu_{\rho}^{\pi}\right)$, i.e, the total divergence between state-action visitation distributions rather than the divergence between state visitation distributions. This is still a reasonable choice
as $D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi}\right)$ is upper bounded by $D_{\mathrm{TV}}\left(\mu_{\rho}^{\pi^{\prime}} \| \mu_{\rho}^{\pi}\right)$ as shown below:

$$
\begin{aligned}
D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi}\right) & =\int_{s}\left|\left(d_{\rho}^{\pi^{\prime}}-d_{\rho}^{\pi}\right)(s)\right| d s \\
& =\int_{s}\left|\int_{a}\left(\mu_{\rho}^{\pi^{\prime}}-\mu_{\rho}^{\pi}\right)(s, a) d a\right| d s \\
& \leq \int_{s} \int_{a}\left|\left(\mu_{\rho}^{\pi^{\prime}}-\mu_{\rho}^{\pi}\right)(s, a)\right| d a d s \\
& =D_{\mathrm{TV}}\left(\mu_{\rho}^{\pi^{\prime}} \| \mu_{\rho}^{\pi}\right)
\end{aligned}
$$

The total variation distance belongs to a broad class of divergences known as $\phi$-divergences [Sriperumbudur et al., 2009]. A $\phi$-divergence is defined as,

$$
\begin{equation*}
D_{\phi}\left(\mu_{\rho}^{\pi^{\prime}} \| \mu_{\rho}^{\pi}\right)=\mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi^{\prime}}}\left[\phi\left(\frac{\mu_{\rho}^{\pi}(s, a)}{\mu_{\rho}^{\pi^{\prime}}(s, a)}\right)\right] \tag{4.5.1}
\end{equation*}
$$

where $\phi:[0, \infty) \rightarrow \mathbb{R}$ is a convex, lower-semicontinuous function and $\phi(1)=0$. Well-known divergences can be obtained by appropriately choosing $\phi$. These include the KL divergence $(\phi(t)=t \log (t))$, total variation distance $(\phi(t)=$ $|t-1|)$, $\chi^{2}$-divergence $\left(\phi(t)=(t-1)^{2}\right)$, etc. Working with the form of $\phi$ divergence given in Equation (4.5.1) requires access to analytic expressions of both $\mu_{\rho}^{\pi}$ and $\mu_{\rho}^{\pi}$ as well as the ability to sample from $\mu_{\rho}^{\pi^{\prime}}$. We have none of these in our problem of interest. To bypass these difficulties, we turn to the alternative variational representation of $\phi$-divergences [Nguyen et al., 2009, Huang et al., 2017b] as

$$
\begin{equation*}
D_{\phi}\left(\mu_{\rho}^{\pi^{\prime}} \| \mu_{\rho}^{\pi}\right)=\sup _{f: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}}\left[\mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi^{\prime}}}[f(s, a)]-\mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi}}\left[\phi^{\star} \circ f(s, a)\right]\right] \tag{4.5.2}
\end{equation*}
$$

where $\phi^{\star}(t)=\sup _{u \in \mathbb{R}}\{t u-\phi(u)\}$ is the convex conjugate of $\phi$. The variational form in Equation (4.5.2) still requires sampling from $\mu_{\rho}^{\pi^{\prime}}$, which we cannot do. To address this issue, we use a clever change of variable trick introduced by Nachum et al. [2019a]. Define $g: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ as the fixed point of the following Bellman equation,

$$
\begin{equation*}
g(s, a)=f(s, a)+\gamma P^{\pi^{\prime}} g(s, a) \tag{4.5.3}
\end{equation*}
$$

where $\mathbb{P}^{\pi^{\prime}}$ is the transition operator induced by $\pi^{\prime}$, defined as $\mathbb{P}^{\pi^{\prime}} g(s, a)=$ $\int \pi^{\prime}\left(a^{\prime} \mid s^{\prime}\right) \mathbb{P}\left(s^{\prime} \mid s, a\right) g\left(s^{\prime}, a^{\prime}\right)$. $g$ may be interpreted as the action-value function of the policy $\pi^{\prime}$ in a modified MDP which shares the same transition model $\mathbb{P}$ as the original MDP, but has $f$ as the reward function instead of $r$. Applying the change of variable (4.5.3) to (4.5.2) and after some algebraic manipulation as
done in Nachum et al. [2019a], we obtain

$$
\begin{gather*}
D_{\phi}\left(\mu_{\rho}^{\pi^{\prime}} \| \mu_{\rho}^{\pi}\right)=\sup _{g: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}}\left[(1-\gamma) \mathbb{E}_{s \sim \rho, a \sim \pi^{\prime}}[g(s, a)]-\right. \\
\left.\mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi}}\left[\phi^{\star}\left(\left(g-\gamma P^{\pi^{\prime}} g\right)(s, a)\right)\right]\right] \tag{4.5.4}
\end{gather*}
$$

Thanks to the change of variable, the first expectation over $\mu_{\rho}^{\pi^{\prime}}$ in (4.5.2) is converted to an expectation over the initial distribution and the policy i.e $s \sim$ $\rho(\cdot), a \sim \pi^{\prime}(\cdot \mid s)$. Therefore, this new form of the $\phi$-divergence in (4.5.4) is completely off-policy and can be estimated using only samples from the policy $\pi$.

Other possible divergence representations: Using the variational representation of $\phi$-divergences was a key step in the derivation of Equation (4.5.4). But in fact any representation that admits a linear term with respect to $\mu_{\rho}^{\pi^{\prime}}$ (i.e $\left.\mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi^{\prime}}}[f(s, a)]\right)$ would work as well. For example, one can use the the Donkser-Varadhan representation [Donsker and Varadhan, 1983] to alternatively express the KL divergence as:

$$
\begin{equation*}
D_{\phi}\left(\mu_{\rho}^{\pi^{\prime}} \| \mu_{\rho}^{\pi}\right)=\sup _{f: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}}\left[\mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi^{\prime}}}[f(s, a)]-\log \left(\mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi}}[\exp (f(s, a))]\right)\right] \tag{4.5.5}
\end{equation*}
$$

The log-expected-exp in this equation makes the Donkser-Varadhan representation (4.5.5) more numerically stable than the variational one (4.5.4) when working with KL divergences. Because of its genericity for $\phi$-divergences, we base the remainder of our exposition on (4.5.4). But it is straightforward to adapt the approach and algorithm to using (4.5.5) for better numerical stability when working with KL divergences specifically. Thus, in practice we will use the latter in our experiments with KL-based regularization, but not in the ones with $\chi^{2}$-based regularization.

### 4.6 A Practical Algorithm Using Adversarial Divergence

We now turn these insights into a practical algorithm. The lower bounds in lemma 2 , suggest using a regularized PPO objective ${ }^{2}$ :

[^3]```
Algorithm 4 PPO-DICE
    Initialisation: random initialize parameters \(\theta_{1}\) (policy), \(\psi_{1}\) (discriminator)
    and \(\omega_{1}\) (value function).
    for \(i=1, \ldots\) do
        Generate a batch of \(M\) rollouts \(\left\{s_{1}^{(j)}, a_{1}^{(j)}, r_{1}^{(j)}, s_{1}^{(j)}, \ldots, s_{T}^{(j)}, a_{T}^{(j)}, r_{T}^{(j)}, s_{T+1}^{(j)}\right\}_{j=1}^{M}\)
        by executing policy \(\pi_{\theta_{i}}\) in the environment for \(T\) steps.
        Estimate Advantage function: \(\hat{A}\left(s_{t}^{(j)}, a_{t}^{(j)}\right)=\sum_{t=1}^{T}(\gamma \lambda)^{t-1}\left(r_{t}^{(j)}+\right.\)
        \(\left.\gamma V_{\omega_{i}}\left(s_{t+1}^{(j)}\right)-V_{\omega_{i}}\left(s_{t}^{(j)}\right)\right)\)
    Compute target value \(y_{t}^{(j)}=r_{t}^{(j)}+\gamma r_{t+1}^{(j)}+\ldots+\gamma^{T+1-t} V_{\omega_{i}}\left(s_{T+1}\right)\)
        \(\omega=\omega_{i} ; \theta=\theta_{i} ; \psi=\psi_{i}\)
        for epoch \(\mathrm{n}=1, \ldots \mathrm{~N}\) do
            for iteration \(\mathrm{k}=1, \ldots \mathrm{~K}\) do
            // Compute discriminator loss:
                    \(\hat{L}_{D}(\psi, \theta)=\frac{1}{M T} \sum_{j=1}^{M} \sum_{t=1}^{T} \phi^{\star}\left(g_{\psi}\left(s_{t}^{(j)}, a_{t}^{(j)}\right)-\gamma g_{\psi}\left(s_{t+1}^{(j)}, a_{t+1}^{\prime(j)}\right)\right)-(1-\gamma) g_{\psi}\left(s_{1}^{(j)}, a_{t}^{\prime(j)}\right)\)
                where \(a_{t}^{\prime(j)} \sim \pi_{\theta}\left(\cdot \mid s_{1}^{(j)}\right), a_{t+1}^{\prime(j)} \sim \pi_{\theta}\left(\cdot \mid s_{t+1}^{(j)}\right)\).
                    // Update discriminator parameters: (using learning rate \(c_{\psi} \eta\) )
                    \(\psi \leftarrow \psi-c_{\psi} \eta \nabla_{\psi} \hat{L}_{D}(\psi, \theta) ;\)
            end for
            // Compute value loss:
            \(\hat{L}_{V}(\omega)=\frac{1}{M T} \sum_{j=1}^{M} \sum_{t=1}^{T}\left(V_{\omega}\left(s_{t}^{(j)}\right)-y_{t}^{(j)}\right)^{2}\)
            // Compute PPO clipped loss:
                \(\hat{L}^{\mathrm{clip}}(\theta)=\frac{1}{M T} \sum_{j=1}^{M} \sum_{t=1}^{T} \min \left\{\hat{A}\left(s_{t}^{(j)}, a_{t}^{(j)}\right) \kappa_{\pi_{\theta} / \pi_{\theta_{i}}}\left(s_{t}^{(j)}, a_{t}^{(j)}\right)\right.\),
\(\left.\hat{A}\left(s_{t}^{(j)}, a_{t}^{(j)}\right) \cdot \operatorname{clip}\left(\kappa_{\pi_{\theta} / \pi_{\theta_{i}}}\left(s_{t}^{(j)}, a_{t}^{(j)}\right), 1-\epsilon, 1+\epsilon\right)\right\}\)
            // Update parameters: (using learning rate \(\eta\) )
            \(\omega \leftarrow \omega-\eta \nabla_{\omega} \hat{L}_{V}(\omega)\);
            \(\theta \leftarrow \theta+\eta \nabla_{\theta}\left(\hat{L}^{\text {clip }}(\theta)+\lambda \cdot \hat{L}_{D}(\psi, \theta)\right)\) (if reparametrization trick applica-
            ble, else gradient step on Eq. (4.6.7))
        end for
        \(\omega_{i+1}=\omega ; \theta_{i+1}=\theta ; \psi_{i+1}=\psi\)
    end for
```

use $L_{\pi_{i}}^{\text {clip }}$ rather than $L_{\pi_{i}}$ because we expect it to work better as the clipping already provides some constraint on action probabilities. Alsggthis will allow a more direct empirical assessment of what the regularization brings compared to vanilla PPO.
$L_{\pi}^{\text {clip }}\left(\pi^{\prime}\right)-\lambda D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi}\right)$, where $\lambda$ is a regularization coefficient. If in place of the total variation we use the off-policy formulation of $\phi$-divergence $D_{\phi}\left(\mu_{\rho}^{\pi^{\prime}} \| \mu_{\rho}^{\pi}\right)$ as detailed in Equation (4.5.4), our main optimization objective can be expressed as the following min-max problem:

$$
\begin{align*}
& \max _{\pi^{\prime}} \min _{g: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}} L_{\pi_{i}}^{\operatorname{clip}}\left(\pi^{\prime}\right)-\lambda\left((1-\gamma) \mathbb{E}_{s \sim \rho, a \sim \pi^{\prime}}[g(s, a)]-\right. \\
& \left.\quad \mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi_{i}}}\left[\phi^{\star}\left(\left(g-\gamma P^{\pi^{\prime}} g\right)(s, a)\right)\right]\right), \tag{4.6.1}
\end{align*}
$$

When the inner minimization over $g$ is fully optimized, it is straightforward to show - using the score function estimator - that the gradient of this objective with respect to $\pi$ is (proof is provided in appendix):

$$
\begin{align*}
& \nabla_{\pi^{\prime}} L_{\pi_{i}}^{\operatorname{clip}^{c}}\left(\pi^{\prime}\right)-\lambda\left((1-\gamma) \mathbb{E}_{\substack{s \sim \rho \\
a \sim \pi^{\prime}}}\left[g(s, a) \nabla_{\pi^{\prime}} \log \pi^{\prime}(a \mid s)\right]\right.  \tag{4.6.2}\\
& \left.+\gamma \mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi_{i}}}\left[\frac{\partial \phi^{\star}}{\partial t}\left(\left(g-\gamma P^{\pi^{\prime}} g\right)(s, a)\right) \mathbb{E}_{\substack{s^{\prime} \sim P(\cdot \mid s, a), a^{\prime} \sim \pi^{\prime}\left(\cdot|\cdot| s^{\prime}\right)}}\left[g\left(s^{\prime}, a^{\prime}\right) \nabla_{\pi^{\prime}} \log \pi^{\prime}\left(a^{\prime} \mid s^{\prime}\right)\right]\right]\right) .
\end{align*}
$$

Furthermore, we can use the reparametrization trick if the policy $\pi$ is parametrized by a Gaussian, which is usually the case in continuous control tasks. We call the resulting new algorithm PPO-DICE, (detailed in Algorithm 4), as it uses the clipped loss of PPO and leverages the DIstribution Correction Estimation idea from Nachum et al. [2019a].

In the min-max objective (4.6.1), $g$ plays the role of a discriminator, as in Generative Adversarial Networks (GAN) [Goodfellow et al., 2014]. The policy $\pi^{\prime}$ plays the role of a generator, and it should balance between increasing the likelihood of actions with large advantage versus inducing a state-action distribution that is close to the one of $\pi_{i}$.

As shown in Algorithm 4, both policy and discriminator are parametrized by neural networks $\pi_{\theta}$ and $g_{\psi}$ respectively. We estimate the objective (4.6.1) with samples from $\pi_{i}=\pi_{\theta_{i}}$ as follows. At a given iteration $i$, we generate a batch of $M$ rollouts $\left\{s_{1}^{(j)}, a_{1}^{(j)}, r_{1}^{(j)}, s_{1}^{(j)}, \ldots, s_{T}^{(j)}, a_{T}^{(j)}, r_{T}^{(j)}, s_{T+1}^{(j)}\right\}_{j=1}^{M}$ by executing the policy $\pi_{i}$ in the environment for $T$ steps. Similarly to the PPO procedure, we learn a value function $V_{\omega}$ by updating its parameters $\omega$ with gradient descent steps, optimizing the following squared error loss:

$$
\begin{equation*}
\hat{L}_{V}(\omega)=\frac{1}{M T} \sum_{j=1}^{M} \sum_{t=1}^{T}\left(V_{\omega}\left(s_{t}^{(j)}\right)-y_{t}^{(j)}\right)^{2} \tag{4.6.3}
\end{equation*}
$$

where $y_{t}^{(j)}=r_{t}^{(j)}+\gamma r_{t+1}^{(j)}+\ldots+\gamma^{T+1-t} V_{\omega}\left(s_{T+1}\right)$. Then, to estimate the advantage, we use the truncated generalized advantage estimate

$$
\begin{equation*}
\hat{A}\left(s_{t}^{(j)}, a_{t}^{(j)}\right)=\sum_{t=1}^{T}(\gamma \lambda)^{t-1}\left(r_{t}^{(j)}+\gamma V_{\omega}\left(s_{t+1}^{(j)}\right)-V_{\omega}\left(s_{t}^{(j)}\right)\right) \tag{4.6.4}
\end{equation*}
$$

This advantage estimate is used to compute an estimate of $L_{\pi_{i}}^{\text {clip }}$ given by:

$$
\begin{align*}
\hat{L}^{\operatorname{clip}}(\theta)= & \frac{1}{M T} \sum_{j=1}^{M} \sum_{t=1}^{T} \min \left\{\hat{A}\left(s_{t}^{(j)}, a_{t}^{(j)}\right) \kappa_{\pi_{\theta} / \pi_{i}}\left(s_{t}^{(j)}, a_{t}^{(j)}\right),\right.  \tag{4.6.5}\\
& \left.\hat{A}\left(s_{t}^{(j)}, a_{t}^{(j)}\right) \cdot \operatorname{clip}\left(\kappa_{\pi_{\theta} / \pi_{i}}\left(s_{t}^{(j)}, a_{t}^{(j)}\right), 1-\epsilon, 1+\epsilon\right)\right\}
\end{align*}
$$

The parameters $\psi$ of the discriminator are learned by gradient descent on the following empirical version of the regularization term in the min-max objective (4.6.1)
$\hat{L}_{D}(\psi, \theta)=\frac{-1}{M T} \sum_{j=1}^{M} \sum_{t=1}^{T}(1-\gamma) g_{\psi}\left(s_{1}^{(j)}, a_{t}^{\prime(j)}\right)-\phi^{\star}\left(g_{\psi}\left(s_{t}^{(j)}, a_{t}^{(j)}\right)-\gamma g_{\psi}\left(s_{t+1}^{(j)}, a_{t+1}^{\prime(j)}\right)\right)$,
where $a_{t}^{\prime(j)} \sim \pi_{\theta}\left(\cdot \mid s_{1}^{(j)}\right)$ and $a_{t+1}^{\prime(j)} \sim \pi_{\theta}\left(\cdot \mid s_{t+1}^{(j)}\right)$.
If the reparametrization trick is applicable (which is almost always the case for continuous control tasks), the parameters $\theta$ of the policy are updated via gradient ascent on the objective $\hat{L}^{\text {clip }}(\theta)+\lambda \hat{L}_{D}(\psi, \theta)$ as we can backpropagate gradient though the action sampling while computing $\hat{L}_{D}(\psi, \theta)$ in Equation (4.6.6). Otherwise, $\theta$ are updated via gradient ascent on the following objective:

$$
\begin{align*}
& \hat{L}^{\operatorname{clip}}(\theta)-\frac{\lambda}{M T} \sum_{j=1}^{M} \sum_{t=1}^{T}(1-\gamma) g_{\psi}\left(s_{1}^{(j)}, a_{t}^{\prime(j)}\right) \log \pi_{\theta}\left(a_{t}^{\prime(j)} \mid s_{1}^{(j)}\right)  \tag{4.6.7}\\
& \left.\left.\quad+\gamma \frac{\partial \phi^{\star}}{\partial t}\left(g_{\psi}\left(s_{t}^{(j)}, a_{t}^{(j)}\right)-\gamma g_{\psi}\left(s_{t+1}^{(j)}, a_{t+1}^{\prime(j)}\right)\right) \cdot g_{\psi}\left(s_{t+1}^{(j)}, a_{t+1}^{\prime(j)}\right) \log \pi_{\theta}\left(a_{t+1}^{\prime(j)}\right) \right\rvert\, s_{t+1}^{(j)}\right)
\end{align*}
$$

Note that the gradient of this equation with respect to $\theta$ corresponds to an empirical estimate of the score function estimator we provided in Equation 4.6.2.

We train the value function, policy, and discriminator for $N$ epochs using $M$ rollouts of the policy $\pi_{i}$. We can either alternate between updating the policy and the discriminator, or update $g_{\psi}$ for a few steps $M$ before updating the policy. We found that the latter worked better in practice, likely due to the fact that the target distribution $\mu_{\rho}^{\pi_{i}}$ changes with every iteration $i$. We also found that increasing the learning rate of the discriminator by a multiplicative factor $c_{\psi}$ of the learning rate for the policy and value function $\eta$ improved performance.

Choice of divergence: The algorithmic approach we just described is valid with any choice of $\phi$-divergence for measuring the discrepancy between statevisitation distributions. It remains to choose an appropriate one. While Lemma 2 advocates the use of total variation distance $(\phi(t)=|t-1|)$, it is notoriously hard to train high dimensional distributions using this divergence (see Kodali
et al. [2017] for example). Moreover, the convex conjugate of $\phi(t)=|t-1|$ is $\phi^{\star}(t)=t$ if $|t| \leq \frac{1}{2}$ and $\phi^{\star}(t)=\infty$ otherwise. This would imply the need to introduce an extra constraint $\left\|g-P^{\pi} g\right\|_{\infty} \leq \frac{1}{2}$ in the formulation (4.5.4), which may be hard to optimize.

Therefore, we will instead use the KL divergence $\left(\phi(t)=t \log (t), \phi^{\star}(t)=\right.$ $\exp (t-1))$. This is still a well justified choice as we know that $D_{\mathrm{TV}}\left(\mu_{\rho}^{\pi^{\prime}} \| \mu_{\rho}^{\pi}\right) \leq$ $\sqrt{\frac{1}{2} D_{\mathrm{KL}}\left(\mu_{\rho}^{\pi^{\prime}} \| \mu_{\rho}^{\pi}\right)}$ thanks to Pinsker's inequality. We will also try $\chi^{2}$-divergence $\left(\phi(t)=(t-1)^{2}\right)$ that yields a squared regularization term.

### 4.7 Related WORK

Constraining policy updates, in order to minimize the information loss due to policy improvement, has been an active area of investigation. Kakade and Langford [2002] originally introduce CPI by maximizing a lower bound on the policy improvement and relaxing the greedification step through a mixture of successive policies. Pirotta et al. [2013] build on Kakade and Langford [2002] refine the lower bounds and introduce a new mixture scheme. Moreover, CPI inspired some popular Deep RL algorithms such as TRPO [Schulman et al., 2015] and PPO [Schulman et al., 2015], Deep CPI [Vieillard et al., 2019] and MPO [Abdolmaleki et al., 2018]. The latter uses similar updates to TRPO/PPO in the parametric version of its E-step. So, our method can be incorporated to it.

Our work is related to regularized MDP literature [Neu et al., 2017, Geist et al., 2019]. Shannon Entropic regularization is used in value iteration scheme [Haarnoja et al., 2017, Dai et al., 2018] and in policy iteration schemes [Haarnoja et al., 2018]. Note that all the mentioned works employ regularization on the action probabilities. Recently, Wang et al. [2019a] introduce divergence-augmented policy optimization where they penalize the policy update by a Bregman divergence on the state visitation distributions, motivated the mirror descent method. While their framework seems general, it doesn't include the divergences we employ in our algorithm. In fact, their method enables the use of the conditional KL divergence between state-action visitations distribution defined by $\int \mu_{\rho}^{\pi}(s, a) \log \frac{\pi(a \mid s)}{\pi^{\prime}(a \mid a)}$ and not the KL divergence $\int \mu_{\rho}^{\pi}(s, a) \log \frac{\mu_{\rho}^{\pi}(s, a)}{\mu_{\rho}^{\pi^{\prime}(s, a)}}$. Note the action probabilities ratio inside the $\log$ in the conditional KL divergence allows them to use the policy gradient theorem, a key ingredient in their framework, which cannot be done for the KL divergence.

Our work builds on recent off-policy approaches: DualDICE [Nachum et al., 2019a] for policy evaluation and ValueDICE [Kostrikov et al., 2019] for imitation learning. Both use the off-policy formulation of KL divergence. The former uses the formulation to estimate the ratio of the state visitation distributions


Figure 4.1: Comparison of $\chi^{2}$ and KL divergences for PPO-DICE for two randomly selected environments in OpenAI Gym MuJoCo and Atari, respectively. We see that KL performs better than $\chi^{2}$ in both settings. Performance plotted across 10 seeds with 1 standard error shaded.
under the target and behavior policies. Whereas, the latter learns a policy by minimizing the divergence.

The closest related work is the recently proposed AlgaeDICE [Nachum et al., 2019b] for off-policy policy optimization. They use the divergence between state-action visitation distribution induced by $\pi$ and a behavior distribution, motivated by similar techniques in Nachum et al. [2019a]. However, they incorporate the regularization to the dual form of policy performance $J(\pi)=$ $\mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi}}[r(s, a)]$ whereas we consider a surrogate objective (lower bound on the policy performance). Moreover, our method is online off-policy in that we collect data with each policy found in the optimization procedure, but also use previous data to improve stability. Whereas, their algorithm is designed to learn a policy from a fixed dataset collected by behaviour policies. Further comparison with AlgaeDICE is provided in appendix.

### 4.8 Experiments And Results

We use the PPO implementation by Kostrikov [2018] as a baseline and modify it to implement our proposed PPO-DICE algorithm. We run experiments on a randomly selected subset of environments in the Atari suite [Bellemare et al., 2013a] for high-dimensional observations and discrete action spaces, as well as on the OpenAI Gym [Brockman et al., 2016] MuJoCo environments, which have continuous state-action spaces. All shared hyperparameters are set at the same
values for both methods, and we use the hyperparameter values recommended by Kostrikov [2018] for each set of environments, Atari and MuJoCo ${ }^{3}$.

### 4.8.1 Important Aspects Of PPO-DICE

## Choice of Divergence

We conducted an initial set of experiments to compare two different choices of divergences, KL and $\chi^{2}$, for the regularization term of PPO-DICE. Figure 4.1 shows training curves for one continuous action and one discrete action environment. There, as in the other environments in which we run this comparison, KL consistently performed better than $\chi^{2}$. We thus opted to use KL divergence in all subsequent experiments.

## Effect of Varying $\lambda$

Next we wanted to evaluate the sensitivity of our method to the $\lambda$ parameter that controls the strength of the regularization. We examine in Figure 4.2 the performance of PPO-DICE when varying $\lambda$. There is a fairly narrow band for Hopper-v2 that performs well, between 0.01 and 1 . Theory indicates that the proper value for $\lambda$ is the maximum of the absolute value of the advantages (see Lemma 2). This prompted us to implement an adaptive approach, where we compute the 90th percentile of advantages within the batch (for stability), which we found performed well across environments. To avoid introducing an additional hyperparameter by tuning $\lambda$, we use the adaptive method for subsequent experiments.

[^4]

Figure 4.2: Varying $\lambda$ in Hopper_v2, 10 seeds, 1 standard error shaded. PPO-DICE is somewhat sensitive to $\lambda$ value, but the theoretically-motivated adaptive version works well.


Figure 4.3: Comparison of PPO-DICE with clipped loss $L^{\text {clip }}$ and without $L$. We see that clipping the action loss is crucial for good performance.

| Game | PPO | PPO-DICE |
| :--- | :--- | :--- |
| AirRaid | $4305.0 \pm 638.15$ | $5217.5 \pm \mathbf{7 6 9 . 1 9}$ |
| Asterix | $4300.0 \pm 169.31$ | $\mathbf{6 2 0 0 . 0} \pm \mathbf{7 5 4 . 1 0}$ |
| Asteroids | $1511.0 \pm 125.03$ | $\mathbf{1 6 5 3 . 0} \pm \mathbf{1 1 2 . 2 0}$ |
| Atlantis | $2120400.0 \pm 471609.93$ | $3447433.33 \pm \mathbf{1 0 0 1 0 5 . 8 2}$ |
| BankHeist | $1247.0 \pm 21.36$ | $\mathbf{1 2 7 3 . 3 3} \pm \mathbf{7 . 8 9}$ |
| BattleZone | $\mathbf{2 9 0 0 0 . 0} \pm \mathbf{2 6 2 0 . 4 3}$ | $19000.0 \pm 2463.06$ |
| Carnival | $3243.33 \pm 369.51$ | $3080.0 \pm 189.81$ |
| ChopperCommand | $566.67 \pm 14.91$ | $\mathbf{9 0 0 . 0} \pm \mathbf{7 7 . 4 6}$ |
| DoubleDunk | $-6.0 \pm 1.62$ | $-4.0 \pm \mathbf{1 . 2 6}$ |
| Enduro | $1129.9 \pm 73.18$ | $\mathbf{1 3 0 8 . 3 3} \pm \mathbf{1 2 0 . 0 9}$ |
| Freeway | $32.33 \pm 0.15$ | $32.0 \pm 0.00$ |
| Frostbite | $\mathbf{6 3 9 . 0} \pm 334.28$ | $296.67 \pm 5.96$ |
| Gopher | $1388.0 \pm 387.65$ | $1414.0 \pm 417.84$ |
| Kangaroo | $4060.0 \pm 539.30$ | $\mathbf{6 6 5 0 . 0} \pm \mathbf{1 5 5 8 . 1 6}$ |
| Phoenix | $\mathbf{1 2 6 1 4 . 0} \pm \mathbf{6 2 1 . 7 1}$ | $11676.67 \pm 588.24$ |
| Robotank | $7.8 \pm 1.33$ | $\mathbf{1 2 . 1} \pm \mathbf{2 . 9 1}$ |
| Seaquest | $1198.0 \pm 128.82$ | $1300.0 \pm 123.97$ |
| TimePilot | $5070.0 \pm 580.53$ | $7000.0 \pm 562.32$ |
| Zaxxon | $7110.0 \pm 841.60$ | $6130.0 \pm 1112.48$ |

Table 4.1: Mean final reward and 1 standard error intervals across 10 seeds for Atari games evaluated at 10M steps.


Figure 4.4: Results from OpenAI Gym MuJoCo suite in more complex domains, with 10 seeds and 1 standard error shaded. Results on the full suite of environments can be found in B.4.

## Importance of Clipping the Action Loss

We earlier mentioned (see footnote 2) two possible forms of our regularized objective: one with clipped action loss $L^{\text {clip }}$ and one without $L$. Clipping the action loss was an extra regularizing measure proposed in PPO [Schulman et al., 2017]. For our algorithm also, we hypothesized that it would be important for providing additional constraints on the policy update to stay within the trust region. Figure 4.3 confirms this empirically: we see the effect on our method of clipping the action loss versus keeping it unclipped. Initially, not having the additional regularization allows it to learn faster, but it soon crashes, showing the need for clipping to reduce variance in the policy update.

### 4.8.2 Results On Atari

Given our above observations we settled on using a KL-regularized $L^{\text {clip }}$, with the adaptive method for $\lambda$ that we explained Section 4.8.1. We run PPODICE on randomly selected environments from Atari. We tuned two additional hyperparameters, the learning rate for the discriminator and the number of discriminator optimization steps per policy optimization step. We found that $K=5$ discriminator optimization steps per policy optimization step performed well.

Fewer steps showed worse performance because the discriminator was not updating quickly enough, while more optimization steps introduced instability from the discriminator overfitting to the current batch. We also found that increasing the discriminator learning rate to be $c_{\psi}=10 \times$ the policy learning rate helped most environments. We used the same hyperparameters across all environments. Results are shown in Figure 4.1. We see that PPO-DICE significantly outperforms PPO on a majority of Atari environments.

### 4.8.3 Results On OpenAI Gym MuJoCo

For the OpenAI Gym MuJoCo suite, we also used $K=5$ discriminator optimization steps per policy optimization step, and $c_{\psi}=10 \times$ learning rate for the discriminator in all environments. We selected 5 of the more difficult environments to showcase in this present (4.4), but additional results on the full suite and all hyperparameters used can be found in B.4. We again see improvement in performance in the majority of environments with PPO-DICE compared to PPO and TRPO.

### 4.9 CONCLUSION

In this work, we have argued that using the action probabilities to constrain the policy update is a suboptimal approximation to controlling the state visitation distribution shift. We then demonstrate that using the recently proposed DIstribution Correction Estimation idea [Nachum et al., 2019a], we can directly compute the divergence between the state-action visitation distributions of successive policies and use that to regularize the policy optimization objective instead. Through carefully designed experiments, we have shown that our method beats PPO in most environments in Atari [Bellemare et al., 2013a] and OpenAI Gym MuJoCo [Brockman et al., 2016] benchmarks.

# Online Learning in Smooth Markov Decision Processes 

### 5.1 Protologue To The Contribution

### 5.1.1 Article Details

This paper is based on the preprint Zooming for Efficient Model-Free Reinforcement Learning in Metric Spaces [Touati et al., 2020b]. This is a joint work with Adrien Ali Taiga and Marc G. Bellemare. I am the main author. I performed the derivation of the new algorithm and its theoretical analysis.

### 5.1.2 Context

In previous chapters, we were concerned with better exploiting the off-policy data but we didn't talk that much about the data collection itself. While recent RL empirical successes have been demonstrated in simulated environments where data can be generated indefinitely up to computational limits, bringing RL to real systems require us to focus on statistical efficiency as real data collection is costly and limited by the physical context. A statistically efficient agent must carefully balance between exploiting its current knowledge and exploring towards informative data.

My interest in this line of research was triggered by an inspiring talk given by Benjamin Van Roy at Mcgill university about what he coined as deep exploration. Deep exploration means exploration which is directed over multiple time steps and takes into account the consequences of an action on future learning. In Osband et al. [2019], his group developed a new approach consisting in choosing actions that are greedy with respect a randomly drawn value function from a proxy of the posterior distribution over value functions. As a first contribution on the topic of exploration, I proposed in Touati et al. [2020a] to scale the idea of randomized value functions to deep RL by leveraging recent advances in variational inference to account for the uncertainty of estimates

Now, the missing piece is that even though randomized value functions offer a promising computationally tractable approach for high dimensional state spaces, the algorithm comes with no theoretical guarantee when combined with function approximation. Especially, it is only shown to be provably efficient with tabular representation. In fact, there was (and it is still true at some extent) a huge gap between the RL theory and practise: despite the wealth of research
into provably efficient reinforcement learning algorithms, most works focus on tabular representation and thus struggle to handle exponentially or infinitely large state-action spaces. This present work is a step towards bridging this gap.

To make exploration with large state-action space intractable, we need to make some structural assumption. In this work, we assume that the MDP is smooth: a natural metric characterizes the proximity between different states and actions and the optimal action-value function is Lipschitz with respect to this metric. This is also known as the metric space assumption.

### 5.1.3 Paper Abstract

Despite the wealth of research into provably efficient reinforcement learning algorithms, most works focus on tabular representation and thus struggle to handle exponentially or infinitely large state-action spaces. In this paper, we consider episodic reinforcement learning with a metric state-action space. We propose ZOOMRL, an online model-free algorithm that learns an adaptive discretization of the joint space by zooming in more promising and frequently visited regions while carefully balancing the exploitation-exploration trade-off. We show that ZOOMRL achieves a worst-case regret $\tilde{O}\left(H^{\frac{5}{2}} K^{\frac{d+1}{d+2}}\right)$ where $H$ is the planning horizon, $K$ is the number of episodes and $d$ is the covering dimension of the space with respect to the metric. Moreover, our algorithm enjoys improved metric-dependent guarantees that reflect the geometry of the underlying space. Finally, we show that our algorithm is robust to small misspecification errors.

### 5.1.4 Recent Developments

The bounds presented in this work are instance-independent guarantees. A recent work Cao and Krishnamurthy [2020] shows that adaptive partitioning approaches achieve problem-dependent guarantees which depend on the zooming dimension instead of the covering dimension of the space. The zooming dimension is a measure of near-optimal regions. In benign instances where the near-optimal regions concentrate to a low dimensional manifold, the zooming dimension could be much smaller than the covering dimension, which enables sharper regret bounds.

Model-based algorithms for metric spaces was proposed in Domingues et al. [2020b] using non-parametric kernel estimators and in Sinclair et al. [2020] using adaptive discretization but their regret bounds have sub-optimal dependence on K.

### 5.2 InTRODUCTION

In the regime of MDPs with a finite state-action space, the OFU principle has been successfully implemented and efficient algorithms typically achieve regret that scales sublinearly with the number of discrete states and the number of discrete actions. This precludes applying them to arbitrarily large state-action spaces. On the other hand, MDPs with continuous state-action spaces have been an active area of investigation [Ortner and Ryabko, 2012, Lakshmanan et al., 2015, Song and Sun, 2019]. A common theme is to assume some structure knowledge, such as the existence of similarity metric between state-action pairs, and then to use a uniform discretisation of the space or nearest-neighbor approximators.

In this work, we focus on the finite-horizon MDP formalism with an unknown transition kernel. We suppose that the state-action space is equipped by a metric that characterizes the proximity between different states and actions. Such metrics have been studied in previous work for state aggregation [Ferns et al., 2004, Ortner, 2007]. We assume that the optimal action-value function is Lipschitz continuous with respect to this metric, which means that state-action pairs that are close to each other have similar optimal values.

We propose an online model-free RL algorithm, ZOOMRL, that actively explores the state-action space by learning on-the-fly an adaptive partitioning. Algorithms based on uniform partitions, such as the works in Ortner and Ryabko [2012] and Song and Sun [2019], disregard the shape of the optimal value function and thus could waste effort in partitioning irrelevant regions of the space. Moreover, the granularity of the partition should be tuned and it depends on the time horizon and the covering dimension of joint space. In contrast, ZOOMRL is able to take advantage of the structure of the problem's instance at hand by adjusting the discretisation to frequently visited and high-rewarding regions to get better estimates. Zooming approaches have been successfully applied in Lipschitz bandits [Kleinberg et al., 2008] and continuous contextual bandits [Slivkins, 2014]. However, in the bandit setting, an algorithm's cumulative regret can be easily decomposed into regret incurred in each sub-partition which is controlled by the size of the sub-partition itself. In contrast, in the reinforcement learning setting, the errors are propagated through iterations and we need to carefully control how they accumulate over iterations and navigate through subpartitions. We show that ZOOMRL achieves a worst-case regret $\tilde{O}\left(H^{\frac{5}{2}} K^{\frac{d+1}{d+2}}\right)$ where $H$ is the planning horizon, $K$ is the number of episodes and $d$ is the covering dimension of the space with respect to the metric. Moreover, ZOOMRL enjoys an improved metric-dependent guarantee that reflects the geometry of the underlying space and whose scaling in terms of $K$ is optimal as it matches the lower bound in continuous contextual bandit [Slivkins, 2014] when $H=1$. Finally, we study how our algorithm cope with the misspecified setting (Assump-
tion 5). We show that it is robust to small misspecification error as it suffers only from an additional regret term $O(H K \epsilon)$ if the true optimal action-value function is Lipschitz up to an additive error uniformly bounded in absolute value by $\epsilon$.

### 5.3 Related Work

Exploration in metric spaces: There have been several recent works that study exploration in continuous state-action MDPs under different structured assumptions. Kakade et al. [2003] assume a local continuity of the reward function and the transition kernel with respect to a given metric. They propose a generalization of the $E^{3}$ algorithm of Kearns and Singh [2002] to metric spaces. Their sample complexity depends on the covering number of the space under the continuity metric instead of the number of the states. However, their algorithm requires access to an approximate planning oracle. Lattimore et al. [2013] assume that the true transition kernel belongs to a finite or compact hypothesis class. Their algorithm consists in maintaining a set of transitions models and pruning it over time by eliminating the provable implausible models. They establish a sample complexity that depends polynomially on the cardinality or covering number of the model class. Pazis and Parr [2013] consider a continuous state-action MDP, develop a nearest-neighbor based algorithm under the assumption that all Qfunctions encountered are Lipschitz continuous, showing a sample complexity that depends on an approximate covering number. Ortner and Ryabko [2012] develop a model-based algorithm that combines state aggregation with the standard UCRL2 algorithm [Jaksch et al., 2010] under the assumption of Lipschitz or Hölder continuity of rewards and transition kernel and they establish a regret bound scaling in $K^{\frac{2 d+1}{2 d+2}}$ where $d$ in the dimension of the state space and $K$ is the number of episodes. Lakshmanan et al. [2015] improve the latter work by considering a kernel density estimator instead of a frequency estimator for the transition probabilities. They achieve a regret bound of $K^{\frac{d+1}{d+2}}$. Yang et al. [2019] consider a deterministic control system under a Lipschitz assumption of the optimal action-value functions and the transition function and they establish a regret of $K^{\frac{d-1}{d}}$ where $d$ here is the doubling dimension. Recently, Song and Sun [2019] extended the tabular Q-learning with upper-confidence bound exploration strategy, developed in Jin et al. [2018], to continuous state-action MDPs using a uniform discretisation of the joint space leading to the regret bound $\tilde{O}\left(H^{\frac{5}{2}} K^{\frac{d+1}{d+2}}\right)$ where $H$ is the planning horizon and $d$ is the covering dimension. They only assume that the optimal action-value function is Lipschitz continuous. This assumption is more general than that used in the aforementioned works as it is known that Lipschitz continuity of the reward function and
the transition kernel leads to Lipschitz continuity of the optimal action-value function [Asadi et al., 2018]. We use the same condition in this present paper.

Adaptive discretization: Our method is closely related to methods that learn partition from continuous bandit literature [Kleinberg et al., 2008, Bubeck et al., 2009, Slivkins, 2014, Azar et al., 2014, Munos et al., 2014]. In particular, our method is inspired by the contextual Zooming algorithm introduced in Slivkins [2014] for contextual bandits, that we extend in non-trivial way to episodic RL setting. Our method is similar to two recently proposed algorithms. Zhu and Dunson [2019] propose and analyze an adaptive partitioning algorithm approach in the specific case where the metric space is a subset of $\mathbb{R}^{d}$ equipped with $l_{\infty}$ distance as similarity metric. Concurrently to our work, Sinclair et al. [2019] extend the latter result to any generic metric space. However, their algorithm ADAPTIVE Q-LEARNING requires, at each re-partition step, a packing oracle that is able to take a region and value $r$ and outputs an $r$-packing of that region. Whereas, our algorithm is oracle-free and creates at most a single subregion when needed. More comparison with this work requires the introduction of some notations and is therefore deferred to Section 5.5.

### 5.4 Problem Statement

### 5.4.1 Episodic Reinforcement Learning and Regret

We consider a finite horizon $\operatorname{MDP}(\mathcal{S}, \mathcal{A}, P, r, H)$. We assume that reward $r_{h}(s, a) \in[0,1]$ for any state-action $(s, a)$ and $h \in[H]$. We recall that for any step $h \in[H]$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$, the state-action value function of a non-stationary deterministic policy $\pi=\left(\pi_{1}, \ldots, \pi_{H}\right)$ is defined as $Q_{h}^{\pi}(s, a)=r_{h}(s, a)+\mathbb{E}\left[\sum_{i=h+1}^{H} r_{i}\left(s_{i}, \pi_{i}\left(s_{i}\right)\right) \mid s_{h}=s, a_{h}=a\right]$, and the value function is $V_{h}^{\pi}(s)=Q_{h}^{\pi}\left(s, \pi_{h}(s)\right)$. As the horizon is finite, under some regularity conditions [Shreve and Bertsekas, 1978], there always exists an optimal policy $\pi^{\star}$ whose value and action-value functions are defined as $V_{h}^{\star}(x) \triangleq V_{h}^{\pi^{\star}}(s)=\max _{\pi} V_{h}^{\pi}(s)$ and $Q_{h}^{\star}(s, a) \triangleq Q_{h}^{\pi^{\star}}(s, a)=\max _{\pi} Q_{h}^{\pi}(s, a)$. Moreover, both $Q^{\pi}$ and $Q^{\star}$ can be conveniently written as the result of the following Bellman equations

$$
\begin{align*}
Q_{h}^{\pi}(s, a) & =r_{h}(s, a)+\left(P_{h} V_{h+1}^{\pi}\right)(s, a)  \tag{5.4.1}\\
Q_{h}^{\star}(s, a) & =r_{h}(s, a)+\left(P_{h} V_{h+1}^{\star}\right)(s, a) \tag{5.4.2}
\end{align*}
$$

where $V_{H+1}^{\tau}(s)=V_{H+1}^{\star}(s)=0$ and $V_{h}^{\star}(s)=\max _{a \in \mathcal{A}} Q_{h}^{\star}(s, a)$, for all $s \in \mathcal{S}$.

We focus on the online episodic reinforcement learning setting in which the reward and the transition kernel are unknown. The learning agent plays the game for $K$ episodes $k=1, \ldots, K$, where each episode $k$ starts from some initial state $s_{1}^{k}$ sampled according to some initial distribution. The agent controls the system by choosing a policy $\pi_{k}$ at the beginning of the $k$-th episode. The total expected regret is defined then

$$
\operatorname{REGRET}(K) \triangleq \sum_{k=1}^{K} V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right)
$$

### 5.4.2 Metric Space

We assume that the state-action space $\mathcal{X} \triangleq \mathcal{S} \times \mathcal{A}$ is compact endowed with a metric dist : $\mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^{+}$. This leads us to state our main assumption:
Assumption 4 (Lipschitz Continuous $Q^{\star}$ ). We assume that for any $h \in[H], Q_{h}^{\star}$ is L-Lipschitz continuous: for all $(s, a),\left(s^{\prime}, a^{\prime}\right) \in \mathcal{S} \times \mathcal{A}$

$$
\left|Q_{h}^{\star}(s, a)-Q_{h}^{\star}\left(s^{\prime}, a^{\prime}\right)\right| \leq L \cdot \operatorname{dist}\left((s, a),\left(s^{\prime}, a^{\prime}\right)\right)
$$

and without loss of generality,

$$
\operatorname{dist}\left((s, a),\left(s^{\prime}, a^{\prime}\right)\right) \leq 1, \forall(s, a),\left(s^{\prime}, a^{\prime}\right) \in \mathcal{S} \times \mathcal{A}
$$

Assumption 4 tells us that the optimal action values of nearby state-action pairs are close.

For a metric space $\mathcal{X}$ and $\epsilon>0$, we denote the $\epsilon$-net, $\mathcal{N}(\epsilon) \subset \mathcal{X}$, as a set such that

$$
\forall x \in \mathcal{X}, \quad \exists x^{\prime} \in \mathcal{N}(\epsilon), \quad \operatorname{dist}\left(x, x^{\prime}\right) \leq \epsilon .
$$

If $\mathcal{X}$ is compact, we denote $N(\epsilon)$ as the minimum size of an $\epsilon$-net for $\mathcal{X}$. The covering dimension $d$ of $\mathcal{X}$ is defined

$$
d \triangleq \inf _{d^{\prime}}\left\{d^{\prime} \geq 0, \forall \epsilon>0 \quad N(\epsilon) \leq \epsilon^{-d^{\prime}}\right\}
$$

In particular, if $\mathcal{X}$ is a subset of Euclidean space equipped with $l_{p}$ distance then its covering dimension is at most the linear dimension of $\mathcal{X}$. In many applications of interests, state-action spaces are commonly thought to be concentrated near a lower-dimensional manifold lying in high-dimensional ambient space. In this case, the covering dimension is much smaller than the linear dimension of the ambient space.

Covering is closely related to packing. We denote an $\epsilon$-packing, $\mathcal{M}(\epsilon) \subset \mathcal{X}$, as a set such that

$$
\forall x, x^{\prime} \in \mathcal{M}(\epsilon), \quad \operatorname{dist}\left(x, x^{\prime}\right)>\epsilon
$$

If $\mathcal{X}$ is compact, we denote $M(\epsilon)$ as the maximum size of an $\epsilon$-packing. $N(\epsilon)$ and $M(\epsilon)$ have the same scaling as we have $M(2 \epsilon) \leq N(\epsilon) \leq M(\epsilon)$.

### 5.5 THE ZOOMRL ALGORITHM

The ZOOMRL algorithm, shown in Algorithm 5, incrementally builds an optimistic estimate of the optimal action-value function over $\mathcal{X}$. The main idea is to estimate $Q$-values precisely in near-optimal regions, while estimating it loosely in sub-optimal regions. To implement this idea, we learn a partition of the space by zooming in more promising and frequently visited regions.

ZOOMRL maintains a partition of the space $\mathcal{X}$ that consists of a growing set of balls, of various sizes. Initially the set contains a single ball which includes the entire state-action space. Over time the set is expanded to include additional balls. The algorithm assigns two quantities to each ball: the number of times the ball is selected and an optimistic estimate of the $Q$-value of its center. By interpolating between these estimates using the Lipschitz structure, the algorithm assigns a tighter upper bound (called index) of the $Q$-value of each ball's center. These indices are then used to select the next ball and the next action to execute (line 8-10 of Algorithm 5). Based on the received reward and the observed next state, the algorithm updates the selected ball's statistics (cf line 14-17 of Algorithm 5). Then, one ball may be created inside the selected ball according to an activation rule that reflects a bias-variance tradeoff (line 19-22 of Algorithm 5).

We denote by $\mathcal{B}_{h}$ the set of balls at step $h \in[H]$ that may change from episode to episode. Each ball $B=\left\{x \in \mathcal{X}\right.$, dist $\left.\left(x_{B}, x\right) \leq \operatorname{rad}(B)\right\}$ has a radius $\operatorname{rad}(B)$, center $x_{B}=\left(s_{B}, a_{B}\right)$ and a domain. The domain of ball $B \in \mathcal{B}_{h}$, denoted by $\operatorname{dom}_{h}(B)$, is defined as the subset of $B$ that excludes all active balls in $\mathcal{B}_{h}$ that have radius strictly smaller than $\operatorname{rad}(B)$, i.e

$$
\operatorname{dom}_{h}(B) \triangleq B \backslash\left(\begin{array}{l}
\left.\cup \begin{array}{c}
B^{\prime} \in \mathcal{B}_{h} \\
\operatorname{rad}\left(B^{\prime}\right)<\operatorname{rad}(B)
\end{array} B^{\prime}\right) . . . . ~ . ~
\end{array}\right.
$$

$B$ is called relevant to a state $s$ at step $h$ if $(s, a) \in \operatorname{dom}_{h}(B)$ for some action $a \in \mathcal{A}$. We denote the set of relevant balls to a given state $s$ at step $h$ by rel $l_{h}(s) \triangleq\{B \in$ $\left.\mathcal{B}_{h}: \exists a \in \mathcal{A},(s, a) \in \operatorname{dom}_{h}(B)\right\}$. For each ball $B \in \mathcal{B}_{h}$ for some $h$, we keep track of the number of times $B$ is selected at step $h\left(\right.$ denoted by $\left.n_{h}(B)\right)$, as well as a high probability upper bound (denoted by $\widehat{\mathrm{Q}}_{h}(B)$ ) for the optimal $Q$-value of the center of $B$ (i.e $Q_{h}^{\star}\left(s_{B}, a_{B}\right)$ ).

Using the Lipschitz continuity assumption, we have that $L \cdot \operatorname{rad}(B)+$ $\widehat{\mathrm{Q}}_{h}\left(B^{\prime}\right)+L \cdot \operatorname{dist}\left(x_{B}, x_{B^{\prime}}\right)$ is a valid high probability upper bound on $Q^{\star}\left(s_{B}, a_{B}\right)$ for any $B^{\prime} \in \mathcal{B}_{h}$. Consequently, we get a tighter (less overoptimistic) upper bound, denoted by index ${ }_{h}(B)$, by taking the minimum of these bounds

$$
\operatorname{index}_{h}(B) \triangleq L \cdot \operatorname{rad}(B)+\min _{\substack{B^{\prime} \in \mathcal{B}_{h} \\ \operatorname{rad}\left(B^{\prime}\right) \geq \operatorname{rad}(B)}}\left\{\widehat{\mathrm{Q}}_{h}\left(B^{\prime}\right)+L \cdot \operatorname{dist}\left(B, B^{\prime}\right)\right\}
$$

where, by abuse of notation, we write dist $\left(B, B^{\prime}\right)=\operatorname{dist}\left(x_{B}, x_{B^{\prime}}\right)$ ).

To facilitate the algorithm's description, we introduce episode-indexed versions of the quantities, as shown in algorithm 5 . We will use $s_{h}^{k}, a_{h}^{k}$ and $B_{h}^{k}$ to represent the state, the action and the ball generated at time step $h$ of the $k$-th episode. Moreover, $\widehat{\mathrm{Q}}_{h}^{k}(B)$ and $n_{h}^{k}(B)$ are the statistics associated with each ball $B$ at time step $h$ at the beginning of the $k$-th episode.

The algorithm proceeds as follows. Initially, ZoOMRL creates a ball centered at arbitrary state-action pair with radius 1 , hence covering the whole space. At step $h$ of the $k$-th episode, a state $s_{h}^{k}$ is observed, the algorithm finds the set of relevant balls to $s_{h}^{k}$ (i.e rel ${ }_{h}^{k}\left(s_{h}^{k}\right)$ ) and picks the ball $B_{h}^{k}$ with the largest index (i.e index ${ }_{h}^{k}$ ) among the relevant balls. Once the ball is selected, an action $a_{h}^{k}$ is chosen randomly among actions $a$ satisfying $\left(s_{h}^{k}, a\right) \in \operatorname{dom}_{h}^{k}\left(B_{h}^{k}\right)$. Action $a_{h}^{k}$ is then executed in the environment, a reward $r_{h}^{k}$ is obtained and next state $s_{h+1}^{k}$ is observed.

Based on the received reward and next state, the algorithm updates the statistics of the selected ball. The number of visits $n_{h}^{k}\left(B_{h}^{k}\right)$ is incremented by 1: $n_{h}^{k+1}\left(B_{h}^{k}\right)=n_{h}^{k}\left(B_{h}^{k}\right)+1$. Let $t=n_{h}^{k+1}\left(B_{h}^{k}\right)$, the $Q$-value estimate is updated as follows

$$
\begin{aligned}
\widehat{\mathrm{Q}}_{h}^{k+1}\left(B_{h}^{k}\right) \leftarrow & \left(1-\alpha_{t}\right) \widehat{\mathrm{Q}}_{h}^{k}\left(B_{h}^{k}\right)+\alpha_{t}\left(r_{h}^{k}+\widehat{\mathrm{V}}_{h+1}^{k}\left(s_{h+1}^{k}\right)\right. \\
& \left.+u_{t}+2 L \cdot \operatorname{rad}\left(B_{h}^{k}\right)\right)
\end{aligned}
$$

$V_{h+1}^{k}\left(x_{h+1}^{k}\right)=\min \left\{H, \max _{B \in \operatorname{rel}_{h+1}^{k}\left(s_{h+1}^{k}\right)}\right.$ index $\left._{h+1}^{k}(B)\right\}$ here is the estimate of the next state's value and $\alpha_{t} \triangleq \frac{H+1}{H+t}$ is a learning rate. The term $u_{t}+2 L \cdot \operatorname{rad}\left(B_{h}^{k}\right)$ corresponds to an exploration bonus used to bound estimation errors on the value function with high probability. The first term of the bonus is set to $u_{t}=$ $4 \sqrt{\frac{H^{3} l}{t}}$ (we use $\imath \triangleq \log \left(4 H K^{2} / p\right)$ for $p \in(0,1)$ to denote the $\log$ factor). It corresponds to a Hoeffding-style bonus which reflects the sample uncertainty due to insufficient number of samples. The second term $2 L \cdot \operatorname{rad}\left(B_{h}^{k}\right)$ accounts for the maximum possible variation of $Q$-values over the selected ball $B_{h}^{k}$.

Contrary to Q-values, the value of next state $\widehat{\mathrm{V}}_{h}^{k}\left(s_{h+1}^{k}\right)$ is defined over the entire state space, we don't need to maintain $\widehat{\mathrm{V}}_{h}$ but we query it whenever we need. In particular, $\widehat{\mathrm{V}}_{h}^{k}\left(s_{h+1}^{k}\right)$ is defined by the largest index among the relevant balls to $s_{h+1}^{k}$ clipped above by $H$. The clipping here is to keep the value estimate into the range of plausible values while preserving the optimism as $H$ is an upper bound on the true optimal value function.

Finally, ZOOMRL may create a new ball according to the following activation rule: If $n_{h}^{k}\left(B_{h}^{k}\right) \geq \frac{1}{\operatorname{rad}\left(B_{h}^{k}\right)^{2}}$, a new ball $B^{\prime}$, centered in $\left(s_{h^{\prime}}^{k} a_{h}^{k}\right)$ and radius $\operatorname{rad}\left(B^{\prime}\right)=\operatorname{rad}\left(B_{h}^{k}\right) / 2$, is created and its $Q$ value is initialized to $H$. The activation criterion (which is equivalent to $\frac{1}{\sqrt{n_{h}^{k}\left(B_{h}^{k}\right)}} \leq \operatorname{rad}\left(B_{h}^{k}\right)$ ) reflects a tradeoff


Figure 5.1: Left: a possible partition for 2-dimensional state-action space. For a given state $s$, we show in orange the three relevant balls $A, B$ and $C$. Right: For the three relevant balls, we show how the index is constructed based on the interpolation between $Q$-value estimates of each ball. The gray piecewise linear curve corresponds to the function : $a \rightarrow \min _{B^{\prime}}\left\{\widehat{\mathrm{Q}}_{h}\left(B^{\prime}\right)+L \cdot \operatorname{dist}\left((s, a), x_{B^{\prime}}\right)\right\}$ for a given state.
between the variance of $Q$-value estimate, due to the number of samples (i.e $\frac{1}{\sqrt{n_{h}^{k}(B)}}$ ) and the bias, corresponding to the radius of the ball.

Comparison with Sinclair et al. [2019]: Concurrently to our work, Sinclair et al. [2019] use a similar approach to learn an adaptive discretization. We highlight here differences between ZoOmRL and their algorithm AdAptive QLEARNING:

- ADAPTIVE Q-LEARNING requires, at each re-partition step, a packing oracle that is able to take a ball $B$ and value $r$ to output an $r$-packing of $B$. Whereas, our algorithm is oracle-free and creates at most a single child ball when needed.
- We leverage more the Lipschitz structure to define the ball's index, which is not used in their algorithm.
- Sinclair et al. [2019] use an exploration bonus $2 \sqrt{\frac{H^{3} \log (4 H K / p)}{t}}+\frac{4 L}{\sqrt{t}}$ where $t=n_{h}^{k}\left(B_{h}^{k}\right)$. The first term looks similar to our term $u_{t}$ with $K^{2}$ instead of $K$ in the $\log$ factor. We think it is due to small issue in their proof because there is a missing union bound over $K$ possible values of the random stop-
ping time $t=n_{h}^{k}\left(B_{h}^{k}\right)$ (cf Proof C.1.3 in appendix). The second term, $\frac{4 L}{\sqrt{t}}$, is different than ours, $L \cdot \operatorname{rad}\left(B_{h}^{k}\right)$.
- Finally, in [Sinclair et al., 2019] each child ball inherents statistics from their parent while in our algorithms the statistics are initialized by zero for $n_{h}$ and $H$ for $\widehat{\mathrm{Q}}_{h}$.

```
Algorithm 5 ZOOMRL
    Data: For \(h \in[H]\), we have a collection \(\mathcal{B}_{h}\) of balls.
    Init: create ball \(B\), with \(\operatorname{rad}(B)=1\) and arbitrary center. \(\mathcal{B}_{h}^{1} \leftarrow\{B\}\) for all
    \(h \in[H]\)
    \(\widehat{Q}_{h}^{1}(B)=H\) and \(n_{h}^{1}(B)=0, \forall h \in[H]\)
    for episode \(k=1, \ldots, K\) do
        Observe \(x_{1}^{k}\)
        for step \(h=1, \ldots, H\) do
            // Select action
            \(B_{h}^{k} \leftarrow \arg \max _{B \in \operatorname{rel}}^{h}{ }_{h}^{k}\left(s_{h}^{k}\right)\) index \({ }_{h}^{k}(B)\)
            \(a_{h}^{k} \leftarrow\) any arm \(a\) such that \(\left(s_{h}^{k}, a\right) \in \operatorname{dom}\left(B_{h}^{k}\right)\)
            Execute action \(a_{h}^{k}\), observe reward \(r_{h}^{k}\) and next state \(s_{h+1}^{k}\)
            // Query the next value function
            \(\widehat{\mathrm{V}}_{h+1}^{k}\left(s_{h+1}^{k}\right) \leftarrow \min \left\{H, \max _{B \in \operatorname{rel}_{h+1}^{k}\left(s_{h+1}^{k}\right)} \operatorname{index}_{h+1}^{k}(B)\right\}\)
            // Update the selected ball's statistics
            \(t=n_{h}^{k+1}\left(B_{h}^{k}\right) \leftarrow n_{h}^{k}\left(B_{h}^{k}\right)+1\)
            \(u_{t} \leftarrow 4 \sqrt{\frac{H^{3} l}{t}}\)
            \(\widehat{\mathrm{Q}}_{h}^{k+1}\left(B_{h}^{k}\right) \leftarrow\left(1-\alpha_{t}\right) \widehat{\mathrm{Q}}_{h}^{k}\left(B_{h}^{k}\right)+\alpha_{t}\left(r_{h}^{k}+\widehat{\mathrm{V}}_{h+1}^{k}\left(s_{h+1}^{k}\right)+u_{t}+2 L \cdot \operatorname{rad}\left(B_{h}^{k}\right)\right)\)
            // New ball's activation step
            if \(n_{h}^{k}(B) \geq \frac{1}{\operatorname{rad}\left(B_{h}^{k}\right)^{2}}\) then
            Create a new ball \(B^{\prime}\) centered in \(\left(s_{h^{\prime}}^{k}, a_{h}^{k}\right)\) and radius \(\operatorname{rad}\left(B^{\prime}\right)=\)
            \(\operatorname{rad}\left(B_{h}^{k}\right) / 2\)
            \(\mathcal{B}_{h}^{k+1}=\mathcal{B}_{h}^{k} \cup B^{\prime}\)
            \(\widehat{\mathrm{Q}}_{h}^{k+1}\left(B^{\prime}\right)=H\) and \(n_{h}^{k+1}\left(B^{\prime}\right)=0, \forall h \in[H]\)
            end if
        end for
    end for
```


### 5.6 MAIN RESUlts

In this section, we present our main theoretical result which is an upper bound on the total regret of ZOOMRL (see Algorithm 5). We start by showing a pessimistic version of the regret bound.

Theorem 1 (Worst case guarantee). For any $p \in(0,1)$, with probability $1-p$, the total regret of ZOOMRL (see Algorithm 5) is at most $O\left(\sqrt{H^{5} \imath} L K^{\frac{d+1}{d+2}}\right)$ where $\imath=$ $\log \left(4 H K^{2} / p\right)$ and $d$ is the covering dimension of the state-action space.

The bound in Theorem 1 matches the regret bound achieved by Net-based Q-learning (NBQL) studied in Song and Sun [2019] which assumes access to an $\epsilon$-net of the whole space as input to the algorithm. Moreover, the $\epsilon$-net should be optimal in the sense that the granularity $\epsilon$ of the covering must be chosen in advance $\left(\epsilon=K^{\frac{-1}{d+2}}\right.$ ). Meanwhile, ZOOMRL builds the partition on the fly and in data-dependent fashion by allocating more effort in promising regions, which would considerably save the memory requirement in favorable problems while preserving the worst-case guarantee (as shown in Theorem 1).

Now, we present a refined regret bound that reflects better the geometry of the underlying space.

Theorem 2 (Refined regret bound). For any $p \in(0,1)$, with probability $1-p$, the total regret of ZOOMRL (see Algorithm 5) is at most

$$
O\left(\left(L+\sqrt{H^{5} \imath}\right) \min _{r_{0} \in(0,1)}\left\{K r_{0}+\sum_{\substack{r=2^{-i} \\ r \geq r_{0}}} \frac{M(r)}{r}\right\}+H^{2}+\sqrt{H^{3} K \imath}\right)
$$

where $\imath=\log \left(4 H K^{2} / p\right), M(r)$ is the $r$-packing number of the state-action space.
Since $M(r)$ is non-increasing in $r$, the leading term (the first term) of the bound of Theorem 2 is upper bounded by $\min _{r_{0} \in(0,1)}\left\{K r_{0}+\frac{M\left(r_{0}\right)}{r_{0}} \log \left(r_{0}\right)\right\}$. By setting $r_{0}=K^{\frac{-1}{d+2}}$, we recover the worst-case bound in Theorem 1. Note that the work of Sinclair et al. [2019] achieves the same regret bound.

The $r$-packing number $M(r)$ is here to uniformly upper bound the number of balls of radius $r$ generated by the algorithm, as we will see in the analysis deferred to the next section. Intuitively, balls with small radius would not cover the whole state-action space but rather would be concentrated around near-optimal regions. We expect that their number would be much smaller that $M(r)$ in practice.

Comparison with contextual bandit setting: We would like to highlight a negative result of the RL setting comparing to the contextual bandit setting. When $H=1$ and if we ignore logarithmic factors in Theorem 2, we obtain a bound in $\min _{r_{0} \in(0,1)}\left\{K r_{0}+\sum_{r=2^{-i} \geq r_{0}} \frac{M(r)}{r}\right\}$. This looks similar to the regret bound of the contextual Zooming algorithm [Slivkins, 2014]. But there is a crucial difference: $M(r)$ here is the $r$-packing of the entire space while it is replaced by the $r$-packing of near-optimal regions in Slivkins [2014]. This follows from the fact that in contextual bandit setting, total regret could be straightforwardly written as sum of instant regrets incurred by each ball. Such regret is bounded, up to a multiplicative constant, by the radius of the ball in which the context falls.

However the dependence of our regret bound on $K$ is still optimal, up to logarithmic factor, with respect to the worst Lipschitz structure. In fact, theorem 8 in Slivkins [2014] states that there exists a distribution $\mathcal{I}$ over problem instances on $(\mathcal{S} \times \mathcal{A}$,dist) such that for any contextual bandit algorithm, the expected regret over $\mathcal{I}$ is lower bounded by $\Omega\left(\min _{r_{0} \in(0,1)}\left\{K r_{0}+\sum_{r=2^{-i} \geq r_{0}} \frac{M(r)}{r}\right\} / \log (K)\right)$.

Tabular MDP In the case of finite state-action MDP without any structural knowledge, one can pick the metric to be dist $\left((s, a),\left(s^{\prime}, a^{\prime}\right)\right)=H, \quad \forall(s, a) \neq$ $\left(s^{\prime}, a^{\prime}\right)$. It is obvious that the optimal action-value function is 1-Lipschitz (i.e $L=1$ ) with respect to this metric and that the packing number is at most equal to $|\mathcal{S} \| \mathcal{A}|$. In this case, if we set $r_{0}$ to $\sqrt{\frac{|\mathcal{S}||\mathcal{A}|}{K}}$, the regret bound in Theorem 2 becomes $O\left(\sqrt{|\mathcal{S}||\mathcal{A}| H^{5} \mathrm{~K}}\right)$. Hence, we recover exactly the regret bound of $Q$ learning with UCB-Hoeffding algorithm of Jin et al. [2018].

### 5.6.1 Result For The Misspecified Case

We study now how ZoomRL deals with misspecification error. First, we present a formal definition for an approximate Lipscthiz $Q$-value.

Assumption 5 (Approximately Lipschitz $Q^{\star}$ ). We assume that for any $h \in[H], Q_{h}^{\star}$ can be decomposed as a L-Lipschitz continuous term and a bounded term as:

$$
\forall(s, a) \in \mathcal{X}, Q_{h}^{\star}(s, a)=f_{h}(s, a)+\Delta_{h}(s, a)
$$

where for all $(s, a),\left(s^{\prime}, a^{\prime}\right) \in \mathcal{X}$

$$
\begin{aligned}
\left|f_{h}(s, a)-f_{h}\left(s^{\prime}, a^{\prime}\right)\right| & \leq L \cdot \operatorname{dist}\left((s, a),\left(s^{\prime}, a^{\prime}\right)\right) \text { and } \\
\left|\Delta_{h}(s, a)\right| & \leq \epsilon .
\end{aligned}
$$

A straightforward consequence of Assumption 4 is:

$$
\left|Q_{h}^{\star}(s, a)-Q_{h}^{\star}\left(s^{\prime}, a^{\prime}\right)\right| \leq L \cdot \operatorname{dist}\left((s, a),\left(s^{\prime}, a^{\prime}\right)\right)+2 \epsilon .
$$

The next theorem states that our algorithm, without any modification, is robust to small misspecification error $\epsilon$.

Theorem 3 (Regret bound in the misspecified case). Suppose that Assumption 5 holds. For any $p \in(0,1)$, with probability $1-p$, the total regret of ZOOMRL (see Algorithm 5) is at most

$$
O\left(\left(L+\sqrt{H^{5} \imath}\right) \min _{r_{0} \in(0,1)}\left\{K r_{0}+\sum_{\substack{r=2^{-i} \\ r \geq r_{0}}} \frac{M(r)}{r}\right\}+H^{2}+\sqrt{H^{3} K \imath}+H K \epsilon\right)
$$

where $\tau=\log \left(4 H K^{2} / p\right), M(r)$ is the $r$-packing number of the state-action space.
The Theorem 3 states that ZoomRL incurs at most an extra regret term $O(H K \epsilon)$, comparing to Theorem 2. This term is linear in the number of episodes $K$ as well as the error $\epsilon$. The good news is that our algorithm, without any adaptation, does not break down entirely and it enjoys good guarantees when the optimal $Q$-value is close to a Lipschitz function i.e the error $\epsilon$ is small.

### 5.7 Proof Outline

In this section we outline some key steps in the proof of Theorem 2. All the omitted proofs as well as the analysis of the misspecified case can be found in the appendix. We start by showing two useful properties of our partitioning scheme.

Lemma 3 (Partition's properties). At each step $h \in[H]$ in episode $k \in[k]$, we have

- The domains of balls cover the space-action space i.e $\cup_{B \in \mathcal{B}_{h}^{k}} \operatorname{dom}_{h}^{k}(B)=\mathcal{S} \times \mathcal{A}$.
- For any two balls of radius $r>0$, their centers are at distance at least $r$. In other words, for any $r>0$, the set $\left\{B \in \mathcal{B}_{h^{\prime}}^{k} \operatorname{rad}(B)=r\right\}$ forms an $r$-packing for $\mathcal{S} \times \mathcal{A}$.

We set $\alpha_{t}=\frac{H+1}{H+t}$. This specific choice of learning rate comes from Jin et al. [2018] where they show that this choice is crucial to obtain regret that is not exponential in $H$. We denote $\alpha_{t}^{0}=\prod_{j=1}^{t}\left(1-\alpha_{j}\right)$ and $\alpha_{t}^{i}=\alpha_{i} \prod_{j=i+1}^{t}\left(1-\alpha_{j}\right)$. We have $\alpha_{t}^{0}=0, \forall t \geq 1$ and $\alpha_{t}^{0}=1$ when $t=0$. The lemma below establishes the recursive formula of $Q$-values estimates for the balls.

Lemma 4. At any $(h, k) \in \times[H] \times[K]$ and $B \in \mathcal{B}_{h^{\prime}}^{k}$ let $t=n_{h}^{k}(B)$, and suppose $B$ was previously selected at step $h$ of episodes $k_{1}, k_{2}, \ldots, k_{t}<k$. By the update rule of $\widehat{\mathrm{Q}}$, we have:

$$
\begin{gathered}
\widehat{\mathrm{Q}}_{h}^{k}(B)=\alpha_{t}^{0} \cdot H+\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(r_{h}\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\widehat{\mathrm{V}}_{h+1}^{k_{i}}\left(x_{h+1}^{k_{i}}\right)\right. \\
\left.+u_{i}+2 L \cdot \operatorname{rad}(B)\right)
\end{gathered}
$$

Throughout the learning process, we hope that our estimation $\widehat{\mathrm{Q}}_{h}^{k}$ will get closer to the optimal value $Q_{h}^{\star}$, as $k$ increases while we preserve optimism. Using Azuma-Hoeffding concentration inequality (see Lemma 15 in appendix) together with the Lipschitz assumption, our next lemma shows that $\widehat{Q}_{h}^{k}$ is an upper bound on $Q_{h}^{\star}$ at any episode $k$ with high probability and the difference between $\widehat{\mathrm{Q}}_{h}^{k}$ and $Q_{h}^{\star}$ is controlled by quantities from the next step.

Lemma 5. For any $p \in(0,1)$, we have $\beta_{t}=2 \sum_{i=1}^{t} \alpha_{t}^{i} u_{i} \leq 16 \sqrt{\frac{H^{3} l}{t}}$ and, with probability at least $1-p / 2$, we have that for all $(s, a, h, k) \in \mathcal{S} \times \mathcal{A} \times[H] \times[K]$ and any ball $B \in \mathcal{B}_{h}^{k}$ such that $(s, a) \in \operatorname{dom}_{h}^{k}(B)$ :
(a) $\widehat{Q}_{h}^{k}(B) \geq Q_{h}^{\star}(s, a)$.
(b) $\widehat{Q}_{h}^{k}(B)-Q_{h}^{\star}(x, a) \leq \alpha_{t}^{0} \cdot H+\beta_{t}+4 L \cdot \operatorname{rad}(B)+\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(\widehat{V}_{h+1}^{k_{i}}-V_{h+1}^{*}\right)\left(s_{h+1}^{k_{i}}\right)$.
where $t=n_{h}^{k}(B)$ and $k_{1}, \cdots, k_{t}<k$ are the episodes where $B$ was selected at step $h$.
The next lemma translates the optimism in terms of Q-value estimates to optimism in terms of value function estimates.

Lemma 6 (Optimism). Following the same setting as in Lemma 5, for any ( $h, k$ ), with probability at least $1-p / 2$, we have for any $s \in \mathcal{S}, \widehat{\mathrm{~V}}_{h}^{k}(s) \geq V_{h}^{\star}(s)$.

Proof. Let $s \in \mathcal{S}$, We have $V_{h}^{\star}(s)=Q_{h}^{\star}\left(s, \pi_{h}^{\star}(s)\right)$. As the set of domains of active balls covers the entire space, there exists $B^{\star} \in \mathcal{B}_{h+1}^{k}$ such that $\left(s, \pi_{h}^{\star}(s)\right) \in$ $\operatorname{dom}_{h}^{k}\left(B^{*}\right)$. By the definition of index, we have index ${ }_{h}^{k}\left(B^{\star}\right)=L \cdot \operatorname{rad}\left(B^{\star}\right)+$

$$
\widehat{\mathrm{Q}}_{h}^{k}\left(\tilde{B}^{\star}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right) \text { for some active ball } \tilde{B}^{\star} \text {. We have }
$$

$$
\begin{aligned}
& \widehat{\mathrm{V}}_{h}^{k}(s)-V_{h}^{\star}(s) \\
& =\min \left\{H, \max _{B \in \operatorname{rel}_{h}^{k}(s)} \operatorname{index}_{h}^{k}(B)\right\}-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& \geq \max _{B \in \operatorname{re1_{h}^{k}}(s)} \operatorname{index}_{h}^{k}(B)-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& \geq \operatorname{index}_{h}^{k}\left(B^{\star}\right)-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& =L \cdot \operatorname{rad}\left(B^{\star}\right)+\widehat{Q}_{h}^{k}\left(\tilde{B}^{\star}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& \geq L \cdot \operatorname{rad}\left(B^{\star}\right)+Q_{h}^{\star}\left(s_{\tilde{B}^{\star}}, a_{\tilde{B}^{\star}}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& \geq L \cdot \operatorname{rad}\left(B^{\star}\right)+Q_{h}^{\star}\left(s_{B^{\star}}, a_{B^{\star}}\right)-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& \geq 0
\end{aligned}
$$

Where $\left(s_{B^{\star}}, a_{B^{\star}}\right)$ and $\left(s_{\tilde{B}^{\star}}, a_{\tilde{B}^{\star}}\right)$ denote respectively the centers of balls $B^{\star}$ and $\tilde{B}^{\star}$. The first inequality follows from $Q_{h}^{\star}(s, a) \leq H$ for any state-action pair $(s, a)$. The third inequality follows from lemma 5 . The fourth and the last inequalities follow from Lipschitz assumption 4

### 5.7.1 Regret Analysis

$\pi_{k}$ is the policy executed by the algorithm in step $h$ for $H$ steps to reach the end of the episode. By the optimism of our estimates with respect to the true value function (see lemma 6), we have with probability at least $1-p / 2$

$$
\operatorname{REGRET}(K)=\sum_{k=1}^{K}\left(V_{1}^{*}-V_{1}^{\pi_{k}}\right)\left(x_{1}^{k}\right) \leq \sum_{k=1}^{K}\left(\widehat{\mathrm{~V}}_{1}^{k}-V_{1}^{\pi_{k}}\right)\left(x_{1}^{k}\right)
$$

Denote by $\delta_{h}^{k} \triangleq\left(\widehat{\mathrm{~V}}_{h}^{k}-V_{h}^{\pi_{k}}\right)\left(s_{h}^{k}\right)$ and $\phi_{h}^{k} \triangleq\left(\widehat{\mathrm{~V}}_{h}^{k}-V_{h}^{\star}\right)\left(s_{h}^{k}\right)$. As $V_{h}^{\star} \geq V^{\pi_{k}}$, we have $\phi_{h}^{k} \leq \delta_{h}^{k}$. In the sequel, we aim to upper bound $\delta_{h}^{k}$ as we have $\operatorname{REGRET}(K) \leq$ $\sum_{k=1} \delta_{1}^{k}$.

Let $B_{h}^{k}$ the ball selected at step $h$ of episode $k$ and $B_{\text {init }}$ be the initial ball of radius one that covers the whole space. We denote $B_{h}^{k, \text { pa }}$ the parent of $B_{h}^{k}$. When $B_{h}^{k}$ is the initial ball, we consider that it is parent of itself.

Lemma 7 (Bound on estimation). Let $\hat{P}_{h}$ the empirical transition operator defined as $\left[\hat{P}_{h} V\right]\left(s_{h}^{k}, a_{h}^{k}\right)=V\left(s_{h+1}^{k}\right)$ for any function $V$, any $h$ and $k$. If we denote $\xi_{h+1}^{k}=$
$\left[\left(P_{h}-\hat{P}_{h}\right)\left(V_{h+1}^{\star}-V_{h}^{\pi_{k}}\right)\right]\left(s_{h}^{k}, a_{h}^{k}\right)$, we have

$$
\begin{aligned}
\delta_{h}^{k} \leq H \alpha_{n_{h}^{k}\left(B_{h}^{k}\right)}^{0} \cdot \mathbb{I}_{\left\{B_{h}^{k}=B_{i n i t}\right\}}+\left(11 L+32 \sqrt{H^{3}}\right) & \operatorname{rad}\left(B_{h}^{k}\right) \\
& +\sum_{i=1}^{n_{h}^{k}\left(B_{h}^{k, p a}\right)} \alpha_{n_{h}^{k}\left(B_{h}^{k, p a}\right)}^{i} \phi_{h+1}^{k_{i}\left(\left(_{h}^{k, p a}\right)\right.}-\phi_{h+1}^{k}+\delta_{h+1}^{k}+\xi_{h+1}^{k}
\end{aligned}
$$

where $k_{i}\left(B_{h}^{k, p a}\right)$ is the $i$-th episode where $B_{h}^{k, p a}$ is selected at step $h$.
Taking the sum over $k \in[K]$ of the estimation bound in lemma 7,

$$
\begin{aligned}
\sum_{k=1}^{K} \delta_{h}^{k} & \leq \square+\left(11 L+32 \sqrt{H^{3} \imath}\right) \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right) \\
& +\triangle+\sum_{k=1}^{K}\left(-\phi_{h+1}^{k}+\delta_{h+1}^{k}+\xi_{h+1}^{k}\right)
\end{aligned}
$$

where $\square=H \sum_{k=1}^{K} \alpha_{n_{h}^{k}\left(B_{h}^{k}\right)}^{0} \cdot \mathbb{I}_{\left\{B_{h}^{k}=B_{\text {init }}\right\}}$ and $\triangle=\sum_{k=1}^{K} \sum_{i=1}^{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)} \alpha_{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)} \phi_{h+1}^{k_{i}\left(B_{h}^{k, \text { pa }}\right)}$. For the fist term, we have $\square=H \sum_{k=1}^{K} \mathbb{I}_{\left\{B_{h}^{k}=B_{\text {init }}, n_{h}^{k}\left(B_{\text {init }}\right)=0\right\}}=H$. For the second term $\triangle$, we regroup the summation in a different way. For every $k^{\prime} \in[K]$, the term $\phi_{h+1}^{k^{\prime}}$ appears in the summation with $k>k^{\prime}$ when $B_{h}^{k}$ and $B_{h}^{k^{\prime}}$ share the same parent. The first time it appears we have $n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)=n_{h}^{k^{\prime}}\left(B_{h}^{k^{\prime} \text {, pa }}\right)+1$, the second time it appears we have $n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)=n_{h}^{k^{\prime}}\left(B_{h}^{k^{\prime}, \text { pa }}\right)+2$ and so on. Therefore:

$$
\begin{aligned}
\triangle & \leq \sum_{k^{\prime}=1}^{K} \phi_{h+1}^{k^{\prime}} \sum_{t=n_{h}^{k^{\prime}}\left(B_{h}^{k^{\prime} \text { pa }}\right)+1}^{\infty} \alpha_{t}^{n_{h}^{k^{\prime}}\left(B_{h}^{k^{\prime}, \mathrm{pa}}\right)} \\
& \leq\left(1+\frac{1}{H}\right) \sum_{k=1}^{K} \phi_{h+1}^{k}
\end{aligned}
$$

We use in the last inequality $\sum_{t=i}^{\infty} \alpha_{t}^{i}=1+\frac{1}{H}$ (Lemma 16 in appendix). Therefore, using that $\phi_{h+1}^{k} \leq \delta_{h+1}^{k}$

$$
\begin{aligned}
& \sum_{k=1}^{K} \delta_{h}^{k} \leq H+\left(11 L+32 \sqrt{H^{3} \imath}\right) \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right) \\
& \quad+\left(1+\frac{1}{H}\right) \sum_{k=1}^{K} \phi_{h+1}^{k}+\sum_{k=1}^{K}\left(-\phi_{h+1}^{k}+\delta_{h+1}^{k}+\xi_{h+1}^{k}\right) \\
& \leq H+\left(11 L+32 \sqrt{H^{3} \imath}\right) \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right)+\left(1+\frac{1}{H}\right) \sum_{k=1}^{K} \delta_{h+1}^{k}+\sum_{k=1}^{K} \xi_{h+1}^{k}
\end{aligned}
$$

By unrolling the last inequality for $h \in[H]$ and using the fact $\delta_{H+1}^{k}=0 \quad \forall k \in$ [K], we obtain

$$
\begin{align*}
\sum_{k=1}^{K} \delta_{1}^{k} & \leq \sum_{h=1}^{H}\left(1+\frac{1}{H}\right)^{h-1}\left(H+\left(11 L+32 \sqrt{H^{3} \imath}\right) \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right)+\sum_{k=1}^{K} \xi_{h+1}^{k}\right) \\
& \leq 3 H^{2}+3\left(11 L+32 \sqrt{H^{3}}\right) \sum_{h=1}^{H} \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right)+3 \sum_{h=1}^{H} \sum_{k=1}^{K} \xi_{h+1}^{k} \tag{5.7.1}
\end{align*}
$$

The last inequality follows from the fact $\forall h \in[H],\left(1+\frac{1}{H}\right)^{h-1} \leq\left(1+\frac{1}{H}\right)^{H} \leq$ $\exp (1) \leq 3$.

Now, we proceed to upper bound the two terms $\sum_{h=1}^{H} \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right)$ and $\sum_{h=1}^{H} \sum_{k=1}^{K} \xi_{h+1}^{k}$. Using concentration argument, we show with probability at least $1-p / 2$, we have (see Lemma 12 in appendix)

$$
\begin{equation*}
\sum_{h=1}^{H} \sum_{k=1}^{K} \xi_{h+1}^{k} \leq 4 \sqrt{2 H^{3} K \imath} \tag{5.7.2}
\end{equation*}
$$

Bounding $\sum_{h=1}^{H} \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right)$ : Let's consider all balls of radius $r$ that have been activated at step $h$ throughout the execution of the algorithm. The maximum number of times a ball $B$ of radius $r$ can be selected before it becomes a parent is upper bounded by $\frac{1}{r^{2}}$. After ball $B$ becomes a parent, a new ball of radius $r / 2$ is created every time $B$ is selected. Therefore, we can write the sum over all ball $B \in \mathcal{B}_{h}^{K}$ of radius $r$ as the sum over set of rounds which consists of the round when $B$ was created and all rounds when $B$ was selected before being
a parent. Let $r_{0} \in(0,1)$. we have

$$
\begin{align*}
& \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right)=\sum_{\substack{r=2^{-i}}} \sum_{\substack{B \in \mathcal{B}_{h}^{K} \\
\operatorname{rad}(B)=r}} \sum_{\substack{k \in[K] \\
B_{h}^{K}=B}} r \\
& \quad=\sum_{\substack{r=2^{-i} \\
r<r_{0}}} \sum_{\substack{B \in \mathcal{B}_{h}^{K}}} \sum_{\substack{k \in[K] \\
\operatorname{rad}(B)=r \\
B_{h}^{K}=B}} r+\sum_{\substack{r=2^{-i} \\
r \geq r_{0}}} \sum_{\substack{B \in \mathcal{B}_{h}^{K} \\
\operatorname{rad}(B)=r}} \sum_{\substack{k \in[K] \\
B_{h}^{K}=B}} r \\
& \leq K r_{0}+\sum_{\substack{r=2^{-i} \\
r \geq r_{0}}} \sum_{\substack{B \in \mathcal{B}_{h}^{K} \\
\operatorname{rad}(B)=r}}^{r} \frac{r}{r^{2}}+2 r \\
& \\
& \leq K r_{0}+3 \sum_{\substack{r=2^{-i} \\
r \geq r_{0}}}\left|\left\{B \in \mathcal{B}_{h}^{K}: \operatorname{rad}(B)=r\right\}\right| \frac{1}{r}  \tag{5.7.3}\\
& \leq K r_{0}+3 \sum_{\substack{r=2^{-i} \\
r \geq r_{0}}} \frac{M(r)}{r}
\end{align*}
$$

The last step follows from lemma 3: The set of active balls of radius $r$ is a $r$ packing of $\mathcal{S} \times \mathcal{A}$. Thus, $\left|\left\{B \in \mathcal{B}_{h}^{K}: \operatorname{rad}(B)=r\right\}\right| \leq M(r)$ where $M(r)$ is the $r$-covering number.

Plugging bounds (5.7.2) and (5.7.3) in (5.7.1) and using the union bound, we obtain the desired regret bound in Theorem 2

### 5.8 CONCLUSION

In this paper, we present ZOOMRL, a provably efficient model-free reinforcement learning algorithm in continuous state-action spaces under the assumption that the true optimal action-value function is Lipscthiz with respect to similarity metric between state-action pairs. Our algorithm takes into account the geometry of the action-value function by allocating more attention to relevant regions. We show that our method achieves sublinear regret that depends on the packing number of the state-action space and that it is robust to small misspecification errors.

Our method requires the knowledge of the Lipschitz constant $L$ as well as the metric dist to achieve its performance. A natural future question is whether an RL algorithm can be proved to be efficient without knowing $L$ or dist in advance.

# Online Learning in Non-stationary Linear Markov Decision Processes 

### 6.1 PROLOGUE TO THE CONTRIBUTION

### 6.1.1 Article Details

This chapter is based on the article «Efficient Learning in Non-Stationary Linear Markov Decision Processes» [Touati and Vincent, 2020a], joint work with Pascal Vincent. I am the first author. I performed the derivation of the new algorithm and its theoretical analysis.

### 6.1.2 Context

In order to design both provable sample-efficient and computationally efficient RL algorithms for large-scale problems, an appealing challenge is to combine exploration strategies with generalization methods. In the previous chapter, we ensured generalization over states by aggregating them into a finite set of meta-states and run tabular exploration mechanism on the latter. In particular, we proposed to actively explore the state-action space by learning on-the-fly an adaptive partitioning that takes into account the shape of the optimal value function. When the state-action space is assumed to be a compact metric space, we showed that such adaptive discretization-based algorithm yield sublinear regret but suffer from the curse of dimensionality as their regret scales almost exponentially with the covering dimension of the whole space.

Another structural assumption, that received attention in the recent literature [Yang and Wang, 2019, Jin et al., 2020b, Zanette et al., 2020a], is when both reward and transition dynamics are linear functions with respect to a given feature mapping. This assumption enables the design of efficient algorithms with a linear representation of the action-value function.

Most prior algorithms with linear function approximation assume that the environment is stationary and minimize the regret over the best fixed policy. However, in many problems of interest, we are faced with a changing world, in some cases with substantial non-stationarity. This is a more realistic and challenging setting, since what has been learned in the past may be obsolete in the present.

### 6.1.3 Paper Abstract

We study episodic reinforcement learning in non-stationary linear (a.k.a. low-rank) Markov Decision Processes (MDPs), i.e, both the reward and transition kernel are linear with respect to a given feature map and are allowed to evolve either slowly or abruptly over time. For this problem setting, we propose OPT-WLSVI an optimistic model-free algorithm based on weighted least squares value iteration which uses exponential weights to smoothly forget data that are far in the past. We show that our algorithm, when competing against the best policy at each time, achieves a regret that is upper bounded by $\widetilde{\mathcal{O}}\left(d^{5 / 4} H^{2} \Delta^{1 / 4} K^{3 / 4}\right)$ where $d$ is the dimension of the feature space, $H$ is the planning horizon, $K$ is the number of episodes and $\Delta$ is a suitable measure of non-stationarity of the MDP. Moreover, we point out technical gaps in the study of forgetting strategies in non-stationary linear bandits setting made by previous works and we propose a fix to their regret analysis.

### 6.1.4 Recent Developments

The authors of Cheung et al. [2019] who pioneered the non-stationary linear bandit, released recently a revised version of their AISTATS 2019 paper to acknowledge the mistake in the analysis. In order for their optimal rate $\tilde{\mathcal{O}}\left(T^{2 / 3}\right)$ to hold, they assume that actions are orthogonal.

Zhao and Zhang [2021] identify also the mistake and proposed the same fix than ours in a technical note released slightly after we had made public on arxiv a version of our paper containing the fix.

Faury et al. [2021] build on our proposed fix to provide a correct regret analysis for non-stationary generalized linear bandit.

While our work shows that forgetting strategies achieve the rate of $\tilde{\mathcal{O}}\left(T^{3 / 4}\right)$, the recent work Wei and Luo [2021] follows a substantially different approach to achieve the optimal rate $\tilde{\mathcal{O}}\left(T^{2 / 3}\right)$ for the first time in the setting of non-stationary linear bandit and MDP. Their algorithm detects non-stationarity by running multiple instances of a base (stationary) algorithm with different durations in a randomized schedule.

### 6.2 INTRODUCTION

In the present work, we study the problem of online learning in episodic nonstationary linear Markov Decision Processes (MDP), where both the reward and transition kernel are linear with respect to a given feature map and are allowed to evolve dynamically and even adversarially over time. The interaction of the
agent with the environment is divided into $K$ episodes of fixed length $H$. Moreover, we assume that the total change of the MDP, that we measure by a suitable metric, over the $K$ episodes is upper bounded by $\Delta$, called variation budget.

To address this problem, we propose a computationally efficient model-free algorithm, that we call OPT-WLSVI. We prove that, in the setting described above, its regret when competing against the best policy for each episode is at most $\widetilde{\mathcal{O}}\left(d^{5 / 4} H^{2} \Delta^{1 / 4} K^{3 / 4}\right)$. Concurrently to our work, Zhou et al. [2020] propose to periodically restart LSVI-UCB from scratch, achieving the same regret. By contrast, our algorithm is based on weighted least squares value iteration that uses exponential weights to smoothly forget data that are far in the past, which drives the agent to keep exploring to discover changes. Our approach is motivated by the recent work of Russac et al. [2019] who establish a new deviation inequality to sequential weighted least squares estimator and apply it to the non-stationary stochastic linear bandit problem. However, in contrast to linear bandit, our algorithm handles the additional problem of credit assignment since future states depend in non-trivial way on the agent's policy and thus we need to carefully control how errors are propagated through iterations. Moreover, we discovered technical errors in the regret analysis of forgetting strategies in non-stationary linear bandits made by previous works and we propose a correction.

### 6.3 Problem Statement

### 6.3.1 Notation

Throughout the paper, all vectors are column vectors. We denote by $\|\cdot\|$ the Euclidean norm for vectors and the operator norm for matrices. For positive definite matrix $A$, we use $\|x\|_{A}$ to denote the matrix norm $\sqrt{x^{\top} A x}$. We define $[N]$ to be the set $\{1,2, \ldots, N\}$ for any positive integer $N$.

### 6.3.2 Non-Stationary Reinforcement Learning and Dynamic Regret

We consider a non-stationary undiscounted finite-horizon MDP $(\mathcal{S}, \mathcal{A}, P, r, H)$ where $\mathcal{S}$ and $\mathcal{A}$ are the state and action space, $H$ is the planning horizon i.e number of steps in each episode, $P=\left\{P_{t, h}\right\}_{t>0, h \in[H]}$ and $r=\left\{r_{t, h}\right\}_{t>0, h \in[H]}$ are collections of transition kernels and reward functions, respectively. More precisely, when taking action $a$ in state $s$ at step $h$ of the $t$-th episode, the agent receives a reward $r_{t, h}(s, a)$ and makes a transition to the next state according to the probability measure $P_{t, h}(\cdot \mid s, a)$.

For any step $h \in[H]$ of an episode $t$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$, the stateaction value function of a policy $\pi=\left(\pi_{1}, \ldots, \pi_{H}\right)$ is defined as $Q_{t, h}^{\pi}(s, a)=$ $r_{t, h}(s, a)+\mathbb{E}\left[\sum_{i=h+1}^{H} r_{t, i}\left(s_{i}, \pi_{i}\left(s_{i}\right)\right) \mid s_{h}=s, a_{h}=a\right]$, and the value function is $V_{t, h}^{\pi}(s)=Q_{t, h}^{\pi}\left(s, \pi_{h}(s)\right)$. The optimal value and action-value functions are defined as $V_{t, h}^{\star}(x) \triangleq \max _{\pi} V_{t, h}^{\pi}(s)$ and $Q_{t, h}^{\star}(s, a) \triangleq \max _{\pi} Q_{t, h}^{\pi}(s, a)$. If we denote $\left[P_{t, h} V_{t, h+1}\right](s, a)=\mathbb{E}_{s^{\prime} \sim P_{t, h}(\cdot \mid s, a)}\left[V_{t, h+1}\left(s^{\prime}\right)\right]$, both $Q^{\pi}$ and $Q^{\star}$ can be conveniently written as the result of the following Bellman equations

$$
\begin{align*}
& Q_{t, h}^{\pi}(s, a)=r_{t, h}(s, a)+\left[P_{t, h} V_{t, h+1}^{\pi}\right](s, a),  \tag{6.3.1}\\
& Q_{t, h}^{\star}(s, a)=r_{t, h}(s, a)+\left[P_{t, h} V_{t, h+1}^{\star}\right](s, a), \tag{6.3.2}
\end{align*}
$$

where $V_{t, H+1}^{\pi}(s)=V_{t, H+1}^{\star}(s)=0$ and $V_{t, h}^{\star}(s)=\max _{a \in \mathcal{A}} Q_{t, h}^{\star}(s, a)$, for all $s \in \mathcal{S}$.
Learning problem: We focus on the online episodic reinforcement learning setting in which the rewards and the transition kernels are unknown. The learning agent plays the game for $K$ episodes $t=1, \ldots, K$, where each episode $t$ starts from some initial state $s_{t, 1}$ sampled according to some initial distribution. The agent controls the system by choosing a policy $\pi_{t}$ at the beginning of the $t$-th episode. We measure the agent's performance by the dynamic regret, defined as the sum over all episodes of the difference between the optimal value function in episode $t$ and the value of $\pi_{t}$ :

$$
\operatorname{REGRET}(K)=\sum_{t=1}^{K} V_{t, 1}^{*}\left(s_{t, 1}\right)-V_{t, 1}^{\pi_{t}}\left(s_{t, 1}\right)
$$

### 6.3.3 Linear Markov Decision Processes

In this work, we consider a special class of MDPs called linear MDPs, where both reward function and transition kernel can be represented as a linear function of a given feature mapping $\phi: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}$. Now we present our main assumption

Assumption 6 (Non-stationary linear MDP). $(\mathcal{S}, \mathcal{A}, P, r, H)$ is non-stationary linear MDP with a feature map $\phi: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}$ if for any $(t, h) \in \mathbb{N} \times[H]$, there exist d unknown (signed) measures $\boldsymbol{\mu}_{t, h}=\left(\mu_{t, h}^{(1)}, \ldots, \mu_{t, h}^{(d)}\right)$ over $\mathcal{S}$ and an unknown vector $\boldsymbol{\theta}_{t, h} \in \mathbb{R}^{d}$, such that for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$
\begin{align*}
P_{t, h}(\cdot \mid s, a) & =\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\mu}_{t, h}(\cdot),  \tag{6.3.3}\\
r_{t, h}(s, a) & =\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\theta}_{t, h} . \tag{6.3.4}
\end{align*}
$$

Without loss of generality, we assume ${ }^{1}\|\phi(s, a)\| \leq 1$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, and $\max \left\{\left\|\boldsymbol{\mu}_{t, h}(\mathcal{S})\right\|,\left\|\boldsymbol{\theta}_{t, h}\right\|\right\} \leq \sqrt{d}$ for all $(t, h) \in \mathbb{N} \times[H]$.

Linear MDPs are also know as low-rank MDPs [Zanette et al., 2020a]. In fact, in the case of finite state and action spaces with cardinalities $|\mathcal{S}|$ and $|\mathcal{A}|$ respectively, the transition matrix $P \in \mathbb{R}^{(|\mathcal{S}| \times|\mathcal{A}|) \times|\mathcal{S}|}$ could be expressed by the following low-rank factorization for any $(t, h)$ :

$$
P_{t, h}=\boldsymbol{\Phi} \boldsymbol{M}_{t, h}
$$

where $\boldsymbol{\Phi} \in \mathbb{R}^{(|\mathcal{S}| \times|\mathcal{A}|) \times d}$ such as $\boldsymbol{\Phi}[(s, a),:]=\phi(s, a)^{\top}$ and $\boldsymbol{M}_{t, h} \in \mathbb{R}^{d \times|\mathcal{S}|}$ such that $\boldsymbol{M}_{t, h}[:, s]=\boldsymbol{\mu}_{t, h}(s)$ a discrete measure. Therefore the rank of the matrix $P$ is at most $d$.

An important consequence of Assumption 6 is that the $Q$-function of any policy is linear in the features $\phi$.
Lemma 8. For every policy $\pi$ and any $(t, h) \in \mathbb{N}^{\star} \times[H]$ there exists $\mathbf{w}_{t, h}^{\pi} \in \mathbb{R}^{d}$ such that

$$
\begin{equation*}
Q_{t, h}^{\pi}=\phi(s, a)^{\top} \mathbf{w}_{t, h^{\prime}}^{\pi} \quad \forall(s, a) \in \mathcal{S} \times \mathcal{A} . \tag{6.3.5}
\end{equation*}
$$

### 6.4 The Proposed Algorithm

Algorithm 6, referred as OPT-WLSVI (OPTimistic Weighted Least Squares Value Iteration), parametrizes the Q -values $Q_{t, h}(s, a)$ by a linear form $\phi(s, a)^{\top} \mathbf{w}_{t, h}$ and updates the parameters $\mathbf{w}_{t, h}$ by solving the following regularized weighted least squares problem:

$$
\mathbf{w}_{t, h}=\arg \min _{\mathbf{w} \in \mathbb{R}^{d}}\left\{\sum_{\tau=1}^{t-1} \eta^{-\tau}\left(r_{\tau, h}+V_{t, h+1}\left(s_{\tau, h+1}\right)-\boldsymbol{\phi}_{\tau, h}^{\top} \mathbf{w}\right)^{2}+\lambda \eta^{-(t-1)}\|\mathbf{w}\|^{2}\right\}
$$

where $\eta \in(0,1)$ is a discount factor, $V_{t, h+1}\left(s_{t, h+1}\right)=\max _{a \in \mathcal{A}} Q_{t, h+1}\left(s_{\tau, h+1}, a\right)$ and $r_{\tau, h}$ and $\phi_{\tau, h}$ are shorthand for $r_{\tau, h}\left(s_{\tau, h}, a_{\tau, h}\right)$ and $\phi\left(s_{\tau, h}, a_{\tau, h}\right)$ respectively. The discount factor $\eta$ plays an important role as it gives exponentially increasing weights to recent transitions, hence, the past is smoothly forgotten.

The regularized weighted least-squares estimator of the above problem can be written in closed form

$$
\begin{equation*}
\mathbf{w}_{t, h}=\Sigma_{t, h}^{-1}\left(\sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h}\left(r_{\tau, h}+V_{t, h+1}\left(s_{\tau, h+1}\right)\right)\right) \tag{6.4.1}
\end{equation*}
$$

[^5]where $\Sigma_{t, h}=\sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top}+\lambda \eta^{-(t-1)} \cdot \mathbf{I}$ is the Gram matrix. We further define the matrix
\[

$$
\begin{equation*}
\widetilde{\Sigma}_{t, h}=\sum_{\tau=1}^{t-1} \eta^{-2 \tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top}+\lambda \eta^{-2(t-1)} \cdot \mathbf{I} \tag{6.4.2}
\end{equation*}
$$

\]

The matrix $\widetilde{\Sigma}_{t, h}$ is connected to the variance of the estimator $w_{t, h}$, which involves the squares of the weights $\left\{\eta^{-2 \tau}\right\}_{\tau \geq 0}$. OPT-WLSVI uses both matrices $\Sigma_{t, h}$ and $\widetilde{\Sigma}_{t, h}$ to define a upper confidence bound (UCB) term $\beta\left(\phi^{\top} \Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1} \phi\right)^{1 / 2}$ to encourage exploration, where $\beta$ is a scalar.

The algorithm proceeds as follows. At the beginning of episode $t$, OPT-WLSVI estimates the weighted least square estimator $w_{t, h}$ for each step $h \in[H]$ as given by Equation (6.4.1). Then, the algorithm updates the $Q$-value and the value function estimates as follows:

$$
\begin{aligned}
Q_{t, h}(\cdot, \cdot) & =\phi(\cdot, \cdot)^{\top} \mathbf{w}_{t, h}+\beta\left(\phi(\cdot, \cdot)^{\top} \Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1} \phi(\cdot, \cdot)\right)^{1 / 2} \\
V_{t, h}(\cdot) & =\min \left\{\max _{a \in \mathcal{A}} Q_{t, h}(\cdot, s), H\right\}
\end{aligned}
$$

The UCB term is used to bound the estimation error of the value function, due to an insufficient number of samples, with high probability. The clipping of the value estimate is here to keep $V_{t, h}$ within the range of plausible values while preserving the optimism as $H$ is an upper bound on the true optimal value function. Finally the algorithm collects a new trajectory by following the greedy policy $\pi_{t}$ with respect to the estimated $Q$-values.

Computational complexity: At each step $h \in[H]$ of an episode $t \in[K]$, we need to compute the inverse of $\Sigma_{t, h}$ to solve the weighted least-squares problem. A naive implementation requires $\mathcal{O}\left(d^{3}\right)$ elementary operations, but as $\Sigma_{t, h}$ is essentially a sum of rank-one matrices, we need only $\mathcal{O}\left(d^{2}\right)$ using the ShermanMorrison update formula. Furthermore, $\mathcal{O}\left(d^{2}\right)$ operations are needed to compute the exploration bonus $\left(\phi^{\top} \Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1} \phi\right)^{1 / 2}$ that can be computed using only matrix-vector multiplications. Therefore computing $V_{t, h+1}$ for all the past successor states requires $\mathcal{O}\left(d^{2}|\mathcal{A}| K\right)$ (the $|\mathcal{A}|$ factor is due to the maximization over actions). As we need to do this at all steps and for every episode, the overall computation complexity of our algorithm is $\mathcal{O}\left(d^{2}|\mathcal{A}| H K^{2}\right)$.

### 6.5 NON-STATIONARY LINEAR BANDITS

Before providing the analysis of OPT-WLSVI, we start by examining the linear bandit case when the horizon $H=1$. Let us first recall the non-stationary linear bandit model

```
Algorithm 6 Optimistic Weighted Least-Squares Value Iteration (OPT-WLSVI)
    for episode \(t=1, \ldots, K\) do
        Receive the initial state \(s_{t, 1}\).
        /* Run LSVI procedure
        \(V_{t, H+1}(\cdot) \leftarrow 0\)
        for step \(h=H, \ldots, 1\) do
            \(\mathbf{w}_{t, h} \leftarrow \Sigma_{t, h}^{-1}\left(\sum_{\tau=1}^{k-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h}\left(r_{\tau, h}+V_{t, h+1}\left(s_{\tau, h+1}\right)\right)\right)\)
            \(Q_{t, h}(\cdot, \cdot) \leftarrow \phi(\cdot, \cdot)^{\top} \mathbf{w}_{t, h}+\beta\left(\phi(\cdot, \cdot)^{\top} \Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1} \phi(\cdot, \cdot)\right)^{1 / 2}\)
            \(V_{t, h}(\cdot) \leftarrow \min \left\{\max _{a \in \mathcal{A}} Q_{t, h}(\cdot, s), H\right\}\)
        end for
        end for
        /* Execute greedy policy
        for step \(h=1, \ldots, H\) do
            Execute \(a_{t, h}=\operatorname{argmax}_{a \in \mathcal{A}} Q_{t, h}\left(s_{t, h}, a\right)\)
            receive \(r_{t, h}\) and observe \(s_{t, h+1}\)
            /* Update matrices
            \(\Sigma_{t+1, h} \leftarrow \Sigma_{t, h}+\eta^{-t} \boldsymbol{\phi}_{t, h} \boldsymbol{\phi}_{t, h}^{\top}+\lambda \eta^{-t}(1-\eta) \cdot \mathbf{I}\)
            \(\widetilde{\Sigma}_{t+1, h} \leftarrow \widetilde{\Sigma}_{t, h}+\eta^{-2 t} \boldsymbol{\phi}_{t, h} \boldsymbol{\phi}_{t, h}^{\top}+\lambda \eta^{-2 t}\left(1-\eta^{2}\right) \cdot \mathbf{I}\)
        end for
    end for
    end for
```

Definition 1 (Non-stationary linear bandit). Let $\mathcal{X} \subset \mathbb{R}^{d}$ a set of decisions. At iteration $t$, the player makes a decision $x_{t}$ from a subset set $\mathcal{X}_{t} \subset \mathcal{X}$, then observes the reward $r_{t}$ satisfying:

$$
\begin{equation*}
r_{t}=x_{t}^{\top} \boldsymbol{\theta}_{t}+z_{t} \tag{6.5.1}
\end{equation*}
$$

where $\boldsymbol{\theta}_{t}$ is the unknown regression parameter at iteration $t$ and $z_{t}$ is conditionally $\sigma$ subgaussian noise. We assume further that $\|x\| \leq 1, \forall x \in \mathcal{X}$ and $\left\|\boldsymbol{\theta}_{t}\right\| \leq S, \forall t$.

When $H=1$ linear MDP reduces to linear bandit if we let $\mathcal{X}=\{\phi(s, a), a \in$ $\mathcal{A}, s \in \mathcal{S}\}$ and $\mathcal{X}_{t}=\left\{\phi\left(s_{t}, a\right), a \in \mathcal{A}\right\}$ where $s_{t}$ is sampled from a given fixed distribution over states.

For the bandit setting, forgetting strategies have been proposed such as sliding-window, weighted regression and restarting in [Cheung et al., 2019], [Russac et al., 2019] and Zhao et al. [2020] respectively. Randomized exploration with weighting strategy has also been introduced in Kim and Tewari [2020]. The aforementioned works provide a regret of $\tilde{\mathcal{O}}\left(d^{2 / 3} \Delta^{1 / 2} K^{2 / 3}\right)$ which is optimal as it matches the established lower bound $\Omega\left(d^{2 / 3} \Delta^{1 / 3} K^{2 / 3}\right)$ in [Cheung et al., 2019] up to $\log (K)$ factors. Unfortunately, we find technical gaps in the regret analysis provided by the earliest paper [Cheung
et al., 2019], which were then reproduced by the other three papers. Specifically Cheung et al. [2019] attempted, in their Lemma 1, to upper bound the non-stationarity bias of the reward parameters by controlling the eigenvalues of matrix $M=V_{t}^{-1} \sum_{\tau=t-W}^{p} x_{\tau} x_{\tau}^{\top}$, where $V_{t}=\sum_{\tau=1}^{t-1} x_{\tau} x_{\tau}^{\top}+\lambda \cdot \mathbf{I}$ is the Gram matrix and for any integer $p \in\{t-W, \ldots, t-1\}$. They then needed to prove that $M$ is positive semi-definite, but their argument has technical errors. We precise in the appendix the issue in their argument and we provide concrete counter-examples.

Now, we provide a fix to the original error in the regret analysis of SW-UCB algorithm in Cheung et al. [2019] (see also Appendix D. 2 for the analysis of DLINUCB algorithm proposed by Russac et al. [2019]). At time $t$, SW-UCB selects a decision as follows:

$$
\begin{equation*}
x_{t}=\arg \max _{x \in \mathcal{X}_{t}} x^{\top} \hat{\boldsymbol{\theta}}_{t}+\beta\|x\|_{V_{t}^{-1}} \tag{6.5.2}
\end{equation*}
$$

where $\hat{\boldsymbol{\theta}}_{t}=V_{t}^{-1} \sum_{\tau=\max \{1, t-W\}}^{t-1} x_{\tau} r_{\tau}$ is the solution of the sliding window least squares problem

In their Lemma 1, Cheung et al. [2019] attempts to control the nonstationarity bias $\left\|\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right\|$ where $\overline{\boldsymbol{\theta}}_{t} \triangleq V_{t}^{-1} \sum_{\tau=\max \{1, t-W\}}^{t-1} A_{\tau} A_{\tau}^{\top} \boldsymbol{\theta}_{\tau}+\lambda \boldsymbol{\theta}_{t}$ is the average of the true regression parameters over the sliding window. We propose to control $\left|x^{\top}\left(\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right)\right|$ for any $x \in \mathcal{X}$ and then use the fact that $\left\|\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right\|=\max _{x:\|x\|=1}\left|x^{\top}\left(\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right)\right|$

$$
\begin{aligned}
& \left|x^{\top}\left(\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right)\right| \\
& =\left|x^{\top} V_{t}^{-1} \sum_{\tau=\max \{1, t-W\}}^{t-1} x_{\tau} x_{\tau}^{\top}\left(\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right)\right| \\
& \leq \sum_{\tau=\max \{1, t-W\}}^{t-1}\left|x^{\top} V_{t}^{-1} x_{\tau}\right| \cdot\left|x_{\tau}^{\top}\left(\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right)\right| \\
& =\sum_{\tau=\max \{1, t-W\}}^{t-1}\left|x^{\top} V_{t}^{-1} x_{\tau}\right| \cdot\left|x_{\tau}^{\top}\left(\sum_{s=\tau}^{t-1}\left(\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right)\right)\right| \\
& \leq \sum_{\tau=\max \{1, t-W\}}^{t-1}\left|x^{\top} V_{t}^{-1} x_{\tau}\right| \cdot\left\|x_{\tau}\right\| \cdot\left\|\sum_{s=\tau}^{t-1}\left(\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right)\right\| \\
& \leq \sum_{\tau=\max \{1, t-W\}}^{t-1}\left|x^{\top} V_{t}^{-1} x_{\tau}\right| \cdot \sum_{s=\tau}^{t-1}\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\| \\
& \text { (Cauchy-Schwarz) } \\
& \leq \quad\left(\left\|x_{\tau}\right\| \leq 1\right)
\end{aligned}
$$

$$
\begin{align*}
& \leq \sum_{s=\max \{1, t-W\}}^{t-1} \sum_{\tau=\max \{1, t-W\}}^{s}\left|x^{\top} V_{t}^{-1} x_{\tau}\right| \cdot\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\| \\
& \leq \sum_{s=\max \{1, t-W\}}^{t-1} \sqrt{\left[\sum_{\tau=\max \{1, t-W\}}^{t-1} \sum_{s=\tau}^{t-1}=\sum_{s=\max \{1, t-W\}}^{t-1} \sum_{\tau=\max \{1, t-W\}}^{s}\right)} \\
& \left.\leq \sum_{s=\max \{1, t-W\}}^{s-1} x^{\top} V_{t}^{-1} x\right] \cdot\left[\sum_{\tau=\max \{1, t-W\}}^{s} x_{\tau}^{\top} V_{t}^{-1} x_{\tau}\right] \\
& {\left[\sum_{\tau=\max \{1, t-W\}}^{t-1} \sum^{s} \boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1} \|\right.} \\
& \left.\leq\|x\| \sqrt{d} V_{t}^{-1} x\right] \cdot d
\end{align*}\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\| \quad((\star))
$$

where the inequality $(\star)$ follows from the fact that $\sum_{\tau=\max \{1, t-W\}}^{s} x_{\tau}^{\top} V_{t}^{-1} x_{\tau} \leq$ $d$ that can be proved as follows. We have $\sum_{\tau=\max \{1, t-W\}}^{t-1} x_{\tau}^{\top} V_{t}^{-1} x_{\tau}=$ $\sum_{\tau=\max \{1, t-W\}}^{t-1} \operatorname{tr}\left(x_{\tau}^{\top} V_{t}^{-1} x_{\tau}\right)=\operatorname{tr}\left(V_{t}^{-1} \sum_{\tau=\max \{1, t-W\}}^{t-1} x_{\tau} x_{\tau}^{\top}\right)$. Given the eigenvalue decomposition $\sum_{\tau=\max \{1, t-W\}}^{t-1} x_{\tau} x_{\tau}=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{d}\right)^{\top}$, we have $V_{t}=\operatorname{diag}\left(\lambda_{1}+\lambda, \ldots, \lambda_{d}+\lambda\right)^{\top}$, and $\operatorname{tr}\left(V_{t}^{-1} \sum_{\tau=1}^{t-1} x_{\tau} x_{\tau}^{\top}\right)=\sum_{i=1}^{d} \frac{\lambda_{j}}{\lambda_{j}+\lambda} \leq d$.

Comparing to the bound on $\left\|\overline{\boldsymbol{\theta}}_{t}-\boldsymbol{\theta}_{t}\right\|$ in the Lemma 1 of Cheung et al. [2019], there is an extra factor $\sqrt{\frac{d W}{\lambda}}$ that multiplies the local non-stationarity term $\sum_{s=t-W}^{t-1}\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\|$. This extra factor will consequently multiply the variation budget term in the final regret, as stated in the following proposition:

Proposition 5. Under the assumption that $\sum_{t=1}^{K-1}\left\|\boldsymbol{\theta}_{t}-\boldsymbol{\theta}_{t+1}\right\| \leq \Delta$, for any $\delta \in(0,1)$, if we set $\beta=\sqrt{\lambda} S+\sigma \sqrt{2 \log (K / \delta)+d \log \left(1+\frac{W}{\lambda d}\right)}$ in the algorithm 1 SW-UCB of Cheung et al. [2019], then with probability $1-\delta$, the dynamic regret of SW-UCB is at most

$$
\mathcal{O}\left(\sqrt{\frac{d W}{\lambda}} \Delta W+\beta \sqrt{d K} \sqrt{\lceil K / W\rceil} \sqrt{\log \left(1+\frac{W}{d \lambda}\right)}\right)
$$

Comparing to the regret upper bound in Theorem 3 of Cheung et al. [2019], our fix leads to an extra factor $\sqrt{d W}$ multiplying the variation budget in, which becomes now $\tilde{\mathcal{O}}\left(d^{1 / 2} \Delta W^{3 / 2}+d K W^{-1 / 2}\right)$. Optimizing over the sliding window
size $W$ leads to a final dynamic regret of $\tilde{\mathcal{O}}\left(d^{7 / 8} \Delta^{1 / 4} K^{3 / 4}\right)$. Note that the latter is not optimal since it does not match the lower bound $\Omega\left(d^{2 / 3} \Delta^{1 / 3} K^{2 / 3}\right)$. This leaves the question of whether or not forgetting strategies are optimal to handle non-stationarity in linear bandits as an open research problem.

### 6.6 THEORETICAL GUARANTEE OF OPT-WLSVI

In this section, we present our main theoretical result which is an upper bound on the dynamic regret of OPT-WLSVI (see Algorithm 6). First, we quantify the variations on reward function and transition kernel over time in terms of their respective variation budgets $\Delta_{r}$ and $\Delta_{P}$. The main advantage of using the variation budget is that it accounts for both slowly-varying and abruptlychanging MDPs.

Definition 2 (MDP Variation budget). We define $\Delta=\Delta_{r}+\Delta_{P}$ where

$$
\begin{aligned}
& \Delta_{r} \triangleq \sum_{t=1}^{K} \sum_{h=1}^{H}\left\|\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{t+1, h}\right\| \\
& \Delta_{P} \triangleq \sum_{t=1}^{K} \sum_{h=1}^{H}\left\|\boldsymbol{\mu}_{t, h}(\mathcal{S})-\boldsymbol{\mu}_{t+1, h}(\mathcal{S})\right\|
\end{aligned}
$$

A similar notion has already been proposed in the literature, for instance total variance distance between $P_{t, h}$ and $P_{t+1, h}$ in tabular MDPs [Ortner et al., 2019, Cheung et al., 2020] or Wasserstein distance in smooth MDPs [Domingues et al., 2020a].

Now we present our bound on the dynamic regret for OPT-WLSVI.
Theorem 4 (Regret Bound). Under Assumption 6, there exists an absolute constant $c>0$ such that, for any fixed $\delta \in(0,1)$, if we set $\lambda=1$ and $\beta=c \cdot d H \sqrt{1}$ in Algorithm 6 with $\imath \triangleq \log \left(\frac{2 d H}{\delta(1-\eta)}\right)$, then with probability $1-\delta$, for any $W>0$ the dynamic regret of OPT-WLSVI is at most

$$
\begin{gather*}
\mathcal{O}\left(c d^{3 / 2} H \sqrt{K \imath} \sqrt{2 K \log (1 / \eta)+2 \log \left(1+\frac{1}{d \lambda(1-\eta)}\right)}+\right. \\
H^{3 / 2} \sqrt{K \imath}+\underbrace{\sqrt{\frac{d}{\lambda(1-\eta)}} H W \Delta+\frac{H^{2} K \sqrt{d}}{\lambda} \frac{\eta^{W}}{1-\eta}}_{\text {non-stationarity bias }}) \tag{6.6.1}
\end{gather*}
$$

where $\mathcal{O}(\cdot)$ hides only absolute constants.

| Linear | Stationary | Non-stationary |
| :---: | :---: | :---: |
| Bandits | $\widetilde{\mathcal{O}}\left(d K^{1 / 2}\right)$ | $\widetilde{\mathcal{O}}\left(d^{7 / 8} \Delta^{1 / 4} K^{3 / 4}\right)$ |
|  | [Abbasi-Yadkori et al., 2011] | Cheung et al. [2019] |
| MDPs | $\widetilde{\mathcal{O}}\left(d^{3 / 2} H^{2} K^{1 / 2}\right)[$ Jin et al., 2020b] | Russac et al. [2019] and our work |
|  | $\widetilde{\mathcal{O}}\left(d H^{2} K^{1 / 2}\right)$ [Zanette et al., 2020b] | $\widetilde{\mathcal{O}}\left(d^{5 / 4} H^{2} \Delta^{1 / 4} K^{3 / 4}\right)$ |

Table 6.1: Comparison of our regret bound with state-of-the-art bounds for both linear bandits and linear MDPs. $d$ is the dimension of the features space, $H$ is the planning horizon of the MDP, $K$ is the number of episodes and $\Delta$ is the variation budget. When we go from a bandit setting to MDPs, the work of Jin et al. [2020b] in the stationary case and our work in the nonstationary case incur an extra $d^{1 / 2}$ factor and $d^{3 / 8}$ respectively. Zanette et al. [2020b] achieve a linear dependence on $d$ in the stationary case but their proposed algorithm is computationally intractable.

The last two terms of the of Equation (6.6.1) are the result of the bias due to the non-stationarity of the MDP. In theorem 4 we introduce the parameter $W$ that can be interpreted, at a high level, as the effective temporal window equivalent to a particular choice of discount factor $\eta$ : the bias resulting from transitions that are within the window $W$ may be bounded in term of $W$ while the remaining ones are bounded globally by the last term of Equation (6.6.1).

The following corollary shows that by optimizing the parameters $W$ and $\eta$, our algorithm achieves a sublinear regret.

Corollary 1. If we set $\log (1 / \eta)=\left(\frac{\Delta}{d K}\right)^{1 / 2}$ and $W=\frac{\log (K /(1-\eta))}{\log (1 / \eta)}$; under the same assumptions as in Theorem 4 , for any $\delta \in(0,1)$, we have that with probability $1-\delta$, the dynamic regret of OPT-WLSVI is at most $\widetilde{\mathcal{O}}\left(d^{5 / 4} H^{2} \Delta^{1 / 4} K^{3 / 4}\right)$ where $\widetilde{\mathcal{O}}(\cdot)$ hides logarithmic factors.

In Corollary 1, we rely on the knowledge of the variation budget $\Delta$ (or at least an upper bound on $\Delta$ ) in order to achieve a sublinear regret. We show in the next section how to relax the requirement of knowing the variation budget. In particular, we will describe how to extend our algorithm, using the Bandit-over-Reinforcement-Learning framework [Cheung et al., 2020] in order to deal with an unknown variation budget.

### 6.6.1 Unknown variation budget

Our algorithm OPT-WLSVI needs the variation budget $\Delta$ to set the optimal value of the forgetting parameter as $\log \left(1 / \eta^{\star}\right)=\left(\frac{\Delta}{d K}\right)^{1 / 2}$. We can use the Bandit-over-Reinforcement-Learning framework (BoRL) [Cheung et al., 2020] to tune the forgetting parameter online.

The idea is to run a multi-armed bandit algorithm over a set of sub-algorithm each using a different parameter. In our case, each sub-algorithm is a OPT-WLSVI with a different guess on $\eta^{\star}$. If $\Delta \geq \sqrt{K}$ the regret bound is vacuous (linear regret), we are only interested in problems with $\Delta$ in the range $[1, \sqrt{K}]$. This implies that the set of $\log (1 / \eta)$ only needs to span the range $\left[\frac{1}{\sqrt{d K}}, \frac{1}{\sqrt{d}}\right]$.

We divide the horizon $K$ into $\frac{K}{M}$ equal-length intervals each of length $M$, specified later. In each interval, sub-algorithm $i$ restarts a OPT-WLSVI with $\log \left(1 / \eta_{i}\right)=\frac{2^{i}}{\sqrt{d K}}$. We have in total $I=\left\lfloor\log _{2}(\sqrt{K})\right\rfloor+1$ possible value of $\log (1 / \eta)$ in the form of $\frac{2^{i}}{\sqrt{d K}}$ that spans $\left[\frac{1}{\sqrt{d K}}, \frac{1}{\sqrt{d}}\right]$. We can verify that there exists $i^{\star} \in[A]$ such that $\log \left(1 / \eta_{i^{\star}}\right) \leq \log \left(1 / \eta^{\star}\right) \leq 2 \log \left(1 / \eta_{i^{\star}}\right)$, which wellapproximates the optimal parameter up to constant factors.

On top of these sub-algorithms, we run the EXP3.P [Auer et al., 2002]. The arms are the sub-algorithms. There are $I$ arms and the reward for each arm or sub-algorithm $i$ in interval $m \in\left[\frac{K}{M}\right]$ is the total of reward collected in the MDP during this interval. EXP3.P is called for $\frac{K}{M}$ rounds to select the arm.

Regret Analysis of OPT-WLSVI + BoRL: Let $i_{m}$ the arm selected by EXP3.P for the interval $m \in\left[\frac{K}{M}\right]$ and $\pi^{i}$ is the algorithm followed by a sub-algorithm $i$. The regret of the overall algorithm can be decomposed as the regret of the algorithm $i^{\star}$ that optimally tunes the parameter plus the loss due to learning $i^{\star}$ with the EXP3.P algorithm:

$$
\begin{aligned}
\operatorname{REGRET}(K) & =\left(\sum_{t=1}^{K} V_{t, 1}^{\star}\left(s_{t}^{1}\right)-V_{t, 1}^{\pi_{t}^{i^{\star}}}\left(s_{t}^{1}\right)\right) \\
& +\left(\sum_{m=1}^{\frac{K}{M}} \sum_{t=(m-1) M+1}^{m M} V_{t, 1}^{\pi_{t}^{i^{\star}}}\left(s_{t}^{1}\right)-V_{t, 1}^{\pi_{t}^{i_{m}}}\left(s_{t}^{1}\right)\right)
\end{aligned}
$$

The fist term corresponds to OPT-WLSVI with parameter $\eta_{i^{\star}}$. Therefore, we can bound this term using Theorem 6.2. As $\eta_{i^{\star}}$ differs from $\eta^{\star}$ up to constant factor, we obtain the bound in Corollary 6.3 i.e $\tilde{\mathcal{O}}\left(d^{5 / 4} H^{2} \Delta^{1 / 4} K^{3 / 4}\right)$.

The second term corresponds to the regret of the EXP3.P learner against the sub-algorithm $i^{\star}$. There are $I$ arms, EXP3.P is called for $\frac{K}{M}$ rounds and the rewards collected during each interval is upperbounded by $M H$. Therefore, by a classical regret bound of ExP3.P [Auer et al., 2002], the second term is upperbounded with high probability by:

$$
\tilde{\mathcal{O}}\left(M H \sqrt{I \frac{K}{M}}\right)=\tilde{\mathcal{O}}(H \sqrt{M K})
$$

We obtain that REGRET $(K)=\tilde{\mathcal{O}}\left(d^{5 / 4} H^{2} \Delta^{1 / 4} K^{3 / 4}+H \sqrt{M K}\right)$ and by choosing $M=\sqrt{K}$, we conclude that $\operatorname{REGRET}(K)=\tilde{\mathcal{O}}\left(d^{5 / 4} H^{2} \Delta^{1 / 4} K^{3 / 4}\right)$. Note that we obtain the same regret bound when the variation budget is known.

### 6.7 Technical Highlights

In this section, we give an overview of some key ideas leading to the regret bound in Theorem 4. Inspired by the analysis of weighting approach in bandit [Russac et al., 2019], one can attempt to interpret the algorithm as acting optimistically with respect to the weighted parameters of the optimal Q-value defined as $\overline{\mathbf{w}}_{t, h}(s, a)=\Sigma_{t, h}^{-1}\left(\sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top} \mathbf{w}_{\tau, h}^{\star}+\lambda \eta^{-(t-1)} \mathbf{w}_{t, h}^{\star}\right)$, where $\mathbf{w}_{t, h}^{\star}$ are the true parameters of the optimal Q -value. This first attempt was unsuccessful. Then, we came up with the implicitly defined weighted MDP and we were able to interpret our algorithm as acting optimistically with respect to this weighted MDP. We provide the full proofs and derivations in the appendix. We first translate the parameter update produced by the algorithm into the following compact update of $Q$-value estimates for any $t \in[K]$ and $h \in\{H, \ldots 1\}$ :

$$
\begin{equation*}
Q_{t, h}=\widehat{r}_{t, h}+\widehat{P}_{t, h} V_{t, h+1}+B_{t, h} \tag{6.7.1}
\end{equation*}
$$

where we define the implicitly empirical reward function $\widehat{r}$ and transition measure $\widehat{P}$ as follows:

$$
\begin{aligned}
\widehat{r}_{t, h}(s, a) & \triangleq \phi(s, a)^{\top} \Sigma_{t, h}^{-1}\left(\sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau, h} r_{\tau, h}\right) \\
\widehat{P}_{t, h}(\cdot \mid s, a) & \triangleq \phi(s, a)^{\top} \Sigma_{t, h}^{-1}\left(\sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau, h} \delta\left(\cdot, s_{\tau, h+1}\right)\right),
\end{aligned}
$$

and $B_{t, h}(\cdot, \cdot)=\beta\left(\phi(\cdot, \cdot)^{\top} \Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1} \phi(\cdot, \cdot)\right)^{1 / 2}=\beta\|\phi(\cdot, \cdot)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}$ is the exploration bonus.

We can interpret the Equation (6.7.1) as an approximation of the backward induction in a weighted average MDP defined formally as follows.
Definition 3 (Weighted Average MDP). let for any $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$
\begin{aligned}
& \bar{r}_{t, h}(s, a) \triangleq \phi(s, a)^{\top} \Sigma_{t, h}^{-1}\left(\sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top} \boldsymbol{\theta}_{\tau, h}+\lambda \eta^{-(t-1)} \boldsymbol{\theta}_{t, h}\right) \\
& \bar{P}_{t, h}(\cdot \mid s, a) \triangleq \phi(s, a)^{\top} \Sigma_{t, h}^{-1}\left(\sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top} \boldsymbol{\mu}_{\tau, h}(\cdot)+\lambda \eta^{-(t-1)} \boldsymbol{\mu}_{t, h}(\cdot)\right) .
\end{aligned}
$$

$(\mathcal{S}, \mathcal{A}, \bar{P}, \bar{r})$ is called the weighted average MDP.

We can see that if we ignore the regularization term (we set $\lambda$ to zero), $\widehat{r}_{t, h}$ coincides with $\bar{r}_{t, h}$ and $\widehat{P}_{t, h}$ is an unbiased estimate of $\bar{P}_{t, h}$. Therefore, in contrast with the stationary case, we are tracking the $Q$-value of the weighted average MDP instead of the true MDP at time $t$. The next Lemma quantifies the bias arising from the time variations of the environment.
Lemma 9 (Non-stationarity bias). For any $W \in[t-1]$ and for any bounded function $f: \mathcal{S} \rightarrow \mathbb{R}$ such as $\|f\|_{\infty} \leq H$, we have:

$$
\left|r_{t, h}(s, a)-\bar{r}_{t, h}(s, a)\right| \leq b i a s_{r}(t, h), \quad\left|\left[\left(P_{t, h}-\bar{P}_{t, h}\right) f\right](s, a)\right| \leq \operatorname{Hbias}_{P}(t, h)
$$

where

$$
\begin{aligned}
\operatorname{bias}_{r}(t, h) & =\sqrt{\frac{d}{\lambda(1-\eta)}} \sum_{s=t-W}^{t-1}\left\|\boldsymbol{\theta}_{s, h}-\boldsymbol{\theta}_{s+1, h}\right\|+\frac{2 \sqrt{d} \eta^{W}}{\lambda(1-\eta)^{\prime}} \\
\operatorname{bias}_{P}(t, h) & =\sqrt{\frac{d}{\lambda(1-\eta)}} \sum_{s=t-W}^{t-1}\left\|\boldsymbol{\mu}_{s, h}(\mathcal{S})-\boldsymbol{\mu}_{s+1, h}(\mathcal{S})\right\|+\frac{2 \sqrt{d} \eta^{W}}{\lambda(1-\eta)} .
\end{aligned}
$$

We analyse now the one-step error decomposition of the difference between the estimates $Q_{t, h}$ and $Q_{t, h}^{\pi}$ of a given policy $\pi$. To do that, we use the weighted $\operatorname{MDP}(\mathcal{S}, \mathcal{A}, \bar{P}, \bar{r})$ to isolate the bias term. The decomposition contains four parts: the reward bias and variance, the transition bias and variance, and the difference in value functions at step $h+1$. It can be written as:

$$
\begin{aligned}
& \boldsymbol{\phi}(s, a)^{\top} \mathbf{w}_{t, h}-Q_{t, h}^{\pi}(s, a)= \underbrace{\left(\bar{r}_{t, h}-r_{t, h}\right)(s, a)}_{\text {reward bias }}+\underbrace{\left(\widehat{r}_{t, h}-\bar{r}_{t, h}\right)(s, a)}_{\text {reward variance }}+ \\
& \underbrace{\left[\left(\bar{P}_{t, h}-P_{t, h}\right) V_{t, h+1}^{\pi}\right](s, a)}_{\text {transition bias }}+\underbrace{\left[\left(\widehat{P}_{t, h}-\bar{P}_{t, h}\right) V_{t, h}\right](s, a)}_{\text {transition variance }}+\underbrace{\left[\bar{P}_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)}_{\text {difference in value functions of next step }} .
\end{aligned}
$$

This differs from the error decomposition in the analysis of LSVI-UCB in several aspects: firstly, the variance terms are with respect the newly defined weighted MDP and not the true MPD. Secondly, we have additional reward and transition bias terms. Finally, the difference in the difference in value-functions at step $h+1$ hides also another bias term. Therefore, we need to carefully propagate bias terms through iteration.

The reward and transition bias terms are controlled by Lemma 9 using the fact that $\left\|V_{t, h}^{\pi}\right\|_{\infty} \leq H$. The difference in value-functions at step $h+1$ can be rewritten as $\left[P_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)+\left[\left(\bar{P}_{t, h}-P_{t, h}\right)\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)$. We control the second term by applying again Lemma 9 since $\left\|V_{t, h+1}-V_{t, h+1}^{\pi}\right\|_{\infty} \leq$ H.

It remains now the two variance terms. The reward variance is easy to control and it reduces simply to the bias due to the regularization as we assume
that $r$ is a deterministic function. Note that the assumption of deterministic reward is not a limiting assumption since the contribution of a stochastic reward in the final regret has lower order term than the contribution of a stochastic transition. Controlling the transition variance is more involved. Basically, we would like use the concentration of weighted self-normalized processes [Russac et al., 2019] to get a high probability bound. However, as $V_{t, h+1}$ is estimated from past transitions and thus depends on the latter in a non-trivial way, we show a concentration bound that holds uniformly for all possible value functions generated by the algorithm. This done by using a union bound argument over an $\epsilon$-net of the set of possible value functions with an appropriate value of $\epsilon$.

Lemma 10. For any $\delta \in(0,1)$, with probability at least $1-\delta / 2$, we have for all $(t, h) \in[K] \times[H]$,

$$
\left\|\sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau, h} \epsilon_{\tau, h}\right\|_{\tilde{\Sigma}_{t, h}^{-1}} \leq C d H \sqrt{\log \left(\frac{d H \beta}{\lambda(1-\eta)} \cdot \frac{2}{\delta}\right)}
$$

where $C>0$ is an absolute constant.
Let bias $\triangleq$ bias $_{r}+$ bias $_{P}$ the total non-stationarity bias of the MDP. By deriving an appropriate value of $\beta$ (see Lemma 20) and an induction arguments, we establish the optimism of our value estimates.

Lemma 11 (Optimism). There exists an absolute value c such that $\beta=c d H \sqrt{2}$ where $\imath=\log \left(\frac{2 d H}{(1-\eta) \delta}\right), \lambda=1$ and for all $(s, a, t, h) \in \mathcal{S} \times \mathcal{A} \times[K] \times[H]$, we have with probability at least $1-\delta / 2$

$$
\begin{equation*}
Q_{t, h}(s, a)+2 H \sum_{h^{\prime}=h}^{H} b i a s(t, h) \geq Q_{t, h}^{\star}(s, a) \tag{6.7.2}
\end{equation*}
$$

### 6.8 RELATED WORK

RL with linear function approximation: Provable algorithms with linear function approximation have seen a growing research interest in the recent literature. Under the assumption of stationary linear MDP, Jin et al. [2020b] propose an optimistic version of LSVI (LSVI-UCB) that achieves a regret of $\widetilde{\mathcal{O}}\left(d^{3 / 2} H^{2} K^{1 / 2}\right)$ where the exploration is induced by adding a UCB bonus to the estimate of the action-value function. Whereas Zanette et al. [2020a] introduce a randomized version of LVSI that achieves $\widetilde{\mathcal{O}}\left(d^{2} H^{2} K^{1 / 2}\right)$ regret where the exploration is induced by perturbing the estimate of the action-value function. Lately, Zanette et al. [2020b] consider a more general assumption, zero inherent Bellman error,
which states that the space of linear functions is close with respect to the Bellman operator (Note that linear MDPs have zero inherent Bellman error). Instead of adding UCB bonuses for every experienced states at each step $h \in[H]$, they propose to solve a global planning optimization program that returns an optimistic solution at the initial state, achieving $\widetilde{\mathcal{O}}\left(d H^{2} K^{1 / 2}\right)$ regret. Yang and Wang [2019] study a slightly different assumption where the transition kernel admits a three-factor low-rank factorization $P(\cdot \mid \cdot)=\phi(\cdot)^{\top} M^{\star} \psi(\cdot)$. They propose a model-based algorithm that tries to learn the core matrix $M^{\star}$ and they show that it achieves $\widetilde{\mathcal{O}}\left(d H^{2} K^{1 / 2}\right)$ regret.

Linear function approximations have also been studied in adversarial settings, where the reward function is allowed to change between episodes in an adversarial manner but the transition kernel stays the same. In the fullinformation setting, Cai et al. [2019] propose an optimistic policy optimization algorithm that achieves $\widetilde{\mathcal{O}}\left(d H^{2} K^{1 / 2}\right)$ providing that the transition kernel has a linear structure $P_{h}\left(s^{\prime} \mid s, a\right)=\boldsymbol{\psi}\left(s, a, s^{\prime}\right)^{\top} \theta_{h}$. In the bandit feedback setting, Neu and Olkhovskaya [2020] propose a new algorithm based on adversarial linear bandit that achieves $\widetilde{\mathcal{O}}\left((d|\mathcal{A}|)^{1 / 3} H^{2} K^{2 / 3}\right)$ regret under the assumption that all action-value functions can be represented as linear functions.

Concurrently to our work, Zhou et al. [2020] also study non-stationary linear MDPs. They establish a lower bound of $\Omega\left(d^{2 / 3} H^{2} \Delta^{1 / 3} K^{2 / 3}\right)$ and they propose a restart strategy that achieves the same dynamic regret as in our Corollary 1. Their algorithm consists in restarting periodically LSVI-UCB, and is thus markedly different from our approach. By throwing away historical data from time to time such a restart strategy would be best suited for abrupt changes in the environment, whereas our approach, by smoothly forgetting the past, would be more beneficial for gradually changing environments. Empirical comparison of both strategies in the bandit setting [Zhao et al., 2020] confirm this.

Non-stationary RL: Provably efficient algorithms for non-stationary RL in the tabular case have been introduced in several recent works. While Gajane et al. [2018] and Cheung et al. [2020] use a sliding-window approach, Ortner et al. [2019] implement a restart strategy where at each restart, past observations are discarded and new estimators for the reward and the transition kernel are built from scratch. Very recently, Domingues et al. [2020a] tackle nonstationary RL in continuous environments, where rewards and transition kernel are assumed to be Lipschitz with respect to some similarity metric over the state-action space. They propose a kernel-based algorithm with regret guarantee using time and space dependant smoothing kernels.

### 6.9 CONCLUSION

In this paper, we studied the problem of RL with linear function approximation in a changing environment where the reward and the transition kernel can change from time to time as long as the total changes are bounded by some variation budget. We introduced a provably efficient algorithm in this setting. The algorithm uses a discount factor to reduce the influence of the past and estimates the $Q$-value's parameters through weighted LSVI. We revisited as well the linear bandit setting. We pointed out a serious technical problem in the analysis of all forgetting strategies. Then, we provide a new regret analysis of these algorithms.

Limitations: In order to obtain theoretical guarantees, we need to make some assumptions such as Linear MDPs. Assumptions weaker than linear MDP either result in computationally inefficient algorithms (as in Zanette et al. [2020b]) or require the transition to be deterministic [Du et al., 2020]. Furthermore, our $\widetilde{\mathcal{O}}\left(T^{3 / 4}\right)$ regrets for both bandits and MDPs don't match the $\Omega\left(T^{2 / 3}\right)$ lower bounds for these problems. Forgetting strategies have been mistakenly believed optimal in linear bandits. In contrast, our work shows that the latter is not true and leaves the question of minimax rate open again. It is an interesting direction to explore in future work.

# Learning One Representation to Optimize All Rewards 

### 7.1 Prologue to the Contribution

### 7.1.1 Article Details

This chapter is based on the article «Learning One Representation to Optimize All Rewards» [Touati and Ollivier, 2021], joint work with Yann Ollivier. This paper was presented as long oral in the self-supervision for reinforcement learning workshop at ICLR 2021. It was also accepted at NeurIPS 2021. I am the first author. Yann Ollivier was at the origin of the idea of the project and had already developed a large part of the theory. I was in charge of implementing the new algorithm and performing the experiments. I also contributed to further develop the theoretical analysis.

### 7.1.2 Context

We have focused so far on the classical RL paradigm which aims to maximize the cumulative reward when the agent has access to the reward signal. Despite being able to capture many AI applications, this reward-driven framework leads to a task-specific agent that is only able to solve the task at hand and needs to be trained from the scratch to adapt to new tasks. It is obvious here that we overlook a lot of information about the environment that could be leveraged across tasks. In this project, we are motivated by the following research problem:

Imagine you have a fixed environment where you can perform actions but you receive no reward information. Your mission is to build a summary of the environment using these reward-free interactions, such that as soon as I describe a reward function, you immediately know what to do to maximize your reward.

This is a step towards building controllable agents able to follow instructions.

### 7.1.3 Paper Abstract

We introduce the forward-backward (FB) representation of the dynamics of a reward-free Markov decision process. It provides explicit near-optimal policies for any reward specified a posteriori. During an unsupervised phase, we use
reward-free interactions with the environment to learn two representations via off-the-shelf deep learning methods and temporal difference (TD) learning. In the test phase, a reward representation is estimated either from reward observations or an explicit reward description (e.g., a target state). The optimal policy for that reward is directly obtained from these representations, with no planning. We assume access to an exploration scheme or replay buffer for the first phase.

The unsupervised FB loss is well-principled: if training is perfect, the policies obtained are provably optimal for any reward function. With imperfect training, the sub-optimality is proportional to the unsupervised approximation error. The FB representation learns long-range relationships between states and actions, via a predictive occupancy map, without having to synthesize states as in modelbased approaches.

Our approach compares well to goal-oriented RL algorithms on discrete and continuous mazes, pixel-based MsPacman, and the FetchReach virtual robot arm. We also illustrate how the agent can immediately adapt to new tasks beyond goal-oriented RL.

Our Code is available at: https://github.com/ahmed-touati/ controllable_agent.

### 7.2 INTRODUCTION

We consider one kind of unsupervised reinforcement learning problem: Given a Markov decision process (MDP) but no reward information, is it possible to learn and store a compact object that, for any reward function specified later, provides the optimal policy for that reward, with a minimal amount of additional computation? In a sense, such an object would encode in a compact form the solutions of all possible planning problems in the environment. This is a step towards building agents that are fully controllable after first exploring their environment in an unsupervised way.

Goal-oriented RL methods [Andrychowicz et al., 2017, Plappert et al., 2018] compute policies for a series of rewards specified in advance (such as reaching a set of target states), but cannot adapt in real time to new rewards, such as weighted combinations of target states or dense rewards.

Learning a model of the world is another possibility, but it still requires explicit planning for each new reward; moreover, synthesizing accurate trajectories of states over long time ranges has proven difficult [Talvitie, 2017, Ke et al., 2018].

Instead, we exhibit an object that is both simpler to learn than a model of the world, and contains the information to recover near-optimal policies for any reward provided a posteriori, without a planning phase.

Borsa et al. [2018] learn optimal policies for all rewards that are linear combinations of a finite number of feature functions provided in advance by the user. This limits applications: e.g., goal-oriented tasks would require one feature per goal state, thus using infinitely many features in continuous spaces. We reuse a policy parameterization from Borsa et al. [2018], but introduce a novel representation with better properties, based on state occupancy prediction instead of expected featurizations. We use theoretical advances on successor state learning from Blier et al. [2021]. We obtain the following.

- We prove the existence of a learnable "summary" of a reward-free discrete or continuous MDP, that provides an explicit formula for optimal policies for any reward specified later. This takes the form of a pair of representations $F: S \times A \times Z \rightarrow Z$ and $B: S \times A \rightarrow Z$ from state-actions into a representation space $Z \simeq \mathbb{R}^{d}$, with policies $\pi_{z}(s) \triangleq \operatorname{argmax}_{a} F(s, a, z)^{\top} z$. Once a reward is specified, a value of $z$ is computed from reward values and $B$; then $\pi_{z}$ is used. Rewards may be specified either explicitly as a function, or as target states, or by samples as in usual RL setups.
- We provide a well-principled unsupervised loss for $F$ and $B$. If FB training is perfect, then the policies are provably optimal for all rewards (Theorem 1). With imperfect training, sub-optimality is proportional to the FB training error (Theorems 5-6). In finite spaces, perfect training is possible with large enough dimension $d$ (Proposition 6).
Explicitly, $F$ and $B$ are trained so that $F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)$ approximates the long-term probability to reach $s^{\prime}$ from $s$ if following $\pi_{z}$. This is akin to a model of the environment, without synthesizing state trajectories.
- We provide a TD-like algorithm to train $F$ and $B$ for this unsupervised loss, with function approximation, adapted from recent methods for successor states [Blier et al., 2021]. No sparse rewards are used: every transition reaches some state $s^{\prime}$, so every step is exploited. As usual with TD, learning seeks a fixed point but the loss itself is not observable.
- We prove viability of the method on several environments from mazes to pixel-based MsPacman and a virtual robotic arm. For single-state rewards (learning to reach arbitrary states), we provide quantitative comparisons with goal-oriented methods such as HER. (Our method is not a substitute for HER: in principle they could be combined, with HER improving replay buffer management for our method.) For more general rewards, which cannot be tackled a posteriori by trained goal-oriented models, we provide qualitative examples.
- We also illustrate qualitatively the sub-optimalities (long-range behavior
is preserved but local blurring of rewards occurs) and the representations learned.


### 7.3 Problem and Notation

Let $\mathcal{M}=(S, A, P, \gamma)$ be a reward-free Markov decision process.
For any policy $\pi:: S \rightarrow \operatorname{Prob}(A)$ and state-action $\left(s_{0}, a_{0}\right)$, define the successor measure $M^{\pi}\left(s_{0}, a_{0}, \cdot\right)$ as the measure over $S \times A$ representing the expected discounted time spent in each set $X \subset S \times A$ :

$$
\begin{equation*}
M^{\pi}\left(s_{0}, a_{0}, X\right) \triangleq \sum_{t \geq 0} \gamma^{t} \operatorname{Pr}\left(\left(s_{t}, a_{t}\right) \in X \mid s_{0}, a_{0}, \pi\right) \tag{7.3.1}
\end{equation*}
$$

for each $X \subset S \times A$. Viewing $M$ as a measure deals with both discrete and continuous spaces.

We consider the following informal problem: Given a reward-free MDP $(S, A, P, \gamma)$, can we compute a convenient learnable object $E$ such that, once a reward function $r: S \times A \rightarrow \mathbb{R}$ is specified, we can easily (with no planning) compute, from $E$ and $r$, a policy $\pi$ whose performance is close to maximal?

### 7.4 Encoding All Optimal Policies Via the FORWARD-BACKWARD REPRESENTATION

We first present forward-backward (FB) representations of a reward-free MDP as a way to summarize all optimal policies via explicit formulas. The resulting learning procedure is described in Section 7.5.

Core idea. The main algebraic idea is as follows. Assume, at first, that $S$ is finite. For a fixed policy, the $Q$-function depends lineary on the reward: namely, $Q_{r}^{\pi}(s, a)=\sum_{s^{\prime}, a^{\prime}} M^{\pi}\left(s, a, s^{\prime}, a^{\prime}\right) r\left(s^{\prime}, a^{\prime}\right)$ where $M^{\pi}\left(s, a, s^{\prime}, a^{\prime}\right)=\sum_{t \geq 0} \gamma^{t} \operatorname{Pr}\left(\left(s_{t}, a_{t}\right)=\left(s^{\prime}, a^{\prime}\right) \mid s, a, \pi\right)$. This rewrites as $Q_{r}^{\pi}=M^{\pi} r$ viewing everything as vectors and matrices indexed by state-actions.

Now let $\left(\pi_{z}\right)_{z \in \mathbb{R}^{d}}$ be any family of policies parameterized by $z$. Assume that for each $z$, we can find $d \times(S \times A)$-matrices $F_{z}$ and $B$ such that $M^{\pi_{z}}=F_{z}^{\top} B$. Then $Q_{r}^{\pi_{z}}=F_{z}^{\top} B r$. Specializing to $z_{R} \triangleq B r$, the $Q$-function of policy $\pi_{z_{R}}$ on reward $r$ is $Q_{r}^{\pi_{z_{R}}}=F_{z_{R}}^{\top} z_{R}$. So far $\pi_{z}$ was unspecified; but if we define $\pi_{z}(s) \triangleq$ $\operatorname{argmax}_{a}\left(F_{z}^{\top} z\right)_{s a}$ at each state $s$, then by definition, $\pi_{z_{R}}$ is the greedy policy with respect to $F_{z_{R}}^{\top} z_{R}$. At the same time, $F_{z_{R}}^{\top} z_{R}$ is the $Q$-function of $\pi_{z_{R}}$ for reward $r$ :
thus, $\pi_{z_{R}}$ is the greedy policy of its own $Q$-function, and is therefore optimal for reward $r$.

Thus, if we manage to find $F, B$, and $\pi_{z}$ such that $\pi_{z}=\operatorname{argmax} F_{z}^{\top} z$ and $F_{z}^{\top} B=M^{\pi_{z}}$ for all $z \in \mathbb{R}^{d}$, then we obtain the optimal policy for any reward $r$, just by computing $B r$ and applying policy $\pi_{B r}$.

This criterion on ( $F, B, \pi_{z}$ ) is entirely unsupervised. Since $F$ and $B$ depend on $\pi_{z}$ but $\pi_{z}$ is defined via $F$, this is a fixed point equation. An exact solution exists for $d$ large enough (Appendix, Prop. 6), while a smaller $d$ provides lowerrank approximations $M^{\pi_{z}} \approx F_{z}^{\top} B$. In Section 7.5 we present a well-grounded algorithm to learn such $F, B$, and $\pi_{z}$.

In short, we learn two representations $F$ and $B$ such that $F\left(s_{0}, a_{0}, z\right)^{\top} B\left(s^{\prime}, a^{\prime}\right)$ is approximately the long-term probability $M^{\pi_{z}}\left(s_{0}, a_{0}, s^{\prime}, a^{\prime}\right)$ to reach $\left(s^{\prime}, a^{\prime}\right)$ if starting at $\left(s_{0}, a_{0}\right)$ and following policy $\pi_{z}$. Then all optimal policies can be computed from $F$ and $B$. We think of $F$ as a representation of the future of a state, and $B$ as the ways to reach a state (Appendix 7.8.4): if $F^{\top} B$ is large, then the second state is reachable from the first. This is akin to a model of the environment, without synthesizing state trajectories.

General statement. In continuous spaces with function approximation, $F_{z}$ and $B$ become functions $S \times A \rightarrow \mathbb{R}^{d}$ instead of matrices; since $F_{z}$ depends on $z, F$ itself is a function $S \times A \times \mathbb{R}^{d} \rightarrow \mathbb{R}^{d}$. The sums over states will be replaced with expectations under the data distribution $\rho$.

Definition 4 (Forward-backward representation). Let $Z=\mathbb{R}^{d}$ be a representation space, and let $\rho$ be a measure on $S \times A$. A pair of functions $F: S \times A \times Z \rightarrow Z$ and $B: S \times A \rightarrow Z$, together with a parametric family of policies $\left(\pi_{z}\right)_{z \in Z}$, is called a forward-backward representation of the MDP with respect to $\rho$, if the following conditions hold for any $z \in Z$ and $(s, a),\left(s_{0}, a_{0}\right) \in S \times A$ :
$\pi_{z}(s)=\underset{a}{\operatorname{argmax}} F(s, a, z)^{\top} z, \quad M^{\pi_{z}}\left(s_{0}, a_{0}, \mathrm{~d} s, \mathrm{~d} a\right)=F\left(s_{0}, a_{0}, z\right)^{\top} B(s, a) \rho(\mathrm{d} s, \mathrm{~d} a)$
where $M^{\pi}$ is the successor measure defined in (7.3.1), and the last equality is between measures.

Theorem 1 (FB representations encode all optimal policies). Let $\left(F, B,\left(\pi_{z}\right)\right)$ be a forward-backward representation of a reward-free MDP with respect to some measure $\rho$.

Then, for any bounded reward function $r: S \times A \rightarrow \mathbb{R}$, the following holds. Set

$$
\begin{equation*}
z_{R} \triangleq \int_{s, a} r(s, a) B(s, a) \rho(\mathrm{d} s, \mathrm{~d} a) \tag{7.4.2}
\end{equation*}
$$

assuming the integral exists. Then $\pi_{z_{R}}$ is an optimal policy for reward $r$ in the MDP. Moreover, the optimal $Q$-function $Q^{\star}$ for reward $r$ is $Q^{\star}(s, a)=F\left(s, a, z_{R}\right)^{\top} z_{R}$.

For instance, for a single reward located at state-action ( $s, a$ ), the optimal policy is $\pi_{z_{R}}$ with $z_{R}=B(s, a)$. (In that case the factor $\rho(\mathrm{d} s, \mathrm{~d} a)$ does not matter because scaling the reward does not change the optimal policy.)

We present in Section 7.5 an algorithm to learn FB representations. The measure $\rho$ will be the distribution of state-actions visited in a training set or under an exploration policy: then $z_{R}=\mathbb{E}_{(s, a) \sim \rho}[r(s, a) B(s, a)]$ can be obtained by sampling from visited states.

In finite spaces, exact FB representations exist, provided the dimension $d$ is larger than $\# S \times \# A$ (Appendix, Prop. 6). In infinite spaces, arbitrarily good approximations can be obtained by increasing $d$, corresponding to a rank-d approximation of the cumulated transition probabilities $M^{\pi}$. Importantly, the optimality guarantee extends to approximate $F$ and $B$, with optimality gap proportional to $F^{\top} B-M^{\pi_{z}} / \rho$ (Appendix, Theorems 5-6 with various norms on $F^{\top} B-M^{\pi} / \rho$ ). For instance, if, for some reward $r$, the error $\left|F\left(s_{0}, a_{0}, z_{R}\right)^{\top} B(s, a)-M^{\pi_{z_{R}}}\left(s_{0}, a_{0}, \mathrm{~d} s, \mathrm{~d} a\right) / \rho(\mathrm{d} s, \mathrm{~d} a)\right|$ is at most $\epsilon$ on average over $(s, a) \sim \rho$ for every $\left(s_{0}, a_{0}\right)$, then $\pi_{z_{R}}$ is $3 \epsilon\|r\|_{\infty} /(1-\gamma)$-optimal for $r$.

These results justify using some norm over $\left|F^{\top} B-M^{\pi_{z}} / \rho\right|$, averaged over $z \in \mathbb{R}^{d}$, as a training loss for unsupervised reinforcement learning. (Below, we average over $z \in \mathbb{R}^{d}$ from a fixed rescaled Gaussian. If prior information is available on the rewards $r$, the corresponding distribution of $z_{R}$ may be used instead.)

If $B$ is fixed in advance and only $F$ is learned, the method has similar properties to successor features based on $B$ (Appendix 7.8.4). But one may set a large $d$ and let $B$ be learned: arguably, by Theorem 1, the resulting features "linearize" optimal policies as much as possible. The features learned in $F$ and $B$ may have broader interest.

### 7.5 LEARNING AND USING FORWARD-BACKWARD Representations

Our algorithm starts with an unsupervised learning phase, where we learn the representations $F$ and $B$ in a reward-free way, by observing state transitions in the environment, generated from any exploration scheme. Then, in a reward estimation phase, we estimate a policy parameter $z_{R}=\mathbb{E}[r(s, a) B(s, a)]$ from some reward observations, or directly set $z_{R}$ if the reward is known (e.g., set $z_{R}=$ $B(s, a)$ to reach a known target $(s, a))$. In the exploitation phase, we directly use the policy $\pi_{z_{R}}(s)=\operatorname{argmax}_{a} F\left(s, a, z_{R}\right)^{\top} z_{R}$.

The unsupervised learning phase. No rewards are used in this phase, and no family of tasks has to be specified manually. $F$ and $B$ are trained off-policy from
observed transitions in the environment, to approximate the successor density: $F^{\top}\left(s_{0}, a_{0}, z\right) B\left(s^{\prime}, a^{\prime}\right) \approx m^{\pi_{z}}\left(s_{0}, a_{0}, s^{\prime}, a^{\prime}\right)$ for every $z$. Training is based on the Bellman equation for the successor measure $M^{\pi}$,

$$
\begin{equation*}
M^{\pi}\left(s_{0}, a_{0},\left\{\left(s^{\prime}, a^{\prime}\right)\right\}\right)=\mathbb{1}_{s_{0}=s^{\prime}, a_{0}=a^{\prime}}+\gamma \mathbb{E}_{s_{1} \sim P\left(\mathrm{~d} s_{1} \mid s_{0}, a_{0}\right)} M^{\pi}\left(s_{1}, \pi\left(s_{1}\right),\left\{\left(s^{\prime}, a^{\prime}\right)\right\}\right) \tag{7.5.1}
\end{equation*}
$$

We leverage a well-principled algorithm from Blier et al. [2021] in the singlepolicy setting: it learns the successor density $m^{\pi}$ of a policy $\pi$ without using the sparse reward $\mathbb{1}_{s_{0}=s^{\prime}, a_{0}=a^{\prime}}$ (which would vanish in continuous spaces). This algorithm uses a parametric model $m_{\theta}^{\pi}\left(s_{0}, a_{0}, s^{\prime}, a^{\prime}\right)$. Given an observed transition $\left(s_{0}, a_{0}, s_{1}\right)$ from the training set, generate an action $a_{1} \sim \pi\left(a_{1} \mid s_{1}\right)$, and sample another state-action $\left(s^{\prime}, a^{\prime}\right)$ from the training set, independently from $\left(s_{0}, a_{0}, s_{1}\right)$. Then update the parameter $\theta$ by $\theta \leftarrow \theta+\eta \delta \theta$ with learning rate $\eta$ and
$\delta \theta \triangleq \partial_{\theta} m_{\theta}^{\pi}\left(s_{0}, a_{0}, s_{0}, a_{0}\right)+\partial_{\theta} m_{\theta}^{\pi}\left(s_{0}, a_{0}, s^{\prime}, a^{\prime}\right) \times\left(\gamma m_{\theta}^{\pi}\left(s_{1}, a_{1}, s^{\prime}, a^{\prime}\right)-m_{\theta}^{\pi}\left(s_{0}, a_{0}, s^{\prime}, a^{\prime}\right)\right)$
This computes the density $m^{\pi}$ of $M^{\pi}$ with respect to the distribution $\rho$ of stateactions in the training set. Namely, the true successor state density $m^{\pi}$ is a fixed point of (7.5.2) in expectation Blier et al. [2021]. Variants exist, such as using a target network for $m_{\theta}^{\pi}\left(s_{1}, a_{1}, s^{\prime}, a^{\prime}\right)$ on the right-hand side, as in DQN.

Thus, we first choose a parametric model $F_{\theta}, B_{\theta}$ for the representations $F$ and $B$, and set $m_{\theta}^{\pi_{z}}\left(s_{0}, a_{0}, s^{\prime}, a^{\prime}\right) \triangleq F_{\theta}\left(s_{0}, a_{0}, z\right)^{\top} B_{\theta}\left(s^{\prime}, a^{\prime}\right)$. Then we iterate the update (7.5.2) over many state-actions and values of $z$. This results in Algorithm 7. At each step, a value of $z$ is picked at random, together with a batch of transitions $\left(s_{0}, a_{0}, s_{1}\right)$ and a batch of state-actions $\left(s^{\prime}, a^{\prime}\right)$ from the training set, with $\left(s^{\prime}, a^{\prime}\right)$ independent from $z$ and $\left(s_{0}, a_{0}, s_{1}\right)$.

For sampling $z$, we use a fixed distribution (rescaled Gaussians, see Appendix E.2). Any number of values of $z$ may be sampled: this does not use up training samples. We use a target network with soft updates (Polyak averaging) as in DDPG. For training we also replace the greedy policy $\pi_{z}=$ $\operatorname{argmax}_{a} F(s, a, z)^{\top} z$ with a regularized version $\pi_{z}=\operatorname{softmax}\left(F(s, a, z)^{\top} z / \tau\right)$ with fixed temperature $\tau$ (Appendix E.2). Since there is unidentifiability between $F$ and $B$ (Remark 4), we normalize $B$ via an auxiliary loss in Algorithm 7.

For exploration in this phase, we use the policies being learned: the exploration policy chooses a random value of $z$ from some distribution (e.g., Gaussian), and follows $\pi_{z}$ for some time ( Algorithm 7). However, the algorithm can also work from an existing dataset of off-policy transitions.

The reward estimation phase. Once rewards are available, we estimate a reward representation (policy parameter) $z_{R}$ by weighing the representation $B$ by the reward:

$$
\begin{equation*}
z_{R} \triangleq \mathbb{E}[r(s, a) B(s, a)] \tag{7.5.3}
\end{equation*}
$$

where the expectation must be computed over the same distribution $\rho$ of stateactions ( $s, a$ ) used to learn $F$ and $B$ (see Section 7.8 .5 for using a different distribution). Thus, if the reward is black-box as in standard RL algorithms, then the exploration policy has to be run again for some time, and $z_{R}$ is obtained by averaging $r(s, a) B(s, a)$ over the states visited.

If the reward is known explicitly, this phase is unnecessary. For instance, if the reward is to reach a target state-action $\left(s_{0}, a_{0}\right)$ while avoiding some forbidden state-actions $\left(s_{1}, a_{1}\right), \ldots,\left(s_{k}, a_{k}\right)$, one may directly set

$$
\begin{equation*}
z_{R}=B\left(s_{0}, a_{0}\right)-\lambda \sum B\left(s_{i}, a_{i}\right) \tag{7.5.4}
\end{equation*}
$$

where the constant $\lambda$ adjusts the negative reward for visiting a forbidden state. This can be used for goal-oriented RL.

If the reward is known algebraically as a function $r(s, a)$, then $z_{R}$ may be computed by averaging the function $r(s, a) B(s, a)$ over a replay buffer from the unsupervised training phase. We may also use a reward model $\hat{r}(s, a)$ of $r(s, a)$ trained on some reward observations from any source. An approximate value for $z_{R}$ still provides an approximately optimal policy (Proposition 7 and Theorem. 7).

The exploitation phase. Once the reward representation $z_{R}$ has been estimated, the $Q$-function is estimated as

$$
\begin{equation*}
Q(s, a)=F\left(s, a, z_{R}\right)^{\top} z_{R} . \tag{7.5.5}
\end{equation*}
$$

The corresponding policy $\pi_{z_{R}}(s)=\operatorname{argmax}_{a} Q(s, a)$ is used for exploitation.
Fine-tuning was not needed in our experiments, but it is possible to fine-tune the $Q$-function using actual rewards, by setting $Q(s, a)=F\left(s, a, z_{R}\right)^{\top} z_{R}+q_{\theta}(s, a)$ where the fine-tuning model $q_{\theta}$ is initialized to 0 and learned via any standard $Q$-learning method.

Incorporating prior information on rewards in B. Trying to plan in advance for all possible rewards in an arbitrary environment may be too generic and problem-agnostic, and become difficult in large environments, requiring long exploration and a large $d$ to accommodate all rewards. In practice, we are often interested in rewards depending, not on the full state, but only on a part or some features of the state (e.g., a few components of the state, such as the position of an agent, or its neighbordhood, rather than the full environment).

If this is known in advance, the representation $B$ can be trained on that part of the state only, with the same theoretical guarantees (Appendix, Theorem 2). F still needs to use the full state as input. This way, the FB model of the transition probabilities (7.3.1) only has to learn the future probabilities of the part of interest in $\left(s^{\prime}, a^{\prime}\right)$, based on the full initial state $\left(s_{0}, a_{0}\right)$. Explicitly, if $\phi: S \times A \rightarrow G$

```
Algorithm 7 FB algorithm: Unsupervised Phase
    Inputs: replay buffer \(\mathcal{D}\), Polyak coefficient \(\alpha, v\) a probability distribution
    over \(\mathbb{R}^{d}\), randomly initialized networks \(F_{\theta}\) and \(B_{\omega}\), learning rate \(\eta\), mini-
    batch size \(b\), number of episodes \(E\), number of gradient updates \(N\), temper-
    ature \(\tau\) and regularization coefficient \(\lambda\).
    for \(m=1, \ldots\) do
        /* Collect \(E\) episodes
        for episode \(e=1, \ldots E\) do
            Sample \(z \sim v\)
            Observe an initial state \(s_{0}\)
            for \(t=1, \ldots\) do
                Select an action \(a_{t}\) according to some behaviour policy (e.g the \(\epsilon\) -
                greedy with respect to \(\left.F_{\theta}\left(s_{t}, a, z\right)^{\top} z\right)\)
                    Observe next state \(s_{t+1}\)
            Store transition \(\left(s_{t}, a_{t}, s_{t+1}\right)\) in the replay buffer \(\mathcal{D}\)
            end for
        end for
        /* Perform \(N\) stochastic gradient descent updates
        for \(n=1 \ldots N\) do
            Sample a mini-batch of transitions \(\left\{\left(s_{i}, a_{i}, s_{i+1}\right)\right\}_{i \in I} \subset \mathcal{D}\) of size \(|I|=b\).
            Sample a mini-batch of target state-action pairs \(\left\{\left(s_{i}^{\prime}, a_{i}^{\prime}\right)\right\}_{i \in I} \subset \mathcal{D}\) of size
            \(|I|=b\).
            Sample a mini-batch of \(\left\{z_{i}\right\}_{i \in I} \sim v\) of size \(|I|=b\).
            Set \(\pi_{z_{i}}\left(\cdot \mid s_{i+1}\right)=\operatorname{softmax}\left(F_{\theta^{-}}\left(s_{i+1}, \cdot, z_{i}\right)^{\top} z_{i} / \tau\right)\)
            \(\mathscr{L}(\theta, \omega)=-\frac{1}{b} \sum_{i \in I} F_{\theta}\left(s_{i}, a_{i}, z_{i}\right)^{\top} B_{\omega}\left(s_{i}, a_{i}\right)+\)
            \(\frac{1}{2 b^{2}} \sum_{i, j \in I^{2}}\left(F_{\theta}\left(s_{i}, a_{i}, z_{i}\right)^{\top} B_{\omega}\left(s_{j}^{\prime}, a_{j}^{\prime}\right)-\gamma \sum_{a \in A} \pi_{z_{i}}\left(a \mid s_{i+1}\right) \cdot F_{\theta^{-}}\left(s_{i+1}, a, z_{i}\right)^{\top} B_{\omega^{-}}\left(s_{j^{\prime}}^{\prime}, a_{j}^{\prime}\right)\right)^{2}\)
            /* Compute orthonormality regularization loss
            \(\mathscr{L}_{\text {reg }}(\omega)=-\frac{1}{b} \sum_{i \in I} B_{\omega}\left(s_{i}, a_{i}\right)^{\top}\) stop-gradient \(\left(B_{\omega}\left(s_{i}, a_{i}\right)\right)+\)
            \(\frac{1}{b^{2}} \sum_{i, j \in I^{2}} B_{\omega}\left(s_{i}, a_{i}\right)^{\top}\) stop-gradient \(\left(B_{\omega}\left(s_{j}^{\prime}, a_{j}^{\prime}\right)\right) \cdot\) stop-gradient \(\left(B_{\omega}\left(s_{i}, a_{i}\right)^{\top} B_{\omega}\left(s_{j}^{\prime}, a_{j}^{\prime}\right)\right)\)
            Update \(\theta \leftarrow \theta-\eta \nabla_{\theta} \mathscr{L}(\theta, \omega)\) and \(\omega \leftarrow \omega-\eta \nabla_{\omega}(\mathscr{L}(\theta, \omega)+\lambda\).
            \(\left.\mathscr{L}_{\text {reg }}(\omega)\right)\)
        end for
        /* Update target network parameters
        \(\theta^{-} \leftarrow \alpha \theta^{-}+(1-\alpha) \theta\)
        \(\omega^{-} \leftarrow \alpha \omega^{-}+(1-\alpha) \omega\)
    end for
```

is a feature map to some features $g=\phi(s, a)$, and if we know that the reward will be a function $R(g)$, then Theorem 1 still holds with $B(g)$ everywhere instead
of $B(s, a)$, and with the successor density $m^{\pi}\left(s_{0}, a_{0}, g\right)$ instead of $m^{\pi}\left(s_{0}, a_{0}, s^{\prime}, a^{\prime}\right)$ ( Theorem 2). Rewards can be arbitrary functions of $g$, so this is more general than Borsa et al. [2018] which only considers rewards linear in $g$. For instance, in MsPacman below, we let $g$ be the 2D position $(x, y)$ of the agent, so we can optimize any reward function that depends on this position.

Limitations. First, this method does not solve exploration: it assumes access to a good exploration strategy. (Here we used the learned $\pi_{z}$ with random values of $z$, corresponding to random rewards.)

Next, this task-agnostic approach is relevant if the reward is not known in advance, but may not bring the best performance on a particular reward. Mitigation strategies include: increasing $d$; using prior information on rewards by including relevant variables into $B$, as discussed above; and fine-tuning the $Q$ function at test time based on the initial $F^{\top} B$ estimate.

Indeed, as reward functions are represented by a $d$-dimensional vector $z_{R}=$ $\mathbb{E}[r . B]$, some information about the reward is necessarily lost. Any reward uncorrelated to $B$ is treated as 0 . The necessary dimension $d$ for good behavior may be large. Still, $d \approx 100$ worked in our experiments, and Section 7.8.2 argues theoretically that $d=O(n)$ is enough for navigation on an $n$-dimensional grid.

We expect this method to have an implicit bias for long-range behavior (spatially smooth rewards), while local details of the reward function may be blurred. Indeed, $F^{\top} B$ is optimized to approximate the successor measure $M^{\pi}=\sum_{t} \gamma^{t}\left(P^{\pi}\right)^{t}$ with $\left(P^{\pi}\right)^{t}$ the $t$-step transition kernel for each policy $\pi$. The rank- $d$ approximation will favor large eigenvectors of $P^{\pi}$, i.e., small eigenvectors of the Markov chain Laplacian Id $-\gamma P^{\pi}$. These loosely correspond to long-range (low-frequency) behavior Mahadevan and Maggioni [2007]: presumably, $F$ and $B$ will learn spatially smooth rewards first. Indeed, experimentally, a small $d$ leads to spatial blurring of rewards and $Q$-functions (Fig. 7.3). Arguably, without any prior information this is a reasonable prior; Stachenfeld et al. [2017] have argued for the cognitive relevance of low-dimensional approximations of successor representations.

Variance is a potential issue in larger environments, although this did not arise in our experiments. Learning $m^{\pi}$ requires sampling a state-action $\left(s_{0}, a_{0}\right)$ and an independent state-action $\left(s^{\prime}, a^{\prime}\right)$. In large spaces, most state-action pairs will be unrelated. A possible mitigation is to use strategies such as Hindsight Experience Replay Andrychowicz et al. [2017] to select goals related to the current state-action. The following may help a lot: the update of $F$ and $B$ decouples as an expectation over $\left(s_{0}, a_{0}\right)$, times an expectation over $\left(s^{\prime}, a^{\prime}\right)$. Thus, by estimating these expectations by a moving average over a dataset, it is easy to have many pairs $\left(s_{0}, a_{0}\right)$ interact with many $\left(s^{\prime}, a^{\prime}\right)$. The cost is handling full $d \times d$ matrices. This will be explored in future work.

### 7.6 EXPERIMENTS

We first consider the task of reaching arbitrary goal states. For this, we can make quantitative comparisons to existing goal-oriented baselines. Next, we illustrate qualitatively some tasks that cannot be tackled a posteriori by goaloriented methods, such as introducing forbidden states. Finally, we illustrate some of the representations learned.

### 7.6.1 Environments and Experimental Setup

We run our experiments on a selection of environments that are diverse in term of state space dimensionality, stochasticity and dynamics.

- Discrete Maze is the classical gridworld with four rooms. States are represented by one-hot unit vectors.
- Continuous Maze is a two dimensional environment with impassable walls. States are represented by their Cartesian coordinates $(x, y) \in[0,1]^{2}$. The execution of one of the actions moves the agent in the desired direction, but with normal random noise added to the position of the agent.
- FetchReach is a variant of the simulated robotic arm environment from Plappert et al. [2018] using discrete actions instead of continuous actions. States are 10-dimensional vectors consisting of positions and velocities of robot joints.
- Ms. Pacman is a variant of the Atari 2600 game Ms. Pacman, where an episode ends when the agent is captured by a monster Rauber et al. [2018]. States are obtained by processing the raw visual input directly from the screen. Frames are preprocessed by cropping, conversion to grayscale and downsampling to $84 \times 84$ pixels. A state $s_{t}$ is the concatenation of $\left(x_{t-12}, x_{t-8}, x_{t-4}, x_{t}\right)$ frames, i.e. an $84 \times 84 \times 4$ tensor. An action repeat of 12 is used. As Ms. Pacman is not originally a multi-goal domain, we define the goals as the 148 reachable coordinates $(x, y)$ on the screen; these can be reached only by learning to avoid monsters.

For all environments, we run algorithms for 800 epochs. Each epoch consists of 25 cycles where we interleave between gathering some amount of transitions, to add to the replay buffer, and performing 40 steps of stochastic gradient descent on the model parameters. To collect transitions, we generate episodes using some behavior policy. For both mazes, we use a uniform policy while for FetchReach and Ms. Pacman, we use an $\epsilon$-greedy policy with respect to the current approximation $F(s, a, z)^{\top} z$ for a sampled $z$. At evaluation time, $\epsilon$-greedy policies are also used, with a smaller $\epsilon$. More details are given in Appendix E.2.


Figure 7.1: Comparative performance of FB for different dimensions and DQN in the FetchReach. Left: success rate averaged over 20 randomly selected goals as function of the first 100 training epochs. Right: success rate averaged over 20 random goals after 800 training epochs.

### 7.6.2 Goal-Oriented Setting: Quantitative Comparisons

We investigate the FB representation over goal-reaching tasks and compare it to goal-oriented baselines: $\mathrm{DQN}^{1}$, and DQN with HER when needed. We define sparse reward functions. For Discrete Maze, the reward function is equal to one when the agent's state is equal exactly to the goal state. For Discrete Maze, we measured the quality of the obtained policy to be the ratio between the true expected discounted reward of the policy for its goal and the true optimal value function, on average over all states. For the other environments, the reward function is equal to one when the distance of the agent's position and the goal position is below some threshold, and zero otherwise. We assess policies by computing the average success rate, i.e the average number of times the agent successfully reaches its goal.

Figs. 7.1 and 7.2 show the comparative performance of FB for different dimensions $d$, and DQN respectively in FetchReach and Ms. Pacman (similar results in Discrete and Continuous Mazes are provided in Appendix E.2). In Ms. Pacman, DQN totally fails to learn and we had to add HER to make it work. The performance of FB consistently increases with the dimension $d$ and the best dimension matches the performance of the goal-oriented baseline.

In Discrete Maze, we observe a drop of performance for $d=25$ (Appendix E.2, Fig. E.1): this is due to the spatial smoothing induced by the small

[^6]

Figure 7.2: Comparative performance of FB for different dimensions and DQN in Ms. Pacman. Left: success rate averaged over 20 randomly selected goals as function of the first 200 training epochs. Right: success rate averaged over the goal space after 800 training epochs.
rank approximation and the reward being nonzero only if the agent is exactly at the goal. This spatial blurring is clear on heatmaps for $d=25$ vs $d=75$ (Fig. 7.3). With $d=25$ the agent often stops right next to its goal.

To evaluate the sample efficiency of FB, after each epoch, we evaluate the agent on 20 randomly selected goals. Learning curves are reported in Figs. 7.1 and 7.2 (left). In all environments, we observe no loss in sample efficiency compared to the goal-oriented baseline. In Ms. Pacman, FB even learns faster than DQN+HER.

### 7.6.3 More Complex Rewards: Qualitative Results

We now investigate FB's ability to generalize to new tasks that cannot be solved by an already trained goal-oriented model: reaching a goal with forbidden states imposed a posteriori, reaching the nearest of two goals, and choosing between a small, close reward and a large, distant one.

First, for the task of reaching a target position $g_{0} \hat{\sim}$ while avoiding some forbidden positions $g_{1}, \ldots g_{k} \bullet$, we set $z_{R}=B\left(g_{1}\right)-\lambda \sum_{i=1}^{k} B\left(g_{i}\right)$ and run the corresponding $\epsilon$-greedy policy defined by $F\left(s, a, z_{R}\right)^{\top} z_{R}$. Fig. 7.4 shows the resulting trajectories, which succeed at solving the task for the different domains. In Ms. Pacman, the path is suboptimal (though successful) due to the sudden appearance of a monster along the optimal path. (We only plot the initial frame; see the full series of frames along the trajectory in Appendix E.2, Fig. E.9.) Fig. 7.6


Figure 7.3: Heatmap of $\max _{a} F\left(s, a, z_{R}\right)^{\top} z_{R}$ for $z_{R}=B(\Re)$ Left: $d=25$. Right: $d=75$.


Figure 7.4: Trajectories generated by the $F^{\top} B$ policies for the task of reaching a target position (star shape $\underset{\sim}{ }$ ) while avoiding forbidden positions (red circle).
(left) provides a contour plot of $\max _{a \in A} F\left(s, a, z_{R}\right)^{\top} z_{R}$ for the continuous maze and shows the landscape shape around the forbidden regions.

Next, we consider the task of reaching the closest target among two equally rewarding positions $g_{0}$ and $g_{1}$, by setting $z_{R}=B\left(g_{0}\right)+B\left(g_{1}\right)$. The optimal Qfunction is not a linear combination of the $Q$-functions for $g_{0}$ and $g_{1}$. Fig. 7.5 shows successful trajectories generated by the policy $\pi_{z_{R}}$. On the contour plot of $\max _{a \in A} F\left(s, a, z_{R}\right)^{\top} z_{R}$ in Fig. 7.6 (right), the two rewarding positions appear as basins of attraction. Similar results for a third task are shown in Appendix E.2: introducing a "distracting" small reward next to the initial position of the agent, with a larger reward further away.

### 7.6.4 Embedding Visualizations

We visualize the learned FB state embeddings for Continuous Maze by projecting them into 2-dimensional space using t-SNE [Van der Maaten and Hinton, 2008] in Fig. 7.7. For the forward embeddings, we set $z=0$ corresponding to the uniform policy. We can see that FB partitions states according to the topology induced by the dynamics: states on opposite sides of walls are separated in the representation space and states on the same side lie together. Appendix E. 2 in-


Figure 7.5: Trajectories generated by the $F^{\top} B$ policies for the task of reaching the closest among two equally rewarding positions (star shapes $\hat{\sim}$ ). (Optimal $Q$-values are not linear over such mixtures.)


Figure 7.6: Contour plot of $\max _{a \in A} F\left(s, a, z_{R}\right)^{\top} z_{R}$ in Continuous Maze. Left: for the task of reaching a target while avoiding a forbidden region, Right: for two equally rewarding targets.
cludes embedding visualizations for different $z$ and for Discrete Maze and Ms. Pacman.

### 7.7 RELATED WORK

Borsa et al. [2018] learn optimal policies for rewards that are linear combinations of a finite number of feature functions provided in advance by the user. This approach cannot tackle generic rewards or goal-oriented RL: this would require introducing one feature per possible goal state, requiring infinitely many features in continuous spaces.

Our approach does not require user-provided features describing the future tasks, thanks to using successor states [Blier et al., 2021] where Borsa et al. [2018] use successor features. Schematically, and omitting actions, successor features start with user-provided features $\phi$, then learn $\psi$ such that $\psi\left(s_{0}\right)=$ $\sum_{t \geq 0} \gamma^{t} \mathbb{E}\left[\phi\left(s_{t}\right) \mid s_{0}\right]$. This limits applicability to rewards that are linear combinations of $\phi$. Here we use successor state probabilities, namely, we learn two


Figure 7.7: Visualization of FB embedding vectors on Continuous Maze after projecting them in two-dimensional space with t-SNE. Left: the states to be mapped. Middle: the $F$ embedding. Right: the $B$ embedding. The walls appear as large dents; the smaller dents correspond to the number of steps needed to get past a wall.
representations $F$ and $B$ such that $F\left(s_{0}\right)^{\top} B\left(s^{\prime}\right)=\sum_{t \geq 0} \gamma^{t} \operatorname{Pr}\left(s_{t}=s^{\prime} \mid s_{0}\right)$. This does not require any user-provided input.

We use a similar parameterization of policies by $F(s, a, z)^{\top} z$ as in Borsa et al. [2018], for similar reasons, although $z$ encodes a different object.

Successor representations where first defined in Dayan [1993] for finite spaces, corresponding to an older object from Markov chains, the fundamental matrix [Kemeny and Snell, 1960, Brémaud, 1999, Grinstead and Snell, 1997]. For successor representations in continuous spaces, a finite number of features $\phi$ are specified first; this can be used for generalization within a family of tasks, e.g., Barreto et al. [2017], Zhang et al. [2017], Grimm et al. [2019], Hansen et al. [2019]. Blier et al. [2021] moves from successor features to successor states by providing pointwise occupancy map estimates even in continuous spaces, without using the sparse reward $\mathbb{1}_{s_{t}=s^{\prime}}$. We borrow a successor state learning algorithm from Blier et al. [2021]. Blier et al. [2021] also introduced simpler versions of $F$ and $B$ for a single, fixed policy; Blier et al. [2021] does not consider the every-optimal-policy setting.

There is a long literature on goal-oriented RL. For instance, Schaul et al. [2015] learn goal-dependent value functions. Goal-dependent value functions have been investigated in earlier works such as Foster and Dayan [2002] and Sutton et al. [2011]. Hindsight experience replay (HER) [Andrychowicz et al., 2017] improve the sample efficiency of multiple goal learning with sparse rewards. A family of rewards has to be specified beforehand, such as reaching arbitrary target states. Specifying rewards a posteriori is not possible: for instance, learning to reach target states does not extend to reaching the nearest among several goals, reaching a goal while avoiding forbidden states, or maximizing any dense reward.

Hierarchical methods such as options [Sutton et al., 1999b] can be used for multi-task RL problems. However, RL training and planning on top of the options is still needed after the task is known.

For finite state spaces, Jin et al. [2020a] use reward-free interactions to build a training set that summarizes a finite environment, in the sense that any optimal policies later computed on this training set instead of the true environment are provably $\epsilon$-optimal, for any reward. They prove tight bounds on the necessary set size. Planning still has to be done afterwards for each reward.

### 7.8 EXTENDED RESUlTS: APPROXIMATE SOlutions and General Goals

This section provides additional theoretical insights and formal statements on approximate FB representations. All the proofs are provided in the appendix E.1. It is organized as follows:

- Section 7.8.1 formalizes the forward-backward representation with a goal or feature space.
- Section 7.8.2 establishes the existence of exact FB representations in finite spaces, and discusses the influence of the dimension $d$.
- Section 7.8.3 shows how approximate solutions provide approximately optimal policies.
- Section 7.8 .4 shows how $F$ and $B$ are successor and predecessor features of each other, and how the policies are optimal for rewards linearly spanned by $B$.
- Section 7.8.5 explains how to estimate $z_{R}$ at test time from a state distribution different from the training distribution.
- Section 7.8.6 presents a note of the measure $M^{\pi}$ and its density $m^{\pi}$.

Notation. In general, we denote by $M^{\pi}$ the successor measure of policy $\pi$ as defined in (7.3.1), and by $m^{\pi}$ its density, if it exists, with respect to a reference measure $\rho$. Namely,

$$
\begin{equation*}
M^{\pi}\left(s_{0}, a_{0}, \mathrm{~d} s, \mathrm{~d} a\right)=m^{\pi}\left(s_{0}, a_{0}, s, a\right) \rho(\mathrm{d} s, \mathrm{~d} a) \tag{7.8.1}
\end{equation*}
$$

Thus, the defining property of forward-backward representations (Definition 4) is $F\left(s_{0}, a_{0}, z\right)^{\top} B(s, a)=m^{\pi_{z}}\left(s_{0}, a_{0}, s, a\right)$.

We use the same convention for parametric models, with $M_{\theta}^{\pi}$ a measure and $m_{\theta}^{\pi}$ its density. The reference measure $\rho$ is fixed and may be unknown (typically $\rho$ is the distribution of state-actions in a training set or under an exploration policy).

### 7.8.1 The Forward-Backward Representation With a Goal or Feature Space

Here we state a generalization of Theorem 1 covering some extensions mentioned in the text.

First, this covers rewards known to only depend on certain features $g=$ $\phi(s, a)$ of the state-action $(s, a)$, where $\phi$ is a known function with values in some goal state $G$ (for instance, rewards depending only on some components of the state). Then it is enough to compute $B$ as a function of the goal $g$. Theorem 1 corresponds to $\phi=\mathrm{Id}$. This is useful to introduce prior information when available, resulting in a smaller model $(F, B)$.

This also recovers successor features as in Borsa et al. [2018], defined by userprovided features $\phi$. Indeed, fixing $B$ to Id and setting our $\phi$ to the $\phi$ of Borsa et al. [2018] (or fixing $B$ to their $\phi$ and our $\phi$ to Id) will represent the same set of rewards and policies as in Borsa et al. [2018], namely, optimal policies for rewards linear in $\phi$ (although with a slightly different learning algorithm and up to a linear change of variables for $F$ and $z$ given by the covariance of $\phi$, see Appendix 7.8.4). More generally, keeping the same $\phi$ but letting $B$ free (with larger $d$ ) can provide optimal policies for rewards that are arbitrary functions of $\phi$, linear or not.

For this, we extend successor state measures to values in goal spaces, representing the discounted time spent at each goal by the policy. Namely, given a policy $\pi$, let $M^{\pi}$ be the the successor state measure of $\pi$ over goals $g$ :

$$
\begin{equation*}
M^{\pi}(s, a, \mathrm{~d} g) \triangleq \sum_{t \geq 0} \gamma^{t} \operatorname{Pr}\left(\phi\left(s_{t}, a_{t}\right) \in \mathrm{d} g \mid s_{0}=s, a_{0}=a, \pi\right) \tag{7.8.2}
\end{equation*}
$$

for each state-action $(s, a)$ and each measurable set $\mathrm{d} g \subset G$. This will be the object approximated by $F(s, a, z)^{\top} B(g)$.

Second, we use a more general model of successor states: instead of $m \approx F^{\top} B$ we use $m \approx F^{\top} B+\bar{m}$ where $\bar{m}$ does not depend on the action, so that the $F^{\top} B$ part only computes advantages. This lifts the constraint that the model of $m$ has rank at most $d$, because there is no restriction on the rank on $\bar{m}$ : the rank restriction only applies to the advantage function.

Third, we state a form of policy improvement for the FB representation. Namely, the $Q$-function $F\left(s, a, z_{R}\right)^{\top} z_{R}$ for a given reward can be computed as a supremum over all values of $z$.

For simplicity we state the result with deterministic rewards, but this extends to stochastic rewards, because the expectation $z_{R}$ will be the same.

Definition 5 (Extended forward-backward representation of an MDP). Consider an MDP with state space $S$ and action space $A$. Let $\phi: S \times A \rightarrow G$ be a function from state-actions to some goal space $G=\mathbb{R}^{k}$.

Let $Z=\mathbb{R}^{d}$ be some representation space. Let

$$
\begin{equation*}
F: S \times A \times Z \rightarrow Z, \quad B: G \rightarrow Z, \quad \bar{m}: S \times Z \times G \rightarrow \mathbb{R} \tag{7.8.3}
\end{equation*}
$$

be three functions. For each $z \in Z$, define the policy

$$
\begin{equation*}
\pi_{z}(a \mid s) \triangleq \underset{a}{\operatorname{argmax}} F(s, a, z)^{\top} z . \tag{7.8.4}
\end{equation*}
$$

Let $\rho$ be any measure over the goal space $G$.
We say that $F, B$, and $\bar{m}$ are an extended forward-backward representation of the MDP with respect to $\rho$, if the following holds: for any $z \in Z$, any state-actions $(s, a)$, and any goal $g \in G$, one has

$$
\begin{equation*}
M^{\pi_{z}}(s, a, \mathrm{~d} g)=\left(F(s, a, z)^{\top} B(g)+\bar{m}(s, z, g)\right) \rho(\mathrm{d} g) \tag{7.8.5}
\end{equation*}
$$

where $M^{\pi_{z}}$ is the successor state measure (7.8.2) of policy $\pi_{z}$.
Theorem 2 (Forward-backward representation of an MDP, with features as goals). Consider an MDP with state space $S$ and action space $A$. Let $\phi: S \times A \rightarrow G$ be a function from state-actions to some goal space $G=\mathbb{R}^{k}$.

Let $F, B$, and $\bar{m}$ be an extended forward-backward representation of the MDP with respect to some measure $\rho$ over $G$.

Then the following holds. Let $R: S \times A \rightarrow \mathbb{R}$ be any bounded reward function, and assume that this reward function depends only on $g=\phi(s, a)$, namely, that there exists a function $r: G \rightarrow \mathbb{R}$ such that $R(s, a)=r(\phi(s, a))$. Set

$$
\begin{equation*}
z_{R} \triangleq \int_{g \in G} r(g) B(g) \rho(\mathrm{d} g) \tag{7.8.6}
\end{equation*}
$$

assuming the integral exists.
Then:

1. $\pi_{z_{R}}$ is an optimal policy for reward $R$ in the MDP.
2. For any $z \in Z$, the $Q$-function of policy $\pi_{z}$ for the reward $R$ is equal to

$$
\begin{equation*}
Q^{\pi_{z}}(s, a)=F(s, a, z)^{\top} z_{R}+\bar{V}^{z}(s) \tag{7.8.7}
\end{equation*}
$$

and the optimal $Q$-function $Q_{R}^{\star}$ is obtained when $z=z_{R}$ :

$$
\begin{equation*}
Q_{R}^{\star}(s, a)=F\left(s, a, z_{R}\right)^{\top} z_{R}+\bar{V}^{z_{R}}(s) . \tag{7.8.8}
\end{equation*}
$$

Here

$$
\begin{equation*}
\bar{V}^{z}(s) \triangleq \int_{g \in G} \bar{m}(s, z, g) r(g) \rho(\mathrm{d} g) \tag{7.8.9}
\end{equation*}
$$

and in particular $\bar{V}=0$ if $\bar{m}=0$.
The advantages $Q^{\pi_{z}}(s, a)-Q^{\pi_{z}}\left(s, a^{\prime}\right)$ do not depend on $\bar{V}$, so computing $\bar{V}$ is not necessary to obtain the policies.
3. If $\bar{m}=0$, then for any state-action $(s, a)$ one has

$$
\begin{equation*}
Q_{R}^{\star}(s, a)=F\left(s, a, z_{R}\right)^{\top} z_{R}=\sup _{z \in Z} F(s, a, z)^{\top} z_{R} . \tag{7.8.10}
\end{equation*}
$$

(We do not claim that $\bar{V}$ is the value function and $F^{\top} z_{R}$ the advantage function, only that the sum is the $Q$-function. When $\bar{m}=0$, the term $F^{\top} z_{R}$ is the whole $Q$-function.)

The last point of the theorem is a form of policy improvement. Indeed, by the second point, $F(s, a, z)^{\top} z_{R}$ is the estimated $Q$-function of policy $\pi_{z}$ for rewards $r$. This may be useful if $z_{R}$ falls outside of the training distribution for $F$ : then the values of $F\left(s, a, z_{R}\right)$ may not be safe to use. In that case, it may be useful to use a finite set $Z^{\prime} \subset Z$ of values of $z$ closer to the training distribution, and use the estimate $\sup _{z \in Z^{\prime}} F(s, a, z)^{\top} z_{R}$ instead of $F\left(s, a, z_{R}\right)^{\top} z_{R}$ for the optimal $Q$-function. A similar option has been used, e.g., in Borsa et al. [2018], but in the end it was not necessary in our experiments.

Remark 3. Formally, the statement holds for arbitrary $\rho$, but it only makes sense if $\rho$ has full support (or at least covers all reachable parts of the state space): (7.8.5) requires the support of $M^{\pi_{z}}$ to be included in that of $\rho$. Otherwise, FB representations may not exist and the statement is empty.

### 7.8.2 Existence of Exact $F B$ Solutions, Influence of Dimension $d$, Uniqueness

Existence of exact $F B$ representations in finite spaces. We now prove existence of an exact solution for finite spaces if the representation dimension $d$ is at least $\# S \times \# A$. Solutions are never unique: one may always multiply $F$ by an invertible matrix $C$ and multiply $B$ by $\left(C^{\top}\right)^{-1}$, see Remark 4 below (this allows us to impose orthonormality of $B$ in the experiments).

The constraint $d \geq \# S \times \# A$ can be largely overestimated depending on the tasks of interest, though. For instance, we prove below that in an $n$-dimensional toric grid $S=\{1, \ldots, k\}^{n}, d=2 n$ is enough to obtain optimal policies for reaching every target state (a set of tasks smaller than optimizing all possible rewards).

Proposition 6 (Existence of an exact $F B$ representation for finite state spaces). Assume that the state and action spaces $S$ and $A$ of an MDP are finite. Let $Z=\mathbb{R}^{d}$ with $d \geq \# S \times \#$. Let $\rho$ be any probability distribution on $S \times A$, with $\rho(s, a)>0$ for any $(s, a)$.

Then there exists $F: S \times A \times Z \rightarrow Z$ and $B: S \times A \rightarrow Z$, such that $F^{\top} B$ is equal to the successor state density of $\pi_{z}$ with respect to $\rho$ :

$$
\begin{equation*}
F^{\top}(s, a, z) B\left(s^{\prime}, a^{\prime}\right)=\sum_{t \geq 0} \gamma^{t} \frac{\operatorname{Pr}\left(\left(s_{t}, a_{t}\right)=\left(s^{\prime}, a^{\prime}\right) \mid s_{0}=s, a_{0}=a, \pi_{z}\right)}{\rho\left(s^{\prime}, a^{\prime}\right)} \tag{7.8.11}
\end{equation*}
$$

for any $z \in Z$ and any state-actions $(s, a)$ and $\left(s^{\prime}, a^{\prime}\right)$, where $\pi_{z}$ is defined as in Theorem 1 by $\pi_{z}(s)=\operatorname{argmax}_{a} F(s, a, z)^{\top} z$.

A small dimension $d$ can be enough for navigation: examples. In practice, even a small $d$ can be enough to get optimal policies for reaching arbitrary many states (as opposed to optimizing all possible rewards). Let us give an example with $S$ a toric $n$-dimensional grid of size $k$.

Let us start with $n=1$. Take $S=\{0, \ldots, k-1\}$ to be a length- $k$ cycle with three actions $a \in\{-1,0,1\}$ (go left, stay in place, go right). Take $d=2$, so that $Z=\mathbb{R}^{2} \simeq \mathbb{C}$.

We consider the tasks of reaching an arbitrary target state $s^{\prime}$, for every $s^{\prime} \in S$. Thus the goal state is $G=S$ in the notation of Theorem 2 , and $B$ only depends on $s^{\prime}$. The policy for such a reward is $\pi_{z_{R}}=\pi_{B\left(s^{\prime}\right)}$.

For a state $s \in\{0, \ldots, k-1\}$ and action $a \in\{-1,0,1\}$, define

$$
\begin{equation*}
F(s, a, z) \triangleq e^{2 i \pi(s+a) / k}, \quad B(s) \triangleq e^{2 i \pi s / k} \tag{7.8.12}
\end{equation*}
$$

Then one checks that $\pi_{B\left(s^{\prime}\right)}$ is the optimal policy for reaching $s^{\prime}$, for every $s^{\prime} \in S$. Indeed, $F\left(s, a, z_{R}\right)^{\top} z_{R}=\cos \left(2 \pi\left(s+a-s^{\prime}\right) / k\right)$. This is maximized for the action $a$ that brings $s$ closer to $s^{\prime}$.

So the policies will be optimal for reaching every target $s^{\prime} \in S$, despite the dimension being only 2.

By taking the product of $n$ copies of this example, this also works on the $n$ dimensional toric grid $S=\{0, \ldots, k-1\}^{n}$ with $2 n+1$ actions (add $\pm 1$ in each direction or stay in place), with a representation of dimension $d=2 n$ in $\mathbb{C}^{n}$, namely, by taking $B(s)_{j} \triangleq e^{2 i \pi s_{j} / k}$ for ecah direction $j$ and likewise for $F$. Then $\pi_{B\left(s^{\prime}\right)}$ is the optimal policy for reaching $s^{\prime}$ for every $s^{\prime} \in S$.

More generally, if one is only interested in the optimal policies for reaching states, then it is easy to show that there exist functions $F: S \times A \rightarrow \mathrm{Z}$ and $B: S \rightarrow Z$ such that the policies $\pi_{z}$ describe the optimal policies to reach each state: it is enough that $B$ be injective (typically requiring $d=\operatorname{dim}(S)$ ). Indeed, for any state $s \in S$, let $\pi_{s}^{\star}$ be the optimal policy to reach $s$. We want $\pi_{z}$ to be equal to $\pi_{s}^{\star}$ for $z=B(s)$ (the value of $z_{R}$ for a reward located at $s$ ). This translates as argmax ${ }_{a} F\left(s^{\prime}, a, B(s)\right)^{\top} B(s)=\pi_{s}^{\star}\left(s^{\prime}\right)$ for every other state $s^{\prime}$. This is realized just by letting $F$ be any function such that $F\left(s^{\prime}, \pi_{s}^{\star}\left(s^{\prime}\right), B(s)\right) \triangleq B(s)$ and $F\left(s^{\prime}, a, B(s)\right) \triangleq-B(s)$ for every other action $a$. As soon as $B$ is injective, there exists such a function $F$.

Let us turn to uniqueness of $F$ and $B$.
Remark 4. Let $C$ be an invertible $d \times d$ matrix. Given $F$ and $B$ as in Theorem 1, define

$$
\begin{equation*}
B^{\prime}(s, a) \triangleq C B(s, a), \quad F^{\prime}(s, a, z) \triangleq\left(C^{\top}\right)^{-1} F\left(s, a, C^{-1} z\right) \tag{7.8.13}
\end{equation*}
$$

together with the policies $\pi_{z}^{\prime}(s) \triangleq \operatorname{argmax}_{a} F^{\prime}(s, a, z)^{\top} z$. For each reward $r$, define $z_{R}^{\prime} \triangleq \mathbb{E}_{(s, a) \sim \rho}\left[r(s, a) B^{\prime}(s, a)\right]$.

Then this operation does not change the policies or estimated Q-values: for any reward, we have $\pi_{z_{R}^{\prime}}^{\prime}=\pi_{z_{R}}$, and $F^{\prime}\left(s, a, z_{R}^{\prime}\right)^{\top} z_{R}^{\prime}=F\left(s, a, z_{R}\right)^{\top} z_{R}$.

In particular, assume that the components of $B$ are linearly independent. Then, taking $C=\left(\mathbb{E}_{(s, a) \sim \rho} B(s, a) B(s, a)^{\top}\right)^{-1 / 2}, B^{\prime}$ is $L^{2}(\rho)$-orthonormal. So up to reducing the dimension d to $\operatorname{rank}(B)$, we can always assume that $B$ is $L^{2}(\rho)$-orthonormal.

Reduction to orthonormal $B$ will be useful in some proofs below. Even after imposing that $B$ be orthonormal, solutions are not unique, as one can still apply a rotation matrix on the variable $z$.

For a single policy $\pi_{z}$, the $F^{\top} B$ decomposition may be further standardized in several ways: after a linear change of variables in $\mathbb{R}^{d}$, and up to decreasing $d$ by removing unused directions in $\mathbb{R}^{d}$, one may assume either that $\operatorname{Cov}_{\rho} B=\mathrm{Id}$ and $\operatorname{Cov}_{\rho} F$ is diagonal, or that $\operatorname{Cov}_{\rho} F=\operatorname{Cov}_{\rho} B$ is diagonal, or write the decomposition as $\tilde{F}^{\top} D \tilde{B}$ with $D$ diagonal and $\operatorname{Cov}_{\rho} \tilde{F}=\operatorname{Cov}_{\rho} \tilde{B}=\mathrm{Id}$, thus corresponding to an approximation of the singular value decomposition of the successor measure $M^{\pi_{z}}$ in $L^{2}(\rho)$. However, since $B$ is shared between all values of $z$ and all policies $\pi_{z}$, it is a priori not possible to realize this for all $z$ simultaneously.

### 7.8.3 Approximate Solutions Provide Approximately Optimal Policies

Here we prove that the optimality in Theorems 1 and 2 is robust to approximation errors during training: approximate solutions still provide approximately optimal policies. We deal first with the approximation errors on ( $F, B$ ) during unsupervised training, then on $z_{R}$ during the reward estimation phase (in case the reward is not known explicity).

## Influence of Approximate $F$ and $B$ : Optimality Gaps are Proportional to $m^{\pi}-$ $F^{\top} B$

In continuous spaces, Theorems 1 and 2 are somewhat spurious: the equality $F^{\top} B=m$ will never hold exactly with finite representation dimension $d$. Instead, $F^{\top} B$ will only be a rank- $d$ approximation of $m$. Even in finite spaces, since $F$ and $B$ are learned by a neural network, we can only expect that $F^{\top} B \approx m$ in general. Therefore, we provide results that extend Theorems 1 and 2 to approximate training.

Optimality gaps are directly controlled by the error $m^{\pi}-F^{\top} B$ on the solution $F^{\top} B$. We provide this result for different notions of approximations between $m^{\pi}$ and $F^{\top} B$. First, in sup norm over $(s, a)$ but in expectation over $\left(s^{\prime}, a^{\prime}\right)$ (so that a
perfect model of the successor states $\left(s^{\prime}, a^{\prime}\right)$ of each $(s, a)$ is not necessary, only an average model). Second, for the weak topology on measures (this is the most relevant in continuous spaces: for instance, a Dirac measure can be approached by a continuous model in the weak topology). Finally, we provide pointwise estimates instead of only norms: for any reward, we show that the optimality gaps at each state can be bounded by an explicit matrix product directly involving the FB error matrix $m^{\pi}-F^{\top} B$ (Theorem 6).
$F$ and $B$ are trained such that $F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)$ approximates the successor state density $m^{\pi_{z}}\left(s, a, s^{\prime}, a^{\prime}\right)$. In the simplest case, we prove that if for some reward $R$,

$$
\begin{equation*}
\mathbb{E}_{\left(s^{\prime}, a^{\prime}\right) \sim \rho} \mid F\left(s, a, z_{R}\right)^{\top} B\left(s^{\prime}, a^{\prime}\right)-m^{\pi_{z_{R}}\left(s, a, s^{\prime}, a^{\prime}\right) \mid \leq \epsilon} \tag{7.8.14}
\end{equation*}
$$

for every $(s, a)$, then the optimality gap of policy $\pi_{z_{R}}$ is at most $(3 \epsilon /(1-$ $\gamma)$ ) sup $|R|$ for that reward (Theorem 5, first case).

In continuous spaces, $m^{\pi}$ is usually a distribution (Appendix 7.8.6), so such an approximation will not hold, and it is better to work on the measures themselves rather than their densities, namely, to compare $F^{\top} B \rho$ to $M^{\pi}$ instead of $F^{\top} B$ to $m^{\pi}$. We prove that if $F^{\top} B \rho$ is close to $M^{\pi}$ in the weak topology, then the resulting policies are optimal for any Lipschitz reward. ${ }^{2}$

Remember that a sequence of nonnegative measures $\mu_{n}$ converges weakly to $\mu$ if for any bounded, continuous function $f, \int f(x) \mu_{n}(d x)$ converges to $\int f(x) \mu(\mathrm{d} x)$ (Bogachev [2007], §8.1). The associated topology can be defined via the following Kantorovich-Rubinstein norm on nonnegative measures (Bogachev [2007], §8.3)

$$
\begin{align*}
\left\|\mu-\mu^{\prime}\right\|_{\mathrm{KR}} \triangleq & \sup \left\{\left|\int f(x) \mu(\mathrm{d} x)-\int f(x) \mu^{\prime}(\mathrm{d} x)\right|:\right. \\
& f \text { 1-Lipschitz function with sup }|f| \leq 1\} \tag{7.8.15}
\end{align*}
$$

where we have equipped the state-action space with any metric compatible with its topology. ${ }^{3}$

The following theorem states that if $F^{\top} B$ approximates the successor state density of the policy $\pi_{z}$ (for various sorts of approximations), then $\pi_{z}$ is approximately optimal. Given a reward function $r$ on state-actions, we denote

$$
\begin{equation*}
\|r\|_{\infty} \triangleq \sup _{(s, a) \in S \times A}|r(s, a)| \tag{7.8.16}
\end{equation*}
$$

[^7]and
\[

$$
\begin{equation*}
\|r\|_{\text {Lip }} \triangleq \sup _{(s, a) \neq\left(s^{\prime}, a^{\prime}\right) \in S \times A} \frac{r(s, a)-r\left(s^{\prime}, a^{\prime}\right)}{d\left((s, a),\left(s^{\prime}, a^{\prime}\right)\right)} \tag{7.8.17}
\end{equation*}
$$

\]

where we have chosen any metric on state-actions.
The first statement is for any bounded reward. The second statement only assumes an $F^{\top} B$ approximation in the weak topology but only applies to Lipschitz rewards. The third statement is more general and is how we prove the first two: weaker assumptions on $F^{\top} B$ work on a stricter class of rewards.

Theorem 5 (If $F$ and $B$ approximate successor states, then the policies $\pi_{z}$ yield approximately optimal returns). Let $F: S \times A \times Z \rightarrow Z$ and $B: S \times A \rightarrow Z$ be any functions, and define the policy $\pi_{z}(s)=\operatorname{argmax}_{a} F(s, a, z)^{\top} z$ for each $z \in Z$.

Let $\rho$ be any positive probability distribution on $S \times A$, and for each policy $\pi$, let $m^{\pi}$ be the density of the successor state measure $M^{\pi}$ of $\pi$ with respect to $\rho$. Let
$\hat{m}^{z}\left(s, a, s^{\prime}, a^{\prime}\right) \triangleq F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right), \quad \hat{M}^{z}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right) \triangleq \hat{m}^{z}\left(s, a, s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)$
be the estimates of $m$ and $M$ obtained via the model $F$ and $B$.
Let $r: S \times A \rightarrow \mathbb{R}$ be any bounded reward function. Let $V^{\star}$ be the optimal value function for this reward $r$. Let $\hat{V}^{\pi_{z}}$ be the value function of policy $\pi_{z}$ for this reward. Let $z_{R}=\mathbb{E}_{(s, a) \sim \rho}[r(s, a) B(s, a)]$.

Then:

1. If $\mathbb{E}_{\left(s^{\prime}, a^{\prime}\right) \sim \rho}\left|\hat{m}^{z_{R}}\left(s, a, s^{\prime}, a^{\prime}\right)-m^{\pi_{z_{R}}}\left(s, a, s^{\prime}, a^{\prime}\right)\right| \leq \epsilon$ for any $(s, a)$ in $S \times A$, then $\left\|\hat{V}^{\pi_{z_{R}}}-V^{\star}\right\|_{\infty} \leq 3 \epsilon\|r\|_{\infty} /(1-\gamma)$.
2. If $r$ is Lipschitz and $\left\|\hat{M}^{z_{R}}(s, a, \cdot)-M^{\pi_{z_{R}}}(s, a, \cdot)\right\|_{K R} \leq \epsilon$ for any $(s, a) \in S \times A$, then $\left\|\hat{V}^{\pi_{z_{R}}}-V^{\star}\right\|_{\infty} \leq 3 \epsilon \max \left(\|r\|_{\infty},\|r\|_{\text {Lip }}\right) /(1-\gamma)$.
3. More generally, let $\|\cdot\|_{A}$ be a norm on functions and $\|\cdot\|_{B}$ a norm on measures, such that $\int f \mathrm{~d} \mu \leq\|f\|_{A}\|\mu\|_{B}$ for any function $f$ and measure $\mu$. Then for any reward function $r$ such that $\|r\|_{A}<\infty$,

$$
\begin{equation*}
\left\|\hat{V}^{\pi_{z_{R}}}-V^{\star}\right\|_{\infty} \leq \frac{3\|r\|_{A}}{1-\gamma} \sup _{s, a}\left\|\hat{M}^{z_{R}}(s, a, \cdot)-M^{\pi_{z_{R}}}(s, a, \cdot)\right\|_{B} . \tag{7.8.19}
\end{equation*}
$$

Moreover, the optimal Q-function is close to $F\left(s, a, z_{R}\right)^{\top} z_{R}$ :

$$
\begin{equation*}
\sup _{s, a}\left|F\left(s, a, z_{R}\right)^{\top} z_{R}-Q^{\star}(s, a)\right| \leq \frac{2\|r\|_{A}}{1-\gamma} \sup _{s, a} \| \hat{M}^{z_{R}}(s, a, \cdot)-M^{\pi_{z_{R}}(s, a, \cdot) \|_{B}} \tag{7.8.20}
\end{equation*}
$$

Pointwise optimality gaps. We now turn to a more precise estimation of the optimality gap at each state, expressed directly as a matrix product involving the FB error $m^{\pi}-F^{\top} B$.

Here we assume that the state space is finite: this is not essential but simplifies notation since everything can be represented as matrices and vectors. Rcall that for each stochastic policy $\pi, P^{\pi}$ denotes the associated stochastic matrix on state-actions: $\left(P_{\pi}\right)_{(s a)\left(s^{\prime} a^{\prime}\right)}=P\left(s^{\prime} \mid s, a\right) \pi\left(s^{\prime}\right)\left[a^{\prime}\right]$. We also view rewards and $Q$-functions as vectors indexed by state-actions.

Theorem 6. Assume that the state space is finite. Let $F: S \times A \times \mathbb{R}^{d} \rightarrow \mathbb{R}^{d}$ and $B: S \times A \rightarrow \mathbb{R}^{d}$ be any functions. Define $\pi_{z}$ as in Theorem 1. Let $\rho>0$ be any distribution over state-actions, and let $m^{\pi}$ be the successor density (7.8.1) of policy $\pi$ with respect to $\rho$.

For each $z \in \mathbb{R}^{d}$, define the $F B$ error $E(z)$ as a matrix over state-actions:

$$
\begin{equation*}
E(z)_{(s a)\left(s^{\prime} a^{\prime}\right)} \triangleq m^{\pi_{z}}\left(s, a, s^{\prime}, a^{\prime}\right)-F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right) . \tag{7.8.21}
\end{equation*}
$$

Let $r$ be any reward function and let $Q^{\star}$ and $\pi^{\star}$ be its optimal $Q$-function and policy. Let $z_{R}=\mathbb{E}_{\rho}[r . B]$ as in Theorem 1, and let $Q^{\pi_{z_{R}}}$ be the Q-function of policy $\pi_{z_{R}}$.

Then we have the componentwise inequality between state-action vectors

$$
\begin{equation*}
0 \leq Q^{\star}-Q^{\pi_{z_{R}}} \leq\left(\sum_{t \geq 0} \gamma^{t+1}\left(P^{\pi^{\star}}\right)^{t}\right)\left(P^{\pi^{\star}}-P^{\pi_{z_{R}}}\right) E\left(z_{R}\right) \operatorname{diag}(\rho) r . \tag{7.8.22}
\end{equation*}
$$

In particular, if $E\left(z_{R}\right)=0$ then $\pi_{z_{R}}$ is optimal for reward $r$.
The matrix $\sum_{t \geq 0} \gamma^{t+1}\left(P^{\pi^{\star}}\right)^{t}$ represents states visited along the optimal trajectory starting at the initial state. Multiplying by $\left(P^{\pi^{\star}}-P^{\pi_{z_{R}}}\right)$ visits states one step away from this trajectory. Therefore, what matters are the values of the FB error matrix $E\left(z_{R}\right)_{(s a)\left(s^{\prime} a^{\prime}\right)}$ at state-actions $(s, a)$ one step away from the optimal trajectories.

In unsupervised RL, we have no control over the first part $\sum_{t \geq 0} \gamma^{t+1}\left(P^{\pi^{\star}}\right)^{t}$, which depends only on the optimal trajectories of the unknown future tasks $r$. Therefore, in general it makes sense to just minimize $E(z)$ in some matrix norm, for as many values of $z$ as possible. (The choice of matrix norm influences which rewards will be best optimized, as illustrated by Theorem 5.)

## An approximate $z_{R}$ yields an approximately optimal policy

We now turn to the second source of approximation: computing $z_{R}$ in the reward estimation phase. This is a problem only if the reward is not specified explicitly.

We deal in turn with the effect of using a model of the reward function, and the effect of estimating $z_{R}=\mathbb{E}_{(s, a) \sim \rho} r(s, a) B(s, a)$ via sampling.

Which of these options is better depends a lot on the situation. If training of $F$ and $B$ is perfect, by construction the policies are optimal for each $z_{R}$ : thus, estimating $z_{R}$ from rewards sampled at $N$ states $\left(s_{i}, a_{i}\right) \sim \rho$ will produce the optimal policy for exactly that empirical reward, namely, a nonzero reward at each $\left(s_{i}, a_{i}\right)$ but zero everywhere else, thus overfitting the reward function. Reducing the dimension $d$ reduces this effect, since rewards are projected on the span of the features in $B$ : $B$ plays both the roles of a transition model and a reward regularizer. This appears as a $\sqrt{d / N}$ factor in Theorem 7 below.

Thus, if both the number of samples to train $F$ and $B$ and the number of reward samples are small, using a smaller $d$ will regularize both the model of the environment and the model of the reward. However, if the number of samples to train $F$ and $B$ is large, yielding an excellent model of the environment, but the number of reward samples is small, then learning a model of the reward function will be a better option than direct empirical estimation of $z_{R}$.

The first result below states that reward misidentification comes on top of the approximation error of the $F^{\top} B$ model. This is relevant, for instance, if a reward model $\hat{r}$ is estimated by an external model using some reward values.

Proposition 7 (Influence of estimating $z_{R}$ by an approximate reward). Let $\hat{r}$ : $S \times$ $A \rightarrow \mathbb{R}$ be any reward function. Let $\hat{z}_{R}=\mathbb{E}_{(s, a) \sim \rho}[\hat{r}(s, a) B(s, a)]$.

Let $\epsilon_{F B}$ be the error attained by the $F^{\top} B$ model in Theorem 5 for reward $\hat{r}$; namely, assume that $\left\|\hat{V}^{\pi_{2_{R}}}-\hat{V}^{\star}\right\|_{\infty} \leq \epsilon_{F B}$ with $\hat{V}^{\star}$ the optimal value function for $\hat{r}$.

Then the policy $\pi_{\hat{z}_{R}}(s)=\operatorname{argmax}_{a} F\left(s, a, \hat{z}_{R}\right)^{\top} \hat{z}_{R}$ defined by the model $\hat{r}$ is $\left(\frac{2\|r-\hat{r}\|_{\infty}}{(1-\gamma)}+\epsilon_{F B}\right)$-optimal for reward $r$.

This still assumes that the expectation $\hat{z}_{R}=\mathbb{E}_{(s, a) \sim \rho}[\hat{r}(s, a) B(s, a)]$ is computed exactly for the model $\hat{r}$. If $\hat{r}$ is given by an explicit model, this expectation can in principle be computed on the whole replay buffer used to train $F$ and $B$, so variance would be low. Nevertheless, we provide an additional statement which covers the influence of the variance of the estimator of $z_{R}$, whether this estimator uses an external model $\hat{r}$ or a direct empirical average reward observations $r(s, a) B(s, a)$.

Definition 6. The skewness $\zeta(B)$ of $B$ is defined as follows. Assume $B$ is bounded. Let $B_{1}, \ldots, B_{d}: S \times A \rightarrow \mathbb{R}$ be the functions of $(s, a)$ defined by each component of $B$. Let $\langle B\rangle$ be the linear span of the $\left(B_{i}\right)_{1 \leq i \leq d}$ as functions on $S \times A$. Set

$$
\begin{equation*}
\zeta(B) \triangleq \sup _{f \in\langle B\rangle, f \neq 0} \frac{\|f\|_{\infty}}{\|f\|_{L^{2}(\rho)}} \tag{7.8.23}
\end{equation*}
$$

Theorem 7 (Influence of estimating $z_{R}$ by empirical averages). Assume that $z_{R}=$ $\mathbb{E}_{(s, a) \sim \rho}[r(s, a) B(s, a)]$ is estimated via

$$
\begin{equation*}
\hat{z}_{R} \triangleq \frac{1}{N} \sum_{i=1}^{N} \hat{r}_{i} B\left(s_{i}, a_{i}\right) \tag{7.8.24}
\end{equation*}
$$

using $N$ independent samples $\left(s_{i}, a_{i}\right) \sim \rho$, where the $r_{i}$ are random variables such that $\mathbb{E}\left[\hat{r}_{i} \mid s_{i}, a_{i}\right]=r\left(s_{i}, a_{i}\right), \operatorname{Var}\left[\hat{r}_{i} \mid s_{i}, a_{i}\right] \leq v$ for some $v \in \mathbb{R}$, and the $\hat{r}_{i}$ are mutually independent given $\left(s_{i}, a_{i}\right)_{i=1, \ldots, N}$.

Let $V^{\star}$ be the optimal value function for reward $r$, and let $\hat{V}$ be the value function of the estimated policy $\pi_{\hat{z}_{R}}$ for reward $r$.

Then, for any $\delta>0$, with probability at least $1-\delta$,

$$
\begin{equation*}
\left\|\hat{V}-V^{\star}\right\|_{\infty} \leq \epsilon_{F B}+\frac{2}{1-\gamma} \sqrt{\frac{\zeta(B) d}{N \delta}\left(v+\left\|r(s, a)-\mathbb{E}_{\rho} r\right\|_{L^{2}(\rho)}^{2}\right)} \tag{7.8.25}
\end{equation*}
$$

which is therefore the bound on the optimality gap of $\pi_{\hat{2}_{R}}$ for $r$. Here $\epsilon_{F B}$ is the error due to the $F^{\top} B$ model approximation, defined as in Proposition 7.

The proofs are not direct, because $F$ is not continuous with respect to $z$. Contrary to $Q$-values, successor states are not continuous in the reward: if an action has reward 1 and the reward for another action changes from $1-\epsilon$ to $1+\epsilon$, the return values change by at most $2 \epsilon$, but the actions and states visited by the optimal policy change a lot. So it is not possible to reason by continuity on each of the terms involved.

### 7.8.4 $F$ and $B$ as Successor and Predecessor Features of Each Other, Optimality for Rewards in the Span of $B$

We now give two statements. The first encodes the idea that $F$ encodes the future of a state while $B$ encodes the past of a state.

The second proves that the $F B$ policies are optimal for any reward that lies in the linear span of the features learned in $B$; for rewards out of this span, it is best if the features in $B$ are spatially smooth.

The intuition that $F$ and $B$ encode the future and past of states is formalized as follows: if $F$ and $B$ minimize their unsupervised loss, then $F$ is equal to the successor features from the dual features of $B$, and $B$ is equal to the predecessor features from the dual features of $F$ (Theorem 8).

This statement holds for a fixed $z$ and the corresponding policy $\pi_{z}$. So, for the rest of this section, $z$ is fixed.

By "dual" features we mean the following. Define the $d \times d$ covariance matrices

$$
\begin{equation*}
\operatorname{Cov} F \triangleq \mathbb{E}_{(s, a) \sim \rho}\left[F(s, a, z) F(s, a, z)^{\top}\right], \quad \operatorname{Cov} B \triangleq \mathbb{E}_{(s, a) \sim \rho}\left[B(s, a) B(s, a)^{\top}\right] . \tag{7.8.26}
\end{equation*}
$$

Then $(\operatorname{Cov} F)^{-1 / 2} F(s, a, z)$ is $L^{2}(\rho)$-orthonormal and likewise for $B$. The "dual" features are $(\operatorname{Cov} F)^{-1} F(s, a, z)$ and $(\operatorname{Cov} B)^{-1} B(s, a)$, without the square root: these are the least square solvers for $F$ and $B$ respectively, and these are the ones that appear below.

The unsupervised FB loss for a fixed $z$ is

$$
\begin{align*}
\ell(F, B) & \triangleq \int\left|F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)-\sum_{t \geq 0} \gamma^{t} \frac{P_{t}\left(\mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime} \mid s, a, \pi_{z}\right)}{\rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)}\right|^{2} \rho(\mathrm{~d} s, \mathrm{~d} a) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)  \tag{7.8.27}\\
& =\left\|F(\cdot, z)^{\top} B(\cdot)-m^{\pi_{z}}(\cdot, \cdot)\right\|_{L^{2}(\rho) \otimes L^{2}(\rho)}^{2} \tag{7.8.28}
\end{align*}
$$

Thus, minimizers in dimension $d$ correspond to an SVD of the successor state density in $L^{2}(\rho)$, truncated to the largest $d$ singular values.

Theorem 8. Consider a smooth parametric model for $F$ and $B$, and assume this model is overparameterized. ${ }^{4}$ Also assume that the data distribution $\rho$ has positive density everywhere.

Let $z \in Z$. Assume that for this $z, F$ and $B$ lie in $L^{2}(\rho)$ and achieve a local extremum of $\ell(F, B)$ within this parametric model. Namely, the derivative $\partial_{\theta} \ell(F, B)$ of the loss with respect to the parameters $\theta$ of $F$ is 0 , and likewise for $B$.

Then $F$ is equal to $(\operatorname{Cov} B)^{-1}$ times the successor features of $B$ : for any $(s, a) \in$ $S \times A$,

$$
\begin{equation*}
(\operatorname{Cov} B) F(s, a, z)=\sum_{t \geq 0} \gamma^{t} \int_{\left(s^{\prime}, a^{\prime}\right)} P_{t}\left(\mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime} \mid s, a, \pi_{z}\right) B\left(s^{\prime}, a^{\prime}\right) \tag{7.8.29}
\end{equation*}
$$

and $B$ is equal to $(\operatorname{Cov} F)^{-1}$ times the predecessor features of $F$ :

$$
\begin{equation*}
(\operatorname{Cov} F) B\left(s^{\prime}, a^{\prime}\right)=\sum_{t \geq 0} \gamma^{t} \int_{(s, a)} \frac{P_{t}\left(\mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime} \mid s, a, \pi_{z}\right)}{\rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)} F(s, a, z) \rho(\mathrm{d} s, \mathrm{~d} a) \tag{7.8.30}
\end{equation*}
$$

$\rho$-almost everywhere. Here the covariances have been defined in (7.8.26), and $P_{t}(\cdot \mid s, a, \pi)$ denotes the law of $\left(s_{t}, a_{t}\right)$ under trajectories starting at $(s, a)$ and following policy $\pi$.

The same result holds when working with features $\phi\left(s^{\prime}, a^{\prime}\right)$, just by applying it to $B \circ \phi$.

[^8]Note that in the FB framework, we may normalize either F or B (Remark 4), but not both.

As a consequence of Theorem 8, we characterize below which kind of rewards we can capture if we fix $B$ and train only $F$ to minimize the unsupervised loss given $B$. Namely, we show that for any reward $r$, the resulting policy is optimal for the $L^{2}(\rho)$-orthogonal projection of $r$ onto the span of $B$.

Theorem 9 (Optimizing $F$ for a given $B$; influence of the span of $B$ ). Let $B$ a fixed function in $L^{2}\left(\rho, \mathbb{R}^{d}\right)$. Define the span of $B$ as the set of functions $w^{\top} B \in L^{2}(\rho, \mathbb{R})$ when $w$ ranges in $\mathbb{R}^{d}$.

Consider a smooth parametric model for $F$, and assume this model is overparameterized. Assume that $F$ lies in $L^{2}(\rho)$ and achieves a local extremum of $\ell(F, B)$ within this parametric model for each $z \in Z$.

Assume that the data distribution $\rho$ has positive density everywhere.
Then for any bounded reward function $r: S \times A \rightarrow \mathbb{R}$ that lies in the span of $B, \pi_{z_{R}}$ is an optimal policy for the reward $r$.

More generally, for any bounded reward function $r: S \times A \rightarrow \mathbb{R}$, the policy $\pi_{z_{R}}$ is an optimal policy for the reward $r_{B}$ defined as the $L^{2}(\rho)$-orthogonal projection of $r$ onto the span of $B$.

Moreover,

$$
\begin{equation*}
\left\|V^{\pi_{z_{R}}}-V^{\star}\right\|_{\infty} \leq \frac{2}{1-\gamma}\left\|r-r_{B}\right\|_{\infty} \tag{7.8.31}
\end{equation*}
$$

with $V^{\star}$ the optimal value function for $r$. More precisely,

$$
\begin{equation*}
\left\|V^{\pi_{z_{R}}-V^{\star}}\right\|_{\infty} \leq\left\|\left(\operatorname{Id}-\gamma P_{\pi^{\star}}\right)^{-1}\left(r-r_{B}\right)\right\|_{\infty}+\left\|\left(\operatorname{Id}-\gamma P_{\pi_{z_{R}}}\right)^{-1}\left(r-r_{B}\right)\right\|_{\infty} \tag{7.8.32}
\end{equation*}
$$

with $\pi^{\star}$ the optimal policy for reward $r$, and notation $P_{\pi}$ as in Theorem 6.
The last bounds implies the previous one, as $\left(\operatorname{Id}-\gamma P_{\pi}\right)^{-1}$ is bounded by $\frac{1}{1-\gamma}$ in $L^{\infty}$ norm for any policy $\pi$.

Discussion: optimal $B$ and priors on rewards. This bound is interesting if $r-r_{B}$ is small, namely, if $B$ captures most of the components of the reward functions we are interested in. But even for rewards not spanned by $B$, the bound is smaller if $r-r_{B}$ avoids the largest eigendirections of $\left(\operatorname{Id}-P_{\pi}\right)^{-1}$ for various policies $\pi$, namely, if $B$ captures these largest eigendirections. These eigendirections are those of $P_{\pi}$ : so the bound will be small if $B$ contains the largest eigendirections of $P_{\pi}$ for various policies $\pi$, corresponding to spatially continuity of functions under transitions in the environment ( $P_{\pi} f$ close to $f$ ).

If we are interested in spatially smooth rewards $r$, then $r-r_{B}$ is small if $B$ if captures smooth functions first. But even for rewards not spanned by $B$, and for non-spatially smooth rewards (e.g., goal-oriented problems with reward
$\mathbb{1}_{s=\text { goal }}$ ), the bound (7.8.32) shows that $B$ should first capture spatially smooth eigenvectors of many policies $\pi$.

Is this a natural consequence of FB training? Up to some approximations, yes. For a single $z$ and policy $\pi_{z}$, the loss (7.8.27) used to train $B$ is optimal when $B$ captures the largest singular directions of $\left(\operatorname{Id}-P_{\pi_{z}}\right)^{-1}$, which is slightly different. The optimal policy $\left(\operatorname{Id}-P_{\pi^{\star}}\right)^{-1}$ is not represented in the criterion, and it is not clear to what extent spatial continuity with respect to $P_{\pi_{z}}$ or $P_{\pi^{\star}}$ differ. Moreover, in the full algorithm, $B$ is shared between several $z$ and several policies. So we have no rigorous result here. Still, the intuition shows that FB training goes in the correct direction.

### 7.8.5 Estimating $z_{R}$ from a Different State Distribution at Test Time

Theorem 1 requires $z_{R}=\mathbb{E}_{(s, a) \sim \rho}[r(s, a) B(s, a)]$ to be estimated using rewards observed from the same state-action distribution $\rho$ as the one used to define $m^{\pi}$ and train $F$ and B. The algorithm of Section 7.5 computes $m^{\pi}$ with respect to the distribution $\rho$ of the training set. So in general, estimating $z_{R}$ requires either being able to run the exploration policy again once reward samples are available, or being able to explicitly estimate $r(s, a)$ on states stored in the training set.

However, at test time, we would generally like to use the policy learned, rather than the exploration policy. This will result in a distribution of statesactions $\rho_{\text {test }}(s, a)$ different from the training distribution $\rho(s, a)$.

If the training set remains accessible, a generic solution is to train a model $\hat{r}(s, a)$ of the reward function, from rewards $r(s, a)$ observed at test time under any distribution $(s, a) \sim \rho_{\text {test }}$. Then one can estimate $z_{R}$ by averaging this model $\hat{r}$ over state-actions sampled from the training set:

$$
\begin{equation*}
\hat{z}_{R} \triangleq \mathbb{E}_{(s, a) \sim \rho}[\hat{r}(s, a) B(s, a)] \tag{7.8.33}
\end{equation*}
$$

However, the training set may not be available anymore at test time. But as it turns out, if we use for $\hat{r}$ a linear model based on the features learned in $B$, then we do not need to store the training set: it is enough to estimate the matrix $\operatorname{Cov}_{\rho}(B)$, which can be pre-computed during training. This is summarized in the following.

Proposition 8. Let $\hat{r}$ be the linear model of rewards computed at test time by linear regression of the reward $r$ over the components $B_{1}, \ldots, B_{d}$ of $B$, with state-actions taken from a test distribution $\rho_{\text {test }}$. Explicitly,

$$
\begin{equation*}
\hat{r}(s, a) \triangleq B(s, a)^{\top} w, \quad w \triangleq\left(\operatorname{Cov}_{\rho_{\text {test }}} B\right)^{-1} \mathbb{E}_{(s, a) \sim \rho_{\text {test }}}[r(s, a) B(s, a)] \tag{7.8.34}
\end{equation*}
$$

Here we assume that $\operatorname{Cov}_{\rho_{\text {test }}} B=\mathbb{E}_{(s, a) \sim \rho_{\text {test }}}\left[B(s, a) B(s, a)^{\top}\right]$ is invertible.
Then the estimate $\hat{z}_{R}$ computed by using this model $\hat{r}$ in (7.8.33) is

$$
\begin{equation*}
\hat{z}_{R}=\left(\operatorname{Cov}_{\rho} B\right)\left(\operatorname{Cov}_{\rho \text { test }} B\right)^{-1} \mathbb{E}_{(s, a) \sim \rho_{\text {test }}}[r(s, a) B(s, a)] . \tag{7.8.35}
\end{equation*}
$$


For this estimate, $\operatorname{Cov}_{\rho_{\text {test }}} B$ and $\mathbb{E}_{(s, a) \sim p_{\text {test }}}[r(s, a) B(s, a)]$ can be computed at test time, while the matrix $\operatorname{Cov}_{\rho} B$ must be computed at training time. This way, the training set can be discarded.

If the estimate (7.8.34) is used, the learned policies correspond to using universal successor features approximators Borsa et al. [2018] on top of the features learned by $B$. Indeed, universal successor features with features $\phi$ use the policies $\operatorname{argmax}_{a} \psi(s, a, w)^{\top} w$ with $w=\left(\operatorname{Cov}_{\rho_{\text {test }}} \phi\right)^{-1} \mathbb{E}_{(s, a) \sim \rho_{\text {test }}}[r(s, a) \phi(s, a)]$ the regression vector of $r$ over the features $\phi$, and with $\psi=\operatorname{Succ}(\phi)$ the successor features of $\phi$. Here we use the policies $\pi_{\hat{z}_{R}}=\operatorname{argmax}{ }_{a} F\left(s, a, \hat{z}_{R}\right)^{\top} \hat{z}_{R}$. Let us set $\phi \triangleq B$ as the base features in successor features. Then the above shows that $\hat{z}_{R}=\left(\operatorname{Cov}_{\rho} B\right) w$ when using the linear model of rewards. Moreover, we proved in Theorem 8 that the optimum for $F$ is $F=\left(\operatorname{Cov}_{\rho} B\right)^{-1} \operatorname{Succ}(B)$. Therefore, the policies coincide in this situation.

Thus, universal successor features based on $B$ appear as a particular case if a linear model of rewards is used at test time, although in general any reward model may be used in (7.8.33).

### 7.8.6 A Note on the Measure $M^{\pi}$ and its Density $m^{\pi}$

In finite spaces, the definition of the successor state density $m^{\pi}$ via

$$
\begin{equation*}
M^{\pi}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right)=m^{\pi}\left(s, a, s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right) \tag{7.8.36}
\end{equation*}
$$

with respect to the data distribution $\rho$ poses no problem, as long as the data distribution is positive everywhere.

In continuous spaces, this can be understood as the (Radon-Nikodym) density of $M^{\pi}$ with respect to $\rho$, assuming $M^{\pi}$ has no singular part with respect to $\rho$. However, this is never the case: in the definition (7.3.1) of the successor state measure $M^{\pi}$, the term $t=0$ produces a Dirac measure $\delta_{s, a}$. So $M^{\pi}$ has a singular component due to $t=0$, and $m^{\pi}$ is better thought of as a distribution.

When $m^{\pi}$ is a distribution, a continuous parametric model $m_{\theta}^{\pi}$ learned by (7.5.2) can approximate $m^{\pi}$ in the weak topology only: $m_{\theta}^{\pi} \rho$ approximates $M^{\pi}$ for the weak convergence of measures. Thus, for the forward-backward representation, $F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)$ weakly approximates $M^{\pi_{z}}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right)$.

We have not found this to be a problem either in theory or practice, and Theorem 5 covers weak approximations.

Alternatively, one may just define successor states starting at $t=1$ in (7.3.1). This only works well if rewards $r(s, a)$ depend on the state $s$ but not the action $a$ (e.g., in goal-oriented settings). If starting the definition at $t=1, m^{\pi}$ is an ordinary function provided the transition kernels $P\left(\mathrm{~d} s^{\prime} \mid s, a\right)$ of the environment are non-singular, $\rho$ has positive density, and $\pi(\mathrm{d} a \mid s)$ is non-singular as well. Starting at $t=1$ induces the following changes in the theorems:

- In the learning algorithm (7.5.2) for successor states, the term $\partial_{\theta} m_{\theta}^{\pi}(s, a, s, a)$ becomes $\gamma \partial_{\theta} m_{\theta}^{\pi}\left(s, a, s^{\prime}, a^{\prime}\right)$.
- The expression for the $Q$-function in Theorem 1 becomes $Q^{\star}(s, a)=$ $r(s, a)+F\left(s, a, z_{R}\right)^{\top} z_{R}$, and likewise in Theorem 2. The $r(s, a)$ term covers the immediate reward at a state, since we have excluded $t=0$ from the definition of successor states.
- In general the expression for optimal policies becomes

$$
\begin{equation*}
\pi_{z}(s) \triangleq \underset{a}{\operatorname{argmax}}\left\{r(s, a)+F(s, a, z)^{\top} z\right\} \tag{7.8.37}
\end{equation*}
$$

which cannot be computed from $z$ and $F$ alone in the unsupervised training phase. The algorithm only makes sense for rewards that depend on $s$ but not on $a$ (e.g., in goal-oriented settings): then the policy $\pi_{z}$ is equal to $\pi_{z}(s) \triangleq \operatorname{argmax}_{a} F(s, a, z)^{\top} z$ again.

### 7.9 CONCLUSION

The FB representation is a learnable mathematical object that "summarizes" a reward-free MDP. It provides near-optimal policies for any reward specified a posteriori, without planning. It is learned from black-box reward-free interactions with the environment. In practice, this unsupervised method performs comparably to goal-oriented methods for reaching arbitrary goals, but is also able to tackle more complex rewards in real time.

## 8 Conclusion

### 8.1 SUMMARY OF CONTRIBUTIONS

The aim of this thesis was to provide theoretical solutions and insights to some fundamental issues in modern RL. This dissertation touched upon three major research areas: off-policy learning, exploration-exploitation dilemma in large state spaces, and self-supervision in RL.

- Off-policy learning is one of the main keys in scaling up RL agents as it enables the use of other streams of experiences than the one generated by themselves. In chapter 3, we showed how we can combine multi-step updates, for a rapid propagation of the errors, with function approximation in a provable way. We also provided a state-of-the-art finite sample analysis of gradient TD methods for policy evaluation. In chapter 4, we showed how to better exploit off-policy data to stabilize the policy improvement step. In particular, we proposed to learn in off-policy way the divergence between state-action visitations, which acts then as regularizer for the policy update. This improves upon conservative approaches which rely only on immediate action distributions for the regularization.
- Balancing between exploration and exploitation is critical for dataefficient learning i.e reaching optimal solution within a decent number of interactions. In this thesis, we focused on large state spaces. In chapter 5, we assumed that the optimal Q-function is smooth and we introduced ZOOMRL, an efficient online algorithm that learns a data-driven partitioning and we proved that its regret scales with the covering dimension and does not depend on the number of states. In chapter 6, we assumed that both reward and transition kernel are linear and non-stationary. We introduced OPT-WLSVI, an efficient online algorithm that deploys a linear function approximation and uses discounting strategy to deal with the non-stationarity. We also fixed a serious flaw in the regret analysis of the non-stationary linear bandits, propagated in the literature.
- We focused in chapter 7 on self-supervision in reward-free environments. We proposed to learn representations, using interactions without reward signal, that encode a predictive occupancy map over states and actions. We
showed that we can read directly on these representation optimal policies for all reward functions.

A recurrent theme in the presented contributions is generalisation. In chapter 3, we devised convergent algorithms for policy evaluation that use off-policy multi-step traces while generalizing over states via a linear function approximator. In chapter 5, under smoothness assumptions, our proposed algorithm ZOOMRL generalizes over large state-action spaces by learning a data-driven partitioning. In chapter 6, our algorithm OPT-WLSVI generalizes over states by using linear function approximation in a non-stationary linear environment. Finally, in chapter 7 we focused on another kind of generalization which is the generalization over reward functions. Our learned forward-backward representations can return optimal behavior for any reward function or any goal state by computing a simple average.

### 8.2 FUTURE RESEARCH

To build reliable RL systems, there is far more to be done and I am eager to explore in the next few years some promising directions that I describe below.

### 8.2.1 Towards Fully Controllable Agents

Unsupervised learning provides a useful paradigm for the autonomous acquisition of task-agnostic knowledge. Our last-mentioned work in Chapter 7 is a step towards building agents that are fully controllable after first exploring their environment in an unsupervised way. While we tackled the representation learning problem, the exploration aspect remains a missing piece. The lack of exploration could limit the applicability of the learned representation to complex downstream tasks.

The open question that I would like to investigate is
Could we devise a principled unsupervised objective to discover behaviors or skills that are useful for both exploration (acquisition task-agnostic knowledge) and exploitation (solving a downstream task)?

### 8.2.2 Optimization for RL

While RL methods are traditionally dynamic-programming (DP) based algorithms, the field witnesses a recent trend that frames RL problems as well-posed optimization problems. Such algorithms have the benefit of avoiding inherent instability of DP and are more compatible with nonlinear/ linear function approximation. My first contribution in Chapter 3 lies in this direction. I would
like to investigate to what extent we can learn from established optimization tools to tackle RL problems.

To give a concrete example, I plan to look at the linear programming (LP) formulation of RL. This LP problem can be cast into a primal-dual problem where primal and dual variables are value functions and occupancy measures respectively. I would like to propose optimization-based algorithms tailored to this problem and study their convergence rate and sample complexity.

### 8.2.3 Statistical RL with General Function Approximation

It has been shown empirically that combining RL algorithms with neural network function approximators could lead to impressive performance on various tasks. However, the design of provable efficient RL algorithm with such general function classes remains elusive. In the past few years, there was some progress on this topic: Jiang et al. [2017] and Dong et al. [2020] introduce algorithms whose performance guarantees can be upper bounded in terms of the Bellman rank. However, these algorithms are not computationally efficient. Ayoub et al. [2020] and Wang et al. [2020] propose algorithms whose regret bounds scale with the Eluder dimension. However, we don't know if there are interesting function classes, beyond linear and generalized linear functions, with finite Eluder dimension.

In order to understand the role of function approximation in RL, the research question I would like to explore is

Could we develop better capacity measures of function class, adapted to the rich temporal structures of $R L$ ?

# Appendix from Chapter 3 

## A. 1 Proof of Proposition 1

We compute $\mathbb{E}\left[A_{k}\right]$ and $\mathbb{E}\left[b_{k}\right]$ where expectation are over trajectories drawn by executing the behavior policy: $s_{k}, a_{k}, r_{k}, s_{k+1}, \ldots s_{t}, a_{t}, r_{t}, s_{t+1} \ldots$ where $s_{k}, a_{k} \sim$ $d, \quad r_{t}=r\left(s_{t}, a_{t}\right), \quad s_{t+1} \sim p\left(\cdot \mid s_{t}, a_{t}\right)$. We note that under stationarity of $d$, $\mathbb{E}\left[A_{k}\right]=E\left[A_{0}\right]$ and $\mathbb{E}\left[b_{k}\right]=\mathbb{E}\left[b_{0}\right]$. Let $\theta, \theta^{\prime} \in \mathbb{R}^{d}$ and let $Q=\Phi \theta$ and $Q^{\prime}=\Phi \theta^{\prime}$ their respective Q -functions.

$$
\begin{aligned}
& \theta^{\prime \top} \mathbb{E}\left[A_{k}\right] \theta \\
& =\mathbb{E}\left[\sum_{t=0}^{\infty}(\lambda \gamma)^{t}\left(\prod_{i=1}^{t} \kappa_{i}\right) Q^{\prime}\left(s_{0}, a_{0}\right)\left[\gamma \mathbb{E}_{\pi} Q\left(s_{t+1}, .\right)-Q\left(s_{t}, a_{t}\right)\right]^{\top}\right] \\
& =\sum_{t=0}^{\infty}(\lambda \gamma)^{t} \mathbb{E}_{s_{0: t+1}}^{a_{0: t}}\left[Q^{\prime}\left(s_{0}, a_{0}\right)\left(\prod_{i=1}^{t} \kappa_{i}\right)\left[\gamma \mathbb{E}_{\pi} Q\left(s_{t+1}, .\right)-Q\left(s_{t}, a_{t}\right)\right]^{\top}\right] \\
& \left.=\sum_{t=0}^{\infty}(\lambda \gamma)^{t} \mathbb{E}_{s_{0: t}}^{\substack{0: t}} \mid Q^{\prime}\left(s_{0}, a_{0}\right)\left(\prod_{i=1}^{t} \kappa_{i}\right)\left(\gamma \mathbb{E}_{s_{t+1}}\left[\mathbb{E}_{\pi} Q\left(s_{t+1}, .\right) \mid s_{t}, a_{t}\right]-Q\left(s_{t}, a_{t}\right)\right)\right] \\
& \left.=\sum_{t=0}^{\infty}(\lambda \gamma)^{t} \mathbb{E}_{s_{0 . t}}^{\substack{0 . t}} \mid Q^{\prime}\left(s_{0}, a_{0}\right)\left(\prod_{i=1}^{t} \kappa_{i}\right)\left(\gamma \sum_{s^{\prime} \in \mathcal{S}} \sum_{a^{\prime} \in \mathcal{A}} p\left(s^{\prime} \mid s_{t}, a_{t}\right) \pi\left(a^{\prime} \mid s^{\prime}\right) Q\left(s^{\prime}, a^{\prime}\right)-Q\left(s_{t}, a_{t}\right)\right)\right] \\
& =\sum_{t=0}^{\infty}(\lambda \gamma)^{t} \mathbb{E}_{\substack{s_{0: t} \\
a_{0: t}}}\left[Q^{\prime}\left(s_{0}, a_{0}\right)\left(\prod_{i=1}^{t} \kappa_{i}\right)\left(\gamma P^{\pi} Q\left(s_{t}, a_{t}\right)-Q\left(s_{t}, a_{t}\right)\right)\right] \\
& =\sum_{t=0}^{\infty}(\lambda \gamma)^{t} \mathbb{E}_{s_{0: t-1}}{ }_{\substack{0: t-1}}\left[Q^{\prime}\left(s_{0}, a_{0}\right)\left(\prod_{i=1}^{t-1} \kappa_{i}\right)\right. \\
& \left.\times \sum_{s^{\prime} \in \mathcal{S}} \sum_{a^{\prime} \in \mathcal{A}} p\left(s^{\prime} \mid s_{t-1}, a_{t-1}\right) \kappa\left(a^{\prime}, s^{\prime}\right) \mu\left(a^{\prime} \mid s^{\prime}\right)\left(\gamma P^{\pi} Q\left(s^{\prime}, a^{\prime}\right)-Q\left(s^{\prime}, a^{\prime}\right)\right)\right] \\
& =\sum_{t=0}^{\infty}(\lambda \gamma)^{t} \mathbb{E}_{\substack{\mathbb{E}_{0 . t-1} \\
a_{0: t-1}}}\left[Q^{\prime}\left(s_{0}, a_{0}\right)\left(\prod_{i=1}^{t-1} \kappa_{i}\right) P^{\kappa \mu}\left(\gamma P^{\pi}-I\right) Q\left(s_{t-1}, a_{t-1}\right)\right] \\
& =\mathbb{E}_{s_{0}, a_{0}}\left[Q^{\prime}\left(s_{0}, a_{0}\right) \sum_{t=0}^{\infty}(\lambda \gamma)^{t}\left(P^{\kappa \mu}\right)^{t}\left(\gamma P^{\pi}-I\right) Q\left(x_{0}, a_{0}\right)\right]
\end{aligned}
$$

$$
\begin{aligned}
& =\mathbb{E}_{s_{0}, a_{0}}\left[Q^{\prime}\left(s_{0}, a_{0}\right)\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1}\left(\gamma P^{\pi}-I\right) Q\left(s_{0}, a_{0}\right)\right] \\
& =\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \xi(s, a) Q^{\prime}(s, a)\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1}\left(\gamma P^{\pi}-I\right) Q(s, a) \\
& =Q^{\prime \top} \Xi\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1}\left(\gamma P^{\pi}-I\right) Q
\end{aligned}
$$

So, $\theta^{\prime \top} \mathbb{E}\left[A_{k}\right] \theta=\theta^{\prime \top} \Phi^{\top} \Xi\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1}\left(\gamma P^{\pi}-I\right) \Phi \theta \quad \forall \theta, \theta^{\prime} \in \mathbb{R}^{d}$, which implies that:

$$
\mathbb{E}\left[A_{k}\right]=\Phi^{\top} \Xi\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1}\left(P^{\pi}-I\right) \Phi
$$

$\theta^{\top} \mathbb{E}\left[b_{k}\right]$
$=\mathbb{E}\left[\sum_{t=0}^{\infty}(\lambda \gamma)^{t}\left(\prod_{i=1}^{t} \kappa_{i}\right) r_{t} Q\left(s_{0}, a_{0}\right)\right]=\sum_{t=0}^{\infty}(\lambda \gamma)^{t} \mathbb{E}_{\substack{\mathbb{E}_{0: t} \\ a_{0}: t}}\left[Q\left(s_{0}, a_{0}\right)\left(\prod_{i=1}^{t} \kappa_{i}\right) r\left(s_{t}, a_{t}\right)\right]$
$=\sum_{t=0}^{\infty}(\lambda \gamma)^{t} \mathbb{E}_{\substack{s_{0: t-1} \\ a_{0: t}-1}}\left[Q\left(s_{0}, a_{0}\right)\left(\prod_{i=1}^{t-1} \kappa_{i}\right) \sum_{s^{\prime} \in \mathcal{S}} \sum_{a^{\prime} \in \mathcal{A}} p\left(s^{\prime} \mid s_{t-1}, a_{t-1}\right) \kappa\left(a^{\prime}, s^{\prime}\right) \mu\left(a^{\prime} \mid s^{\prime}\right) r\left(s^{\prime}, a^{\prime}\right)\right]$
$=\sum_{t=0}^{\infty}(\lambda \gamma)^{t} \mathbb{E}_{\substack{s_{0: t-1} \\ a_{0: t-1}}}\left[Q\left(s_{0}, a_{0}\right)\left(\prod_{i=1}^{t-1} \kappa_{i}\right) P^{\kappa \mu} r\left(s^{\prime}, a^{\prime}\right)\right]$
$=\mathbb{E}_{s_{0}, a_{0}}\left[Q\left(s_{0}, a_{0}\right)\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1} r\left(s_{0}, s_{0}\right)\right]$
$=\sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \xi(s, a) Q(s, a)\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1} r(s, a)=Q^{\top} \Xi\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1} r$
So, $\theta^{\top} \mathbb{E}\left[b_{k}\right]=\theta^{\top} \Phi^{\top} \Xi\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1} r \quad \forall \theta \in \mathbb{R}^{d}$, which implies that:

$$
\mathbb{E}\left[b_{k}\right]=\Phi^{\top} \Xi\left(I-\lambda \gamma P^{\kappa \mu}\right)^{-1} r
$$

## A. 2 Proof of Proposition 2

$\operatorname{MSPBE}(\theta)$

$$
\begin{aligned}
& =\frac{1}{2}\left\|\Pi^{\mu} \mathcal{R}(\Phi \theta)-\Phi \theta\right\|_{\Xi}^{2} \\
& =\frac{1}{2}\left\|\Pi^{\mu}(\mathcal{R}(\Phi \theta)-\Phi \theta)\right\|_{\Xi}^{2} \\
& =\frac{1}{2}\left(\Pi^{\mu}(\mathcal{R}(\Phi \theta)-\Phi \theta)\right)^{\top} \Xi\left(\Pi^{\mu}(\mathcal{R}(\Phi \theta)-\Phi \theta)\right) \\
& =\frac{1}{2}\left(\Phi^{\top} \Xi(\mathcal{R}(\Phi \theta)-\Phi \theta)\right)^{\top}\left(\Phi^{\top} \Xi \Phi\right)^{-1} \Phi^{\top} \Xi\left(\Phi\left(\Phi^{\top} \Xi \Phi\right)^{-1} \Phi^{\top} \Xi(\mathcal{R}(\Phi \theta)-\Phi \theta)\right) \\
& =\frac{1}{2}\left\|\Phi^{\top} \Xi(\mathbb{R}(\Phi \theta)-\Phi \theta)\right\|_{M^{-1}}^{2} \\
& =\frac{1}{2}\left\|\Phi^{\top} \Xi\left(\left(I-\lambda \gamma P^{\mu \pi}\right)^{-1}\left(\mathcal{T}^{\pi}-\lambda \gamma P^{\mu \pi}\right) \Phi \theta-\Phi \theta\right)\right\|_{M^{-1}}^{2} \\
& =\frac{1}{2}\left\|\Phi^{\top} \Xi\left(I-\lambda \gamma P^{\mu \pi}\right)^{-1}\left(\gamma P^{\pi}-I\right) \Phi \theta+\Phi^{\top} \Xi\left(I-\lambda \gamma P^{\mu \pi}\right)^{-1} r\right\|_{M^{-1}}^{2} \\
& =\frac{1}{2}\|A \theta+b\|_{M^{-1}}^{2}
\end{aligned}
$$

## A. 3 Proof of Proposition 3

Let's show that $\mathbb{E}\left[\hat{A}_{k}\right]=A$. Let's $\Delta_{t}$ denotes $\left[\gamma \mathbb{E}_{\pi} \phi\left(s_{t+1}, .\right)^{\top}-\phi\left(s_{t}, a_{t}\right)^{\top}\right]$

$$
\begin{aligned}
A= & \mathbb{E}\left[\sum_{t=k}^{\infty}(\lambda \gamma)^{t-k}\left(\prod_{i=k+1}^{t} \kappa_{i}\right) \phi\left(s_{k}, a_{k}\right) \Delta_{t}\right] \\
= & \mathbb{E}\left[\phi\left(s_{k}, a_{k}\right) \Delta_{k}+\sum_{t=k+1}^{\infty}(\lambda \gamma)^{t-k}\left(\prod_{i=k+1}^{t} \kappa_{i}\right) \phi\left(s_{k}, a_{k}\right) \Delta_{t}\right] \\
= & \mathbb{E}\left[\phi\left(s_{k}, a_{k}\right) \Delta_{k}+\sum_{t=k}^{\infty}(\lambda \gamma)^{t-k+1}\left(\prod_{i=k+1}^{t+1} \kappa_{i}\right) \phi\left(s_{k}, a_{k}\right) \Delta_{t+1}\right] \\
= & \mathbb{E}\left[\phi\left(s_{k}, a_{k}\right) \Delta_{k}+\lambda \gamma \kappa\left(s_{k+1}, a_{k+1}\right) \phi\left(s_{k}, a_{k}\right) \Delta_{k+1}\right. \\
& \left.+\sum_{t=k+1}^{\infty}(\lambda \gamma)^{t-k+1}\left(\prod_{i=k+1}^{t+1} \kappa_{i}\right) \phi\left(s_{k}, a_{k}\right) \Delta_{t+1}\right]
\end{aligned}
$$

$$
\begin{aligned}
\stackrel{(\star)}{=} & \mathbb{E}\left[\phi\left(s_{k}, a_{k}\right) \Delta_{k}+\lambda \gamma \kappa\left(s_{k}, a_{k}\right) \phi\left(s_{k-1}, a_{k-1}\right) \Delta_{k}\right. \\
& \left.+\sum_{t=k+1}^{\infty}(\lambda \gamma)^{t-k+1}\left(\prod_{i=k+1}^{t+1} \kappa_{i}\right) \phi\left(s_{k}, a_{k}\right) \Delta_{t+1}\right] \\
= & \mathbb{E}\left[\Delta _ { k } \left(\phi\left(s_{k}, a_{k}\right)+\lambda \gamma \kappa\left(s_{k}, a_{k}\right) \phi\left(s_{k-1}, a_{k-1}\right)\right.\right. \\
& \left.\left.+(\lambda \gamma)^{2} \kappa\left(s_{k}, a_{k}\right) \kappa\left(s_{k-1}, a_{k-1}\right) \phi\left(s_{k-2}, a_{k-2}\right)+\ldots\right)\right] \\
= & \mathbb{E}\left[\Delta_{k}\left(\sum_{i=0}^{k}(\lambda \gamma)^{k-i}\left(\prod_{j=i+1}^{k} \kappa_{j}\right) \phi\left(x_{i}, a_{i}\right)\right)\right] \\
= & \mathbb{E}\left[\Delta_{k} e_{k}\right]=\mathbb{E}\left[\hat{A}_{k}\right]
\end{aligned}
$$

we have used in the line $(\star)$ the fact that $\mathbb{E}\left[\kappa\left(s_{k+1}, a_{k+1}\right) \phi\left(s_{k} a_{k}\right) \Delta_{k+1}\right]=$ $\mathbb{E}\left[\kappa\left(s_{k}, a_{k}\right) \phi\left(s_{k-1} a_{k-1}\right) \Delta_{k}\right]$ thanks to the stationarity of the distribution $d$. we have also denote by $e_{k}$ the following vector:

$$
\begin{aligned}
e_{k} & =\sum_{i=0}^{k}(\lambda \gamma)^{k-i}\left(\prod_{j=i+1}^{k} \kappa_{j}\right) \phi\left(s_{i}, a_{i}\right) \\
& =\lambda \gamma \kappa_{k}\left(\sum_{i=0}^{k-1}(\lambda \gamma)^{k-1-i}\left(\prod_{j=i+1}^{k-1} \kappa_{j}\right) \phi\left(s_{i}, a_{i}\right)\right)+\phi\left(s_{k}, a_{k}\right) \\
& =\lambda \gamma \kappa_{k} e_{k-1}+\phi\left(s_{k}, a_{k}\right)
\end{aligned}
$$

Vector $e_{k}$ corresponds to the eligibility traces defined in the proposition. Similarly, we could show that $\mathbb{E}_{\mu}\left[\hat{b}_{k}\right]=b$.

## A. 4 Convergence Rate Analysis

Let's recall the key quantities defined in the main article:

$$
\rho \triangleq \lambda_{\max }\left(A^{\top} M^{-1} A\right), \quad \delta \triangleq \lambda_{\min }\left(A^{\top} M^{-1} A\right), \quad L_{G} \triangleq\left\|\mathbb{E}\left[\hat{G}_{k}^{\top} \hat{G}_{k} \mid \mathcal{F}_{k}\right]\right\|
$$

We will make use of spectral properties of the matrix $G$ provided in the appendix A of [Du et al., 2017]. it was shown that if we set $\beta=\frac{8 \rho}{\lambda_{\min }(M)}$, the matrix $G$ is diagonalizable with all its eigenvalues real and positive. It is a straightforward application of result from [Benzi and Simoncini, 2006] Moreover, it was proved that $G$ can be written as: $G=Q \Lambda Q^{-1}$ where $\Lambda$ is a
diagonal matrix whose diagonal entries are the eigenvalues of $G$ and $Q$ consists of it eigenvectors as columns such that the condition number of $Q$ is upper bounded by the one of $M$ as follows:

$$
c(Q)^{2} \leq 8 c(M)
$$

Finally, the paper showed upper and lower bounds for the eigenvalues of G:

$$
\begin{aligned}
\lambda_{\max }(G) & \leq 9 c(M) \rho \\
\lambda_{\min }(G) & \geq \frac{8}{9} \delta
\end{aligned}
$$

Let's recall our updates:

$$
z_{k+1}=z_{k}-\alpha_{k}\left(\hat{G}_{k} z_{k}-\hat{g}_{k}\right)
$$

By subtracting $z^{\star}$ from both sides on the later equation and using the optimality condition $G z^{\star}+g=0$ :

$$
\begin{equation*}
\Delta_{k+1}=\Delta_{k}-\alpha_{k} G \Delta_{k}+\alpha_{k}\left[G z_{k}-g-\left(\hat{G}_{k} z_{k}-\hat{g}_{k}\right)\right] \quad \text { where } \quad \Delta_{k} \triangleq z_{k}-z^{\star} \tag{A.4.1}
\end{equation*}
$$

By multiplying both sides by $Q^{-1}$ and using the fact that $Q^{-1} G=\Lambda Q^{-1}$ :

$$
\begin{aligned}
& Q^{-1} \Delta_{k+1}=Q^{-1} \Delta_{k}-\alpha_{k} Q^{-1} G \Delta_{k}+\alpha_{k} Q^{-1}\left[G z_{k}-g-\left(\hat{G}_{k} z_{k}-\hat{g}_{k}\right)\right] \\
& =\left(I-\alpha_{k} \Lambda\right) Q^{-1} \Delta_{k}+\alpha_{k} Q^{-1}\left[G z_{k}-g-\left(\hat{G}_{k} z_{k}-\hat{g}_{k}\right)\right] \\
& \mathbb{E}\left[\left\|Q^{-1} \Delta_{k+1}\right\|^{2} \mid \mathcal{F}_{k-1}\right] \\
& =\mathbb{E}\left[\left\|\left(I-\alpha_{k}\right) Q^{-1} \Delta_{k}+\alpha_{k} Q^{-1}\left[G z_{k}-g-\left(\hat{G}_{k} z_{k}-\hat{g}_{k}\right)\right]\right\|^{2} \mid \mathcal{F}_{k-1}\right] \\
& =\mathbb{E}\left[\left\|\left(I-\alpha_{k} \Lambda\right) Q^{-1} \Delta_{k}\right\|^{2} \mid \mathcal{F}_{k-1}\right] \\
& \quad+2 \mathbb{E}\left[\left\langle\left(I-\alpha_{k}\right) Q^{-1} \Delta_{k}, \alpha_{k} Q^{-1}\left[G z_{k}-g-\left(\hat{G}_{k} z_{k}-\hat{g}_{k}\right)\right]\right\rangle \mid \mathcal{F}_{k-1}\right] \\
& =\| \alpha_{k}^{2} \mathbb{E}\left[\left\|Q^{-1}\left[G z_{k}-g-\left(\hat{G}_{k} z_{k}-\hat{g}_{k}\right)\right]\right\|^{2} \mid \mathcal{F}_{k-1}\right] \\
& \left.\leq \|-\alpha_{k} \Lambda\right) Q^{-1} \Delta_{k} \|^{2}+\alpha_{k}^{2} \mathbb{E}\left[\left\|Q^{-1}\left[G z_{k}-g-\left(\hat{G}_{k} z_{k}-\hat{g}_{k}\right)\right]\right\|^{2} \mid \mathcal{F}_{k-1}\right] \\
& =\left\|I-\alpha_{k} \Lambda\right\|^{2}\left\|Q^{-1} \Delta_{k}\right\|^{2}+\alpha_{k}^{2} \mathbb{E}\left[\left\|Q^{-1}\left(\hat{G}_{k} z_{k}-\hat{g}_{k}\right)\right\|^{2} \mid \mathcal{F}_{k-1}\right] \\
& \leq\left\|I-\alpha_{k} \Lambda\right\|^{2}\left\|Q^{-1} \Delta_{k}\right\|^{2}+\alpha_{k}^{2} \mathbb{E}\left[\left\|Q^{-1}\left(\hat{G}_{k} \Delta_{k}+\hat{G}_{k} z^{\star}-\hat{g}_{k}\right)\right\|^{2} \mid \mathcal{F}_{k-1}\right] \\
& \leq\left\|Q^{-1} \Delta_{k}\right\|^{2}+2 \alpha_{k}^{2} \mathbb{E}\left[\left\|Q^{-1} \hat{G}_{k} \Delta_{k}\right\|^{2} \mid \mathcal{F}_{k-1}\right]+2 \alpha_{k}^{2} \mathbb{E}\left[\left\|Q^{-1}\left(\hat{G}_{k} z^{\star}+\hat{g}_{k}\right)\right\|^{2} \mid \mathcal{F}_{k-1}\right]
\end{aligned}
$$

we use in the third line the fact that $\mathbb{E}\left[\hat{G}_{k} \mid \mathcal{F}_{k-1}\right]=G$ and $\mathbb{E}\left[\hat{g}_{k-1} \mid \mathcal{F}_{k-1}\right]=g$.

$$
\begin{aligned}
& \left\|I-\alpha_{k} \Lambda\right\|^{2}=\max \left\{\left|1-\alpha_{k} \lambda_{\min }(G)\right|^{2},\left|1-\alpha_{k} \lambda_{\max }(G)\right|^{2}\right\} \\
& \leq 1-2 \alpha_{k} \lambda_{\min }+\alpha_{k}^{2} \lambda_{\max }^{2} \\
& \leq 1-2 \alpha_{k} \frac{8}{9} \delta+\alpha_{k}^{2} 9^{2} c(M)^{2} \rho^{2} \\
& \leq 1-2 \alpha_{k} \delta^{\prime}
\end{aligned}+\alpha_{k}^{2} 9^{2} c(M)^{2} \rho^{2} \quad \text { where } \delta^{\prime} \triangleq \frac{8}{9} \delta ~ 子 \begin{aligned}
\mathbb{E}\left[\left\|Q^{-1} \hat{G}_{k} \Delta_{k}\right\|^{2} \mid \mathcal{F}_{k-1}\right] & \leq\left\|Q^{-1}\right\|^{2} \mathbb{E}\left[\left\|\hat{G}_{k} \Delta_{k}\right\|^{2} \mid \mathcal{F}_{k-1}\right] \\
& =\left\|Q^{-1}\right\|^{2} \mathbb{E}\left[\Delta_{k}^{\top} \hat{G}_{k}^{\top} \hat{G}_{k} \Delta_{k} \mid \mathcal{F}_{k-1}\right] \\
& =\left\|Q^{-1}\right\|^{2} \Delta_{k}^{\top} \mathbb{E}\left[\hat{G}_{k}^{\top} \hat{G}_{k} \mid \mathcal{F}_{k-1}\right] \Delta_{k} \\
& \leq\left\|Q^{-1}\right\|^{2}\left\|\mathbb{E}\left[\hat{G}_{k}^{\top} \hat{G}_{k} \mid \mathcal{F}_{k}\right]\right\|^{2} \Delta_{k}^{\top} \Delta_{k} \\
& \leq\left\|Q^{-1}\right\|^{2} L_{G}\left\|\Delta_{k}\right\|^{2} \\
& =\left\|Q^{-1}\right\|^{2} L_{G}\left\|Q Q^{-1} \Delta_{k}\right\|^{2} \\
& \leq\left\|Q^{-1}\right\|^{2}\|Q\|^{2} L_{G}\left\|Q^{-1} \Delta_{k}\right\|^{2} \\
& \leq c(Q)^{2} L_{G}\left\|Q^{-1} \Delta_{k}\right\|^{2}
\end{aligned}
$$

So, we have:

$$
\begin{aligned}
& \mathbb{E}\left[\left\|Q^{-1} \Delta_{k+1}\right\|^{2}\right] \\
& \left.\leq\left(1-2 \alpha_{k} \delta^{\prime}+\alpha_{k}^{2} 9^{2} c(M)^{2} \rho^{2}+16 \alpha_{k}^{2} c(M) L_{G}\right) \mathbb{E}\left[\left\|Q^{-1} \Delta_{k}\right\|^{2}\right]+2 \alpha_{k}^{2}\left\|Q^{-1}\right\|^{2} \mathbb{E}\left[\| \hat{G}_{k} z^{\star}-\hat{g}_{k}\right) \|^{2}\right]
\end{aligned}
$$

By selecting $\alpha_{k}=\frac{2 \delta^{\prime}}{\delta^{\prime 2}(k+2)+2 \times 9^{2} c(M)^{2} \rho^{2}+32 c(M) L_{G}}=\frac{2 \delta^{\prime}}{\delta^{\prime 2}(k+2)+\zeta}$ with $\zeta=2 \times$

$$
\begin{aligned}
& 9^{2} c(M)^{2} \rho^{2}+32 c(M) L_{G}, \text { we get: } \\
& \mathbb{E}\left[\left\|Q^{-1} \Delta_{k+1}\right\|^{2}\right] \\
& \leq\left(1-\delta^{\prime} \alpha_{k}\right) \mathbb{E}\left[\left\|Q^{-1} \Delta_{k}\right\|^{2}\right]+2 \alpha_{k}^{2}\left\|Q^{-1}\right\|^{2} \mathbb{E}\left[\left\|\hat{G}_{k} z^{\star}-\hat{g}_{k}\right\|^{2}\right] \\
&=\left.\frac{\delta^{\prime 2} k+\zeta}{\delta^{\prime 2}(k+2)+\zeta} \mathbb{E}\left[\left\|Q^{-1} \Delta_{k}\right\|^{2}\right]+\frac{8 \delta^{\prime 2}}{\left(\delta^{\prime 2}(k+2)+\zeta\right)^{2}}\left\|Q^{-1}\right\|^{2} \mathbb{E}\left[\| \hat{G}_{k} z^{\star}+\hat{g}_{k}\right) \|^{2}\right] \\
& \leq\left(\prod_{i=0}^{k} \frac{\delta^{\prime 2} i+\zeta}{\delta^{\prime 2}(i+2)+\zeta}\right) \mathbb{E}\left[\left\|Q^{-1} \Delta_{0}\right\|^{2}\right] \\
&\left.+8 \delta^{\prime 2} \sum_{i=0}^{k}\left(\prod_{j=i}^{k} \frac{\delta^{\prime 2} j+\zeta}{\delta^{\prime 2}(j+2)+\zeta}\right) \frac{1}{\left(\delta^{\prime 2}(i+2)+\zeta\right)^{2}}\left\|Q^{-1}\right\|^{2} \mathbb{E}\left[\| \hat{G}_{i} z^{\star}+\hat{g}_{i}\right) \|^{2}\right] \\
&= \frac{\zeta\left(\delta^{\prime 2}+\zeta\right)}{\left(\delta^{\prime 2}(k+1)+\zeta\right)\left(\left(\delta^{\prime 2}(k+2)+\zeta\right)\right.} \mathbb{E}\left[\left\|Q^{-1} \Delta_{0}\right\|^{2}\right] \\
&+8 \delta^{\prime 2} \sum_{i=0}^{k} \frac{\left(\delta^{\prime 2}(i+1)+\zeta\right)\left(\delta^{\prime 2} i+\zeta\right)}{\left(\delta^{\prime 2}(k+1)+\zeta\right)\left(\left(\delta^{\prime 2}(k+2)+\zeta\right)\right.} \frac{1}{\left(\delta^{\prime 2}(i+2)+\zeta\right)^{2}}\left\|Q^{-1}\right\|^{2} \mathbb{E}\left[\left\|\hat{G}_{i} z^{\star}-\hat{g}_{i}\right\|^{2}\right] \\
& \leq \frac{\zeta\left(\delta^{\prime 2}+\zeta\right)}{\left(\delta^{\prime 2}(k+1)+\zeta\right)\left(\delta^{\prime 2}(k+2)+\zeta\right)} \mathbb{E}\left[\left\|Q^{-1} \Delta_{0}\right\|^{2}\right] \\
&+8 \delta^{\prime 2} \sum_{i=0}^{k} \frac{1}{\left(\delta^{\prime 2}(k+1)+\zeta\right)\left(\delta^{\prime 2}(k+2)+\zeta\right)}\left\|Q^{-1}\right\|^{2} \mathbb{E}\left[\left\|\hat{G}_{i} z^{\star}-\hat{g}_{i}\right\|^{2}\right] \\
& \leq \frac{\left(\delta^{\prime}+\zeta\right)^{2}}{\left(\delta^{\prime 2}(k+1)+\zeta\right)^{2}} \mathbb{E}\left[\left\|Q^{-1} \Delta_{0}\right\|^{2}\right] \\
&\left.+8 \frac{\delta^{\prime 2}(k+1)}{\left(\delta^{\prime 2}(k+1)+\zeta\right)\left(\delta^{\prime 2}(k+2)+\zeta\right)}\left\|Q^{-1}\right\|^{2} \sup _{i=0 . . . k} \mathbb{E}\left[\| \hat{G}_{i} z^{\star}+\hat{g}_{i}\right) \|^{2}\right] \\
& \leq \frac{\left(\delta^{\prime}+\zeta\right)^{2}}{\left(\delta^{\prime 2}(k+1)+\zeta\right)^{2}} \mathbb{E}\left[\left\|Q^{-1} \Delta_{0}\right\|^{2}\right]+\frac{8}{\left(\delta^{\prime 2}(k+1)+\zeta\right)}\left\|Q^{-1}\right\|^{2} \sup _{i=0 \ldots k} \mathbb{E}\left[\left\|\hat{G}_{i} z^{\star}-\hat{g}_{i}\right\|^{2}\right] \\
& \leq \frac{\left(\delta^{\prime}+\zeta\right)^{2}\left\|Q^{-1}\right\|^{2}}{\left(\delta^{\prime 2}(k+1)+\zeta\right)^{2}} \mathbb{E}\left[\left\|\Delta_{0}\right\|^{2}\right]+\frac{8 \sigma^{2}\left\|Q^{-1}\right\|^{2}}{\left(\delta^{\prime 2}(k+1)+\zeta\right)}\left(1+\left\|z^{\star}\right\|^{2}\right)
\end{aligned}
$$

Moreover, we have $\mathbb{E}\left[\left\|\Delta_{k+1}\right\|^{2}\right]=\mathbb{E}\left[\left\|Q Q^{-1} \Delta_{k+1}\right\|^{2}\right] \leq\|Q\|^{2} \mathbb{E}\left[\left\|Q^{-1} \Delta_{k+1}\right\|^{2}\right]$.

Then, we get:

$$
\begin{aligned}
\mathbb{E}\left[\left\|\Delta_{k+1}\right\|^{2}\right] & \leq \frac{\left(\delta^{\prime}+\zeta\right)^{2} c(Q)^{2}}{\left(\delta^{\prime 2}(k+1)+\zeta\right)^{2}} \mathbb{E}\left[\left\|\Delta_{0}\right\|^{2}\right]+\frac{8 \sigma^{2} c(Q)^{2}}{\left(\delta^{\prime 2}(k+1)+\zeta\right)}\left(1+\left\|z^{\star}\right\|^{2}\right) \\
& \leq \frac{8\left(\delta^{\prime}+\zeta\right)^{2} c(M)}{\left(\delta^{\prime 2}(k+1)+\zeta\right)^{2}} \mathbb{E}\left[\left\|\Delta_{0}\right\|^{2}\right]+\frac{8^{2} \sigma^{2} c(M)}{\left(\delta^{\prime 2}(k+1)+\zeta\right)}\left(1+\left\|z^{\star}\right\|^{2}\right) \\
& =\frac{89^{2}(8 \delta+9 \zeta)^{2} c(M)}{\left(8^{2} \delta^{2}(k+1)+9^{2} \zeta\right)^{2}} \mathbb{E}\left[\left\|\Delta_{0}\right\|^{2}\right]+\frac{9^{2} \times 8^{2} \sigma^{2} c(M)}{\left(8^{2} \delta^{2}(k+1)+9^{2} \zeta\right)}\left(1+\left\|z^{\star}\right\|^{2}\right) \\
& =9^{2} \times 8 c(M)\left\{\frac{(8 \delta+9 \zeta)^{2} \mathbb{E}\left[\left\|\Delta_{0}\right\|^{2}\right]}{\left(8^{2} \delta^{2}(k+1)+9^{2} \zeta\right)^{2}}+\frac{8 \sigma^{2}\left(1+\left\|z^{\star}\right\|^{2}\right)}{\left(8^{2} \delta^{2}(k+1)+9^{2} \zeta\right)}\right\}
\end{aligned}
$$

The overall convergence rate is then equal to $O(1 / k)$.

## A. 5 TRUE ON-LINE EQUIVALENCE

In van Hasselt et al. [2014], the authors derived a true on-line update for $\operatorname{GTD}(\lambda)$ that empirically performed better than $\operatorname{GTD}(\lambda)$ with eligibility traces. Based on this work, we derive true on-line updates for our algorithm. The gradient off-policy algorithm was derived by turning the expected forward view into an expected backward view which can be sampled. In order to derive a true on-line update, we sample instead the forward view and then we turn the sampled forward view to an exact backward view using Theorem 1 in van Hasselt et al. [2014]. If $k$ denotes the time horizon, we consider the sampled truncated interim forward return:

$$
\forall t<k, \quad Y_{t}^{k}=\sum_{i=t}^{k-1}(\lambda \gamma)^{i-t}\left(\prod_{j=t+1}^{i} \kappa_{j}\right) \delta_{i}
$$

where $\delta_{i}=r_{i}+\theta_{t}^{\top} \mathbb{E}_{\pi} \phi\left(s_{t+1}, \cdot\right)-\theta_{t}^{\top} \phi\left(s_{t}, a_{t}\right)$, which gives us the sampled forward update of $\omega$ :

$$
\begin{equation*}
\forall k<t, \quad \omega_{t+1}^{k}=\omega_{t}^{k}+\alpha_{t}\left(Y_{t}^{k}-\phi\left(x_{t}, a_{t}\right)^{\top} \omega_{t}^{k}\right) \phi\left(x_{t}, a_{t}\right) \tag{A.5.1}
\end{equation*}
$$

Proposition 9. For any $k$, the parameter $\omega_{k}^{k}$ defined by the forward view (A.5.1) is equal to $\omega_{k}$ defined by the following backward view:

$$
\begin{aligned}
e_{-1}^{\omega} & =0, \quad \forall k \geq 0 \\
e_{k}^{\omega} & =\lambda \gamma \kappa_{k} e_{k-1}^{\omega}+\alpha_{k}\left(1-\lambda \gamma \kappa_{k} \phi\left(s_{k}, a_{k}\right)^{\top} e_{k-1}^{\omega}\right) \phi\left(s_{k}, a_{k}\right) \\
\omega_{k+1} & =\omega_{k}+\delta_{k} e_{k}^{\omega}-\alpha_{t} \phi\left(s_{k}, a_{k}\right)^{\top} \omega_{k} \phi\left(s_{k}, a_{k}\right)
\end{aligned}
$$

Proof. The return's temporal difference $Y_{t}^{k+1}-Y_{t}^{k}$ are related through:

$$
\begin{aligned}
\forall t<k, \quad Y_{t}^{k+1}-Y_{t}^{k} & =\sum_{i=t}^{k}(\lambda \gamma)^{i-t}\left(\prod_{j=t+1}^{i} \kappa_{j}\right) \delta_{i}-\sum_{i=t}^{k-1}(\lambda \gamma)^{i-t}\left(\prod_{j=t+1}^{i} \kappa_{j}\right) \delta_{i} \\
& =(\lambda \gamma)^{k-t}\left(\prod_{j=t+1}^{k} \kappa_{j}\right) \delta_{k} \\
& =\lambda \gamma \kappa_{k+1}\left((\lambda \gamma)^{k-(t+1)}\left(\prod_{j=t+2}^{k} \kappa_{j}\right) \delta_{k}\right) \\
& =\lambda \gamma \kappa_{k+1}\left(Y_{t+1}^{k+1}-Y_{t+1}^{k}\right)
\end{aligned}
$$

We could then apply Theorem 1 of van Hasselt et al. [2014] that give us the following backward view:

$$
\begin{aligned}
e_{0} & =\alpha_{0} \phi\left(x_{0}, a_{0}\right) \\
e_{t} & =\lambda \gamma \kappa_{t} e_{t-1}+\alpha_{t}\left(1-\lambda \gamma \kappa_{k} \phi\left(s_{t}, a_{t}\right)^{\top} e_{t-1}\right) \phi\left(s_{t}, a_{t}\right) \quad \forall t>0 \\
\omega_{t+1} & =\omega_{t}+\left(Y_{t}^{t+1}-Y_{t}^{t}\right) e_{t}+\alpha_{t}\left(Y_{t}^{t}-\phi\left(s_{t}, a_{t}\right)^{\top} \omega_{t}\right) \phi\left(s_{t}, a_{t}\right) \\
& \stackrel{(\star)}{=} \omega_{t}+\delta_{t} e_{t}-\alpha_{t} \phi\left(s_{t}, a_{t}\right)^{\top} \omega_{t} \phi\left(s_{t}, a_{t}\right)
\end{aligned}
$$

We used in the line ( $\star$ ) that $Y_{t}^{t+1}=\delta_{t}$ and $Y_{t}^{t}=0$
The resulting detailed procedure is provided in Algorithm 8.
Note that when $\lambda$ is equal to zero, the Algorithm 1 and 2 both reduce to the same update:

$$
\begin{aligned}
\omega_{k+1} & =\omega_{k}+\alpha_{k}\left(\delta_{k}-\phi\left(s_{k}, a_{k}\right)^{\top} \omega_{k}\right) \phi\left(s_{k}, a_{k}\right) \\
\theta_{k+1} & \left.=\theta_{k}-\alpha_{k} \phi\left(s_{k}, a_{k}\right)^{\top} w_{k}\left(\gamma \mathbb{E}_{\pi}\left[\phi\left(s_{k+1}, \cdot\right)\right]-\phi\left(s_{k}, a_{k}\right)\right]\right)
\end{aligned}
$$

```
Algorithm 8 Gradient Off-policy with eligibility/Dutch traces
Given: target policy \(\pi\), behavior policy \(\mu\)
    Initialize \(\theta_{0}\) and \(\omega_{0}\)
    for \(\mathrm{n}=0 \ldots\) do
        set \(e_{-1}^{\theta}=e_{-1}^{\omega}=0\)
        for \(\mathrm{k}=0 \ldots\) end of episode do
            Observe \(s_{k}, a_{k}, r_{k}, s_{k+1}\) according to \(\mu\)
            /* Update traces
            \(e_{k}=\lambda \gamma \kappa\left(s_{k}, a_{k}\right) e_{k-1}+\phi\left(s_{k}, a_{k}\right)\)
            \(/^{*}\) Update Dutch traces
            \(e_{k}^{\omega}=\lambda \gamma \kappa_{k} e_{k-1}^{\omega}+\alpha_{k}\left(1-\lambda \gamma \kappa_{k} \phi\left(s_{k}, a_{k}\right)^{\top} e_{k-1}^{\omega}\right) \phi\left(s_{k}, a_{k}\right)\)
            /* Update parameters
            \(\delta_{k}=r_{k}+\gamma \theta_{k}^{\top} \mathbb{E}_{\pi} \phi\left(s_{k+1},.\right)-\theta_{k}^{\top} \phi\left(s_{k}, a_{k}\right)\)
            \(\omega_{k+1}=\omega_{k}+\delta_{k} e_{k}^{\omega}-\alpha_{k} \phi\left(s_{k}, a_{k}\right)^{\top} \omega_{k} \phi\left(s_{k}, a_{k}\right)\)
            \(\theta_{k+1}=\theta_{k}-\alpha_{k} \omega_{k}^{\top} e_{k}\left(\gamma \mathbb{E}_{\pi}\left[\phi\left(s_{k+1},.\right)\right]-\phi\left(s_{k}, a_{k}\right)\right)\)
        end for
    end for
```


## Appendix For Chapter 4

## B. 1 PROOF OF LEMMA 2

According to performance difference lemma 1, we have

$$
\begin{aligned}
J\left(\pi^{\prime}\right) & =J(\pi)+\mathbb{E}_{s \sim d_{\rho}^{\pi^{\prime}}} \mathbb{E}_{a \sim \pi^{\prime}(\cdot \mid s)}\left[A^{\pi}(s, a)\right] \\
& =J(\pi)+\left(\mathbb{E}_{s \sim d_{\rho}^{\pi}} \mathbb{E}_{a \sim \pi^{\prime}(\cdot \mid s)}\left[A^{\pi}(s, a)\right]+\int_{s \in \mathcal{S}} \mathbb{E}_{a \sim \pi^{\prime}(\cdot \mid s)}\left[A^{\pi}(s, a)\right]\left(d_{\rho}^{\pi^{\prime}}(s)-d_{\rho}^{\pi}(s)\right) d s\right) \\
& \geq J(\pi)+\left(\mathbb{E}_{s \sim d_{\rho}^{\pi}} \mathbb{E}_{a \sim \pi^{\prime}(\cdot \mid s)}\left[A^{\pi}(s, a)\right]-\int_{s \in \mathcal{S}}\left|\mathbb{E}_{a \sim \pi(\cdot \mid s)}\left[A^{\pi}(s, a)\right]\right| \cdot\left|d_{\rho}^{\pi^{\prime}}(s)-d_{\rho}^{\pi}(s)\right| d s\right) \\
& \geq J(\pi)+\left(\mathbb{E}_{s \sim d_{\rho}^{\pi}} \mathbb{E}_{a \sim \pi^{\prime}(\cdot \mid s)}\left[A^{\pi}(s, a)\right]-\epsilon^{\pi} \int_{s \in \mathcal{S}}\left|d_{\rho}^{\pi^{\prime}}(s)-d_{\rho}^{\pi}(s)\right| d s\right) \\
& \geq J(\pi)+\left(\mathbb{E}_{s \sim d_{\rho}^{\pi}} \mathbb{E}_{a \sim \pi^{\prime}(\cdot \mid s)}\left[A^{\pi}(s, a)\right]-\epsilon^{\pi} D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi}\right)\right) \\
& =L_{\pi}\left(\pi^{\prime}\right)-\epsilon^{\pi} D_{\mathrm{TV}}\left(d_{\rho}^{\pi^{\prime}} \| d_{\rho}^{\pi}\right)
\end{aligned}
$$

where $\epsilon^{\pi}=\max _{s}\left|\mathbb{E}_{a \sim \pi^{\prime}(\cdot \mid s)}\left[A^{\pi}(s, a)\right]\right|$ and $D_{\mathrm{TV}}$ is total variation distance. The first inequality follows from Cauchy-Schwartz inequality.

## B. 2 SCORE FUNCTION ESTIMATOR OF THE GRADIENT WITH RESPECT TO THE POLICY

$$
\begin{aligned}
\nabla_{\pi^{\prime}} \mathbb{E}_{\substack{s \sim \rho \\
a \sim \pi^{\prime}}}[g(s, a)] & =\nabla_{\pi^{\prime}} \int g(s, a) \rho(s) \pi^{\prime}(a \mid s) \\
& =\int g(s, a) \rho(s) \nabla_{\pi^{\prime}} \pi^{\prime}(a \mid s) \\
& =\underset{\substack{s \sim \rho \\
a \sim \pi^{\prime}}}{ }\left[g(s, a) \nabla_{\pi^{\prime}} \log \pi^{\prime}(a \mid s)\right]
\end{aligned}
$$

$$
\begin{aligned}
& \nabla_{\pi^{\prime}} \mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi_{i}}}\left[\phi^{\star}\left(\left(g-\gamma P^{\pi^{\prime}} g\right)(s, a)\right)\right] \\
& =\mathbb{E}_{(s, a) \sim \mu_{\rho}}^{\pi_{i}}\left[\nabla_{\pi^{\prime}} \phi^{\star}\left(\left(g-\gamma P^{\pi^{\prime}} g\right)(s, a)\right)\right] \\
& =\mathbb{E}_{(s, a) \sim \mu_{\rho}^{\pi_{i}}[ }\left[\frac{\partial \phi^{\star}}{\partial t}\left(\left(g-\gamma P^{\pi^{\prime}} g\right)(s, a)\right) \nabla_{\pi^{\prime}}\left(g-\gamma \mathbb{P}^{\pi^{\prime}} g\right)\right] \\
& \left.=-\gamma \mathbb{E}_{(s, a) \sim \mu_{\rho}}^{\pi_{i}}\left[\frac{\partial \phi^{\star}}{\partial t}\left(\left(g-\gamma P^{\pi^{\prime}} g\right)(s, a)\right) \nabla_{\pi^{\prime}} \int g\left(s^{\prime}, a^{\prime}\right) P\left(s^{\prime} \mid s, a\right) \pi^{\prime}\left(a^{\prime} \mid s^{\prime}\right)\right)\right] \\
& =-\gamma \mathbb{E}_{(s, a) \sim \mu_{\rho}}^{\pi_{i}}\left[\frac{\partial \phi^{\star}}{\partial t}\left(\left(g-\gamma P^{\pi^{\prime}} g\right)(s, a)\right) \mathbb{E}_{\substack{s^{\prime} \sim P(\cdot \mid s, a) \\
a^{\prime} \sim \pi^{\prime}\left(\cdot \mid s^{\prime}\right)}}\left[g\left(s^{\prime}, a^{\prime}\right) \nabla_{\pi^{\prime}} \log \pi^{\prime}\left(a^{\prime} \mid s^{\prime}\right)\right]\right]
\end{aligned}
$$

## B. 3 Comparison with Algaedice

Both the recent AlgaeDICE [Nachum et al., 2019b] and our present work propose regularisation based on discounted state-action visitation distribution but in different ways. Firstly, AlgaeDICE is initially designed to find an optimal policy given a batch of training data. They alter the objective function itself i.e the policy performance $J(\pi)$ by adding the divergence between the discounted state-action visitation distribution and training distribution, while our approach adds the divergence term to $L_{\pi}\left(\pi^{\prime}\right)$. The latter is a first order Taylor approximation of the policy performance $J\left(\pi^{\prime}\right)$. Therefore, our approach could be seen as a mirror descent that uses the divergence as a proximity term. Secondly, their training objective is completely different from ours. Their method ends up being an off-policy version of the actor-critic method.

We implemented the AlgaeDICE min-max objective to replace our surrogate min-max objective in the PPO training procedure i.e at each iteration, we sample rollouts from the current policy and update the actor and the critic of AlgaeDICE for 10 epochs. Empirically, we observed that AlgaeDICE objective is very slow to train in this setting. This was expected as it is agnostic to training data while our method leverages the fact that the data is produced by the current policy and estimates advantage using on-policy multi-step Monte Carlo. So our approach is more suitable than AlgaeDICE in this setting. However, AlgaeDICE, as an off-policy method, would be better when storing all history of transitions and updating both actor and critic after each transition, as shown in Nachum et al. [2019b].

## B. 4 ADDITIONAL EMPIRICAL RESUlTS ON MuJoCo

See Figure B. 1


Figure B.1: Our method with KL divergences in comparison to PPO and TRPO on MuJoCo, with 10 seeds. Standard error shaded.

## B. 5 HYPERPARAMETERS

For the OpenAI Gym environments we use the default hyperparameters in Kostrikov [2018].

| Parameter name | MuJoCo | Atari |
| :---: | :---: | :---: |
| Number of minibatches | 4 | 4 |
| Discount $\gamma$ | 0.99 | 0.99 |
| Optimizer | Adam | Adam |
| Learning rate | $3 \mathrm{e}-4$ | $2.5 \mathrm{e}-4$ |
| PPO clip parameter | 0.2 | 0.1 |
| PPO epochs | 10 | 4 |
| Number of processes | 1 | 8 |
| GAE $\lambda$ | 0.95 | 0.95 |
| Entropy coef | 0 | 0.01 |
| Value loss coef | 0.5 | 0.5 |
| Number of forward steps per update | 2048 | 128 |

Table B.1: A complete overview of used hyper parameters for all methods.

## Appendix For Chapter 5

## C. 1 OMITTED PROOFS FOR THE LIPSCHITZ SETTING

## C.1.1 Proof of Lemma 3

(a) It is obvious that $\cup_{B \in \mathcal{B}_{h}^{k}} \operatorname{dom}_{h}^{k}(B) \subset \cup_{B \in \mathcal{B}_{h}^{k}} B$. Let $x \in \cup_{B \in \mathcal{B}_{h}^{k}} B$. Consider a smallest radius ball $B$ in $\mathcal{B}_{h}^{k}$ that contains $x$. Hence, $x \in \operatorname{dom}_{h}^{k}(B)$. This shows that $\cup_{B \in \mathcal{B}_{h}^{k}} B \subset \cup_{B \in \mathcal{B}_{h}^{k}} \operatorname{dom}_{h}^{k}(B)$ and consequently $\cup_{B \in \mathcal{B}_{h}^{k}} B=$ $\cup_{B \in \mathcal{B}_{h}^{k}} \operatorname{dom}_{h}^{k}(B)$. Moreover, $\cup_{B \in \mathcal{B}_{h}^{k}} B=\mathcal{S} \times \mathcal{A}$ as it contains the initial ball which covers the whole space.
(b) Let $\left(B, B^{\prime}\right) \in \mathcal{B}_{h}^{k}$ two balls of radius $r>0$. Without loss of generality, we suppose that $B$ is created in episode $\tau \leq k$ with parent ball $B^{\text {pa }}$ and $B^{\prime}$ is created before $\tau$. According to the activation step in ZOOMRL algorithm, $\left(s_{h}^{\tau}, a_{h}^{\tau}\right)$ is the center of $B$ and $\left(s_{h}^{\tau}, a_{h}^{\tau}\right) \in \operatorname{dom}_{h}^{\tau}\left(B^{\text {pa }}\right)$. By the definition of a ball's domain, $\left(s_{h}^{\tau}, a_{h}^{\tau}\right) \notin B^{\prime}$, which proves that $\operatorname{dist}\left(B, B^{\prime}\right)>r$.

## C.1.2 Proof of Lemma 4

Proof. We fix $B \in \mathcal{B}_{h}^{k}$. For notation simplicity, denote $t=n_{h}^{k}(B)$. We have:

$$
\begin{aligned}
\widehat{\mathrm{Q}}_{h}^{k}(B)= & \left(1-\alpha_{t}\right) \cdot \widehat{\mathrm{Q}}_{h}^{k_{t}}(B)+\alpha_{t} \cdot\left(r_{h}\left(x_{h}^{k_{t}}, a_{h}^{k_{t}}\right)+\widehat{\mathrm{V}}_{h+1}^{k_{t}}\left(x_{h+1}^{k_{t}}\right)+u_{t}+2 L \cdot \operatorname{rad}(B)\right) \\
= & \left(1-\alpha_{t}\right) \cdot\left(\left(1-\alpha_{t-1}\right) \cdot \widehat{\mathrm{Q}}_{h}^{k_{t-1}}(B)\right. \\
& \left.+\alpha_{t-1} \cdot\left(r_{h}\left(x_{h}^{k_{t-1}}, a_{h}^{k_{t-1}}\right)+\widehat{\mathrm{V}}_{h+1}^{k_{t-1}}\left(x_{h+1}^{k_{t-1}}\right)+u_{t-1}+2 L \cdot \operatorname{rad}(B)\right)\right) \\
& +\alpha_{t} \cdot\left(r_{h}\left(x_{h}^{k_{t}}, a_{h}^{k_{t}}\right)+\widehat{\mathrm{V}}_{h+1}^{k_{t}}\left(x_{h+1}^{k_{t}}\right)+u_{i}+2 L \cdot \operatorname{rad}(B)\right) \\
= & \ldots \\
= & \prod_{i=1}^{t}\left(1-\alpha_{i}\right) H+\sum_{i=1}^{t} \alpha_{i} \prod_{j=i+1}^{t}\left(1-\alpha_{j}\right)\left(r_{h}\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\widehat{\mathrm{V}}_{h+1}^{k_{i}}\left(x_{h+1}^{k_{i}}\right)+u_{i}+2 L \cdot \operatorname{rad}(B)\right) \\
= & \alpha_{t}^{0} \cdot H+\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(r_{h}\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\widehat{\mathrm{V}}_{h+1}^{k_{i}}\left(x_{h+1}^{k_{i}}\right)+u_{i}+2 L \cdot \operatorname{rad}(B)\right),
\end{aligned}
$$

where the first step follows from the update rule of $\widehat{Q}_{h}^{k}$, the second step follows from the update rule for $\widehat{\mathrm{Q}}_{h}^{k_{t-1}}$, the third step follows from recursively representing $\widehat{\mathrm{Q}}_{h}^{k_{i}}$ using $\widehat{\mathrm{Q}}_{h}^{k_{i-1}}$ until $i=1$, and the last step follows from the definition of $\alpha_{t}^{0}$ and $\alpha_{t}^{i}$.

## C.1.3 Proof of Lemma 5

Proof. Let $B \in \mathcal{B}_{h}^{k}$ and $(s, a) \in \operatorname{dom}_{h}^{k}(B)$.
Since $\sum_{i=0}^{t} \alpha_{t}^{i}=1$, we have that $Q_{h}^{\star}(s, a)=\alpha_{t}^{0} Q_{h}^{\star}(s, a)+\sum_{i=1}^{t} \alpha_{t}^{i} Q_{h}^{*}(s, a)$.
By the Lipschitz assumption 4 and the fact $\forall i \in[t],\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \in B$ and $(s, a) \in$ $B$, we have:

$$
\left|Q_{h}^{\star}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)-Q_{h}^{\star}(s, a)\right| \leq L \cdot \operatorname{dist}\left(\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right),(x, a)\right) \leq 2 L \cdot \operatorname{rad}(B)
$$

Then we have

$$
\begin{align*}
& Q_{h}^{\star}(s, a) \geq \alpha_{t}^{0} Q_{h}^{\star}(s, a)+\sum_{i=1}^{t} \alpha_{t}^{i}\left(Q_{h}^{\star}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)-2 L \cdot \operatorname{rad}(B)\right)  \tag{C.1.1}\\
& Q_{h}^{\star}(s, a) \leq \alpha_{t}^{0} Q_{h}^{\star}(s, a)+\sum_{i=1}^{t} \alpha_{t}^{i}\left(Q_{h}^{\star}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+2 L \cdot \operatorname{rad}(B)\right) . \tag{C.1.2}
\end{align*}
$$

By Bellman equation, we have $Q_{h}^{\star}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)=r_{h}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\left[P_{h} V_{h+1}^{\star}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)$. Recall $\left[\hat{P}_{h}^{k_{i}} V_{h+1}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)=V_{h+1}\left(s_{h+1}^{k_{i}}\right)$, we have:

$$
Q_{h}^{\star}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)=r_{h}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\left[\left(P_{h}-\widehat{P}_{h}^{k_{i}}\right) V_{h+1}^{\star}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+V_{h+1}^{\star}\left(s_{h+1}^{k_{i}}\right) .
$$

Substitute the above equality into Eq. (C.1.1) and (C.1.2), we have:

$$
\begin{align*}
Q_{h}^{\star}(s, a) & \left.\geq \alpha_{t}^{0} Q_{h}^{\star}(x, a)\right)+\sum_{i=1}^{t} \alpha_{t}^{i}\left(r_{h}\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\left[\left(P_{h}-\widehat{P}_{h}^{k_{i}}\right) V_{h+1}^{\star}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right. \\
& \left.+V_{h+1}^{\star}\left(x_{h+1}^{k_{i}}\right)-2 L \cdot \operatorname{rad}(B)\right) \tag{C.1.3}
\end{align*}
$$

and

$$
\begin{align*}
Q_{h}^{\star}(s, a) & \left.\leq \alpha_{t}^{0} Q_{h}^{\star}(s, a)\right)+\sum_{i=1}^{t} \alpha_{t}^{i}\left(r_{h}\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\left[\left(P_{h}-\widehat{P}_{h}^{k_{i}}\right) V_{h+1}^{\star}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right. \\
& \left.+V_{h+1}^{\star}\left(x_{h+1}^{k_{i}}\right)+2 L \cdot \operatorname{rad}(B)\right) \tag{C.1.4}
\end{align*}
$$

Subtracting the formula in Lemma 4 from the two above inequalities, we have:
$\widehat{\mathrm{Q}}_{h}^{k}(B)-Q_{h}^{\star}(s, a) \geq \sum_{i=1}^{t} \alpha_{t}^{i}\left(\left(\widehat{\mathrm{~V}}_{h+1}^{k_{i}}-V_{h+1}^{\star}\right)\left(x_{h+1}^{k_{i}}\right)+\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+u_{i}\right)$
and

$$
\begin{align*}
\widehat{\mathrm{Q}}_{h}^{k}(B)-Q_{h}^{\star}(s, a) & \leq \alpha_{t}^{0} H+\sum_{i=1}^{t} \alpha_{t}^{i}\left(\left(\widehat{\mathrm{~V}}_{h+1}^{k_{i}}-V_{h+1}^{\star}\right)\left(x_{h+1}^{k_{i}}\right)+\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right. \\
& \left.+u_{i}+4 L \cdot \operatorname{rad}(B)\right) \tag{C.1.6}
\end{align*}
$$

## High probability bounds on the sampling noise

To ensure that our estimates concentrate around the true optimal $Q$-values, we need to ensure that the noise terms $\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)$, due to the next states sampling, are not large.

For each ball $B \in \mathcal{B}_{h}^{k}, k_{i}$ is the episode of which $B$ was selected as step $h$ for the $i$-th time. Let $\mathcal{F}_{t}$ be the $\sigma$-field generated by all the random variables until episode $t$, step $h$. As $\left\{k_{i}=t\right\} \in \mathcal{F}_{t}$, the random variable $k_{i}$ is a stopping time. By definition for any $i \geq 0, k_{i} \leq k_{i+1}$ so the $\sigma$-algebra $\mathcal{F}_{k_{i}}$ at time $k_{i}$ satisfies $\mathcal{F}_{k_{i}} \subset \mathcal{F}_{k_{i+1}}$ (see Lemma 17). Let's denote $\mathcal{G}_{i}=\mathcal{F}_{\tau_{i+1}}$. Then, $\left(\mathcal{G}_{i}\right)_{i}$ is a filtration. Moreover, via optional stopping [Chow and Teicher, 1998], $\left(\mathbb{I}\left(k_{i} \leq K\right)\left[\left(\widehat{P}_{h}^{k_{i}}-P\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right)_{i=1}^{\tau}$ is a $2 H$-bounded martingale difference sequence w.r.t the filtration $\left(\mathcal{G}_{i}\right)_{i \geq 0}$. By Azuma-Hoeffding 15, we have $\forall t>0, \tau \in[K]$
$\operatorname{Pr}\left[\left|\sum_{i=1}^{\tau} \alpha_{\tau}^{i} \cdot \mathbb{I}\left(k_{i} \leq K\right) \cdot\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right| \geq t\right] \leq 2 \exp \left(\frac{-t^{2}}{8 H^{2} \sum_{i=1}^{\tau}\left(\alpha_{\tau}^{i}\right)^{2}}\right)$

Let $p \in(0,1)$, by setting $2 \exp \left(\frac{-t^{2}}{8 H^{2} \sum_{i=1}^{\tau}\left(\alpha_{\tau}^{i}\right)^{2}}\right)=\frac{p}{2 H K^{2}}$, we have that for all $\tau \in[k]$ with probability at least $1-\frac{p}{2 H K^{2}}$ :

$$
\begin{aligned}
\left|\sum_{i=1}^{\tau} \alpha_{\tau}^{i} \cdot \mathbb{I}\left(k_{i} \leq K\right) \cdot\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right| & \leq 2 \sqrt{2} H \sqrt{\sum_{i=1}^{\tau}\left(\alpha_{\tau}^{i}\right)^{2} \ln \left(4 H K^{2} / p\right)} \\
& \leq 4 \sqrt{\frac{H^{3} \ln \left(4 H K^{2} / p\right)}{\tau}}=4 \sqrt{\frac{H^{3} \imath}{\tau}}
\end{aligned}
$$

where the second inequality follows from $\sum_{i=1}^{\tau}\left(\alpha_{\tau}^{i}\right)^{2} \leq \frac{2 H}{\tau}$ for any $\tau>0$ (see lemma 16). Then by union bound over $\tau \in[K]$, we have with probability at least $1-\frac{p}{2 H K}$

$$
\forall \tau \in[K], \quad\left|\sum_{i=1}^{\tau} \alpha_{\tau}^{i} \cdot \mathbb{I}\left(k_{i} \leq K\right) \cdot\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right| \leq 4 \sqrt{\frac{H^{3} l}{\tau}}
$$

Since the above inequality holds for all $\tau \in[K]$, it also holds for $\tau=t=n_{h}^{h}(B) \leq$ $K$. We also have that $\mathbf{I}\left(k_{i} \leq K\right)=1$ for any $i \leq n_{h}^{h}(B)$. As $\left|\mathcal{B}_{h}^{k}\right| \leq K$ for all $(h, k) \in[H] \times[K]$, using union bound for all balls and for all steps, we have with probability at least $1-p / 2: \forall(h, k) \in[H] \times[K]$ and for all ball $B \in \mathcal{B}_{h^{\prime}}^{k}$

$$
\begin{equation*}
\left|\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right| \leq 4 \sqrt{\frac{H^{3} \imath}{t}}, \text { where } t=n_{h}^{k}(B) \tag{C.1.7}
\end{equation*}
$$

According to Lemma 16, we have $1 / \sqrt{t} \leq \sum_{i=1} \frac{\alpha_{t}^{i}}{t} \leq 2 / \sqrt{t}$. This implies

$$
4 \sqrt{\frac{H^{3} l}{t}} \leq 4 \sqrt{H^{3} \imath} \cdot \sum_{i=1}^{t} \frac{\alpha_{t}^{i}}{t}=\sum_{i=1}^{t} \alpha_{t}^{i} u_{i}=\beta_{t} / 2 \leq 8 \sqrt{\frac{H^{3} l}{t}}
$$

Then Eq (C.1.7) gives that, with probability at least $1-p / 2: \forall(h, k) \in[H] \times$ $[K]$ and for all ball $B \in \mathcal{B}_{h^{\prime}}^{k}$

$$
\begin{equation*}
\left|\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right| \leq \beta_{t} / 2, \text { where } t=n_{h}^{k}(B) \tag{C.1.8}
\end{equation*}
$$

## Optimism of $Q$-values: Lemma 5 (a)

We proceed by induction. By definition, we have $\widehat{Q}_{H+1}^{k}=Q_{H+1}^{*}=0$ which implies $Q_{H+1}^{k}(B)-Q_{H+1}^{*}(s, a)=0$. Assume that $Q_{h+1}^{k}(B)-Q_{h+1}^{*}(s, a) \geq 0$.

Let $i \in[1, t]$, We have $V_{h+1}^{\star}\left(s_{h+1}^{k_{i}}\right)=Q_{h+1}^{\star}\left(s_{h+1}^{k_{i}}, \pi_{h+1}^{\star}\left(s_{h+1}^{k_{i}}\right)\right)$. As the set of domains of active balls covers the entire space, there exists $B^{\star} \in \mathcal{B}_{h+1}^{k}$ such that $\left(s_{h+1}^{k_{i}}, \pi_{h+1}^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \in \operatorname{dom}_{h+1}^{k}\left(B^{*}\right)$. By the definition of index, we have index ${ }_{h+1}^{k}\left(B^{\star}\right)=L \cdot \operatorname{rad}\left(B^{\star}\right)+\widehat{\mathrm{Q}}_{h+1}^{k}\left(\tilde{B}^{\star}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)$ for some active ball $\tilde{B}^{\star}$.

We have

$$
\begin{aligned}
& \widehat{\mathrm{V}}_{h+1}^{k}\left(s_{h+1}^{k_{i}}\right)-V_{h+1}^{\star}\left(s_{h+1}^{k_{i}}\right) \\
& =\min \left\{H, \max _{B \in \operatorname{rel}_{h+1}^{k}\left(s_{h+1}^{k_{i}}\right)}^{\operatorname{mindex}} \operatorname{ind}_{h+1}^{k}(B)\right\}-Q_{h+1}^{\star}\left(s_{h+1}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& \geq \max _{B \in \operatorname{rel}_{h+1}^{k}\left(s_{h+1}^{k_{i}}\right)} \operatorname{index}_{h+1}^{k}(B)-Q_{h+1}^{\star}\left(s_{h+1^{k}}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& \geq \operatorname{index}{ }_{h+1}^{k}\left(B^{\star}\right)-Q_{h+1}^{\star}\left(s_{h+1^{\prime}}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& =L \cdot \operatorname{rad}\left(B^{\star}\right)+\widehat{Q}_{h+1}^{k}\left(\tilde{B}^{\star}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)-Q_{h+1}^{\star}\left(s_{h+1}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& \geq L \cdot \operatorname{rad}\left(B^{\star}\right)+Q_{h+1}^{\star}\left(s_{\tilde{B}^{\star}}, a_{\tilde{B}^{\star}}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)-Q_{h+1}^{\star}\left(s_{h+1}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& \geq L \cdot \operatorname{rad}\left(B^{\star}\right)+Q_{h+1}^{\star}\left(s_{B^{\star}}, a_{B^{\star}}\right)-Q_{h+1}^{\star}\left(s_{h+1}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& \geq 0
\end{aligned}
$$

where $\left(s_{B^{\star}}, a_{B^{\star}}\right)$ and $\left(s_{\tilde{B}^{\star}}, a_{\tilde{B}^{\star}}\right)$ denote respectively the centers of balls $B^{\star}$ and $\tilde{B}^{\star}$. The first inequality follows from $Q_{h+1}^{\star}(s, a) \leq H$ for any state-action pair $(s, a)$. The third inequality follows from the induction hypothesis. The fourth and the last inequalities follow from Lipschitz assumption 4

Therefore, we have

$$
\begin{aligned}
Q_{h}^{k}(B)-Q_{h}^{*}(s, a) & \geq \sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(\left(\widehat{\mathrm{~V}}_{h+1}^{k_{i}}-V_{h+1}^{\star}\right)\left(s_{h+1}^{k_{i}}\right)+\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{\star}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+u_{i}\right) \\
& \geq-\beta_{t} / 2+\beta_{t} / 2=0
\end{aligned}
$$

## Upper bound: lemma 5 (b)

We have:

$$
\begin{aligned}
& \widehat{\mathrm{Q}}_{h}^{k}(B)-Q_{h}^{\star}(s, a) \\
& \leq \alpha_{t}^{0} \cdot H+\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(\left(V_{h+1}^{k_{i}}-V_{h+1}^{*}\right)\left(x_{h+1}^{k_{i}}\right)+\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+u_{i}+4 r(B)\right) \\
& \leq \alpha_{t}^{0} H+\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(\widehat{\mathrm{~V}}_{h+1}^{k_{i}}-V_{h+1}^{*}\right)\left(x_{h+1}^{k_{i}}\right)+\beta_{t} / 2+\sum_{i=1}^{t} \alpha_{t}^{i} u_{i}+4 L \sum_{i=1}^{t} \alpha_{t}^{i} \operatorname{rad}(B) \\
& \leq \alpha_{t}^{0} H+\beta_{t}+4 L \cdot \operatorname{rad}(B)+\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(\widehat{\mathrm{~V}}_{h+1}^{k_{i}}-V_{h+1}^{*}\right)\left(x_{h+1}^{k_{i}}\right)
\end{aligned}
$$

where the second inequality follows from the inequality C.1.8. The third inequality follows from $\sum_{i=1}^{t} \alpha_{t}^{i} \leq 1$

## C.1.4 Proof of lemma 7

Let $B_{h}^{k}$ the ball selected at step $h$ of episode $k$ and $B_{\text {init }}$ be the initial ball of radius one that covers the whole space. We need to distinguish between cases where $B_{h}^{k}=B_{\text {init }}$ or not. By the selection step in ZOOMRL algorithm, we have $\max _{B \in \operatorname{rel}_{h}^{k}\left(s_{h}^{k}\right)}$ index $_{h}^{k}(B)=\operatorname{index}{ }_{h}^{k}\left(B_{h}^{k}\right)$ and $\pi_{k}\left(s_{h}^{k}\right)=a_{h}^{k}$.

1. Case of $B_{h}^{k} \neq B_{\text {init }}$ : We denote $B_{h}^{k, \text { pa }}$ the parent of $B_{h}^{k}$.

$$
\begin{aligned}
\delta_{h}^{k} & =\left(\widehat{\mathrm{V}}_{h}^{k}-V_{h}^{\pi_{k}}\right)\left(s_{h}^{k}\right) \\
& \leq \max _{B \in \operatorname{rel}_{h}^{k}\left(s_{h}^{k}\right)} \operatorname{index}_{h}^{k}(B)-V_{h}^{\pi_{k}}\left(s_{h}^{k}\right) \\
& =\operatorname{index} k\left(B_{h}^{k}\right)-Q_{h}^{\pi_{k}}\left(s_{h}^{k}, a_{h}^{k}\right) \\
& \leq L \cdot \operatorname{rad}\left(B_{h}^{k}\right)+\widehat{\mathrm{Q}}_{h}^{k}\left(B_{h}^{k, \mathrm{pa}}\right)+L \cdot \operatorname{dist}\left(B_{h}^{k, \mathrm{pa}}, B_{h}^{k}\right)-Q_{h}^{\pi_{k}}\left(s_{h}^{k}, a_{h}^{k}\right) \\
& \leq L \cdot \operatorname{rad}\left(B_{h}^{k}\right)+\widehat{\mathrm{Q}}_{h}^{k}\left(B_{h}^{k, \mathrm{pa}}\right)+L \cdot \operatorname{rad}\left(B_{h}^{k, \mathrm{pa}}\right)-Q_{h}^{\pi_{k}}\left(s_{h^{\prime}}^{k} a_{h}^{k}\right) \\
& =3 L \cdot \operatorname{rad}\left(B_{h}^{k}\right)+\underbrace{\widehat{\mathrm{Q}}_{h}^{k}\left(B_{h}^{k, \mathrm{par}}\right)-Q_{h}^{*}\left(s_{h}^{k}, a_{h}^{k}\right)}_{q_{1}}+\left(Q_{h}^{*}-Q_{h}^{\pi_{k}}\right)\left(s_{h}^{k}, a_{h}^{k}\right)
\end{aligned}
$$

The third inequality follows from the fact that the center of $B_{h}^{k}$ is in $B_{h}^{k, p a}$ and the last equality follows from $\operatorname{rad}\left(B_{h}^{k, \mathrm{pa}}\right)=\operatorname{rad}\left(B_{h}^{k}\right)$.

Since $\left(x_{h}^{k}, a_{h}^{k}\right) \in \operatorname{dom}\left(B_{h}^{k}\right) \subset B_{h}^{k, \text { pa }}$, we have by Lemma 5

$$
\begin{aligned}
q_{1} & \leq \alpha_{n_{h}^{k}\left(B_{h}^{k, \mathrm{pa}}\right)}^{0} H+\beta_{n_{h}^{k}\left(B_{h}^{k, \mathrm{pa}}\right)}+4 L \cdot \operatorname{rad}\left(B_{h}^{k, \mathrm{pa}}\right) \\
& +\sum_{i=1}^{n_{h}^{k}\left(B_{h}^{k, \mathrm{par}}\right)} \alpha_{n_{h}^{k}\left(B_{h}^{k, \mathrm{pa}}\right)}^{i}\left(V_{h+1}^{k_{i}\left(B_{h}^{k, \mathrm{pa}}\right)}-V_{h+1}^{*}\right)\left(x_{h+1}^{k_{i}\left(B_{h}^{k, \mathrm{pa}}\right)}\right)
\end{aligned}
$$

where we denote by $k_{i}(B) \in\left[1, n_{h}^{k}(B)\right]$ the $i$-th episode where $B$ was selected by the algorithm at step $h$. As $B_{h}^{k, \text { pa }}$ is a parent, we have $n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)>0$ implying that $\alpha_{n_{h}^{k}\left(B_{h}^{k, p a}\right)}^{0}=\mathbb{I}\left\{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)=0\right\}=0$. Moreover, by the activation rule, we have $\frac{1}{\sqrt{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)}} \leq \operatorname{rad}\left(B_{h}^{k, \text { pa }}\right)$, implying that $\beta_{n_{h}^{k}\left(B_{h}^{k, p a}\right)} \leq$ $16 \sqrt{\frac{H^{3} \imath}{n_{h}^{k}\left(B_{h}^{k, p a}\right)}} \leq 16 \sqrt{H^{3}{ }_{\imath}} \operatorname{rad}\left(B_{h}^{k, \text { pa }}\right)=32 \sqrt{H^{3}}{ }_{\imath} \operatorname{rad}\left(B_{h}^{k}\right)$. Consequently,

$$
q_{1} \leq\left(8 L+32 \sqrt{H^{3}}\right) \operatorname{rad}\left(B_{h}^{k}\right)+\sum_{i=1}^{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)} \alpha_{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)}^{i} \phi_{h+1}^{k_{i}\left(B_{h}^{k, \text { pa }}\right)}
$$

and therefore,

$$
\begin{equation*}
\delta_{h}^{k} \leq\left(11 L+32 \sqrt{H^{3} \imath}\right) \operatorname{rad}\left(B_{h}^{k}\right)+\sum_{i=1}^{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)} \alpha_{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)}^{i} \phi_{h+1}^{k_{i}\left(B_{h}^{k, p a}\right)}+\left(Q_{h}^{*}-Q_{h}^{\pi_{k}}\right)\left(s_{h}^{k}, a_{h}^{k}\right) \tag{C.1.9}
\end{equation*}
$$

2. Case of $B_{h}^{k}=B_{\text {init }}$ :

$$
\begin{align*}
\delta_{h}^{k} & \leq \max _{B \in \operatorname{rel}_{h}^{k}\left(s_{h}^{k}\right)} \operatorname{index}_{h}^{k}(B)-V_{h}^{\pi_{k}}\left(s_{h}^{k}\right) \\
& =\operatorname{index}_{h}^{k}\left(B_{h}^{k}\right)-Q_{h}^{\pi_{k}}\left(s_{h}^{k} a_{h}^{k}\right) \\
& \leq L \cdot \operatorname{rad}\left(B_{h}^{k}\right)+\widehat{Q}_{h}^{k}\left(B_{h}^{k}\right)-Q_{h}^{\pi_{k}}\left(s_{h^{\prime}}^{k} a_{h}^{k}\right) \\
& =L \cdot \operatorname{rad}\left(B_{h}^{k}\right)+\widehat{Q}_{h}^{k}\left(B_{h}^{k}\right)-Q_{h}^{*}\left(s_{h^{\prime}}^{k} a_{h}^{k}\right)+\left(Q_{h}^{*}-Q_{h}^{\pi_{k}}\right)\left(s_{h^{\prime}}^{k} a_{h}^{k}\right) \\
& \leq 5 L \cdot \operatorname{rad}\left(B_{h}^{k}\right)+\alpha_{n_{h}^{k}\left(B_{h}^{k}\right)}^{0} H+\beta_{n_{h}^{k}\left(B_{h}^{k}\right)}+\sum_{i=1}^{n_{h}^{k}\left(B_{h}^{k}\right)} \alpha_{n_{h}^{k}\left(B_{h}^{k}\right)}^{i} \phi_{h+1}^{k_{i}\left(B_{h}^{k}\right)}+\left(Q_{h}^{*}-Q_{h}^{\pi_{k}}\right)\left(s_{h^{\prime}}^{k} a_{h}^{k}\right) \\
& \leq \alpha_{n_{h}^{k}\left(B_{h}^{k}\right)}^{0} H+\left(5 L+16 \sqrt{H^{3}}\right) \operatorname{rad}\left(B_{h}^{k}\right)+\sum_{i=1}^{n_{h}^{k}\left(B_{h}^{k}\right)} \alpha_{n_{h}^{k}\left(B_{h}^{k}\right)}^{i} \phi_{h+1}^{k_{i}\left(B_{h}^{k}\right)}+\left(Q_{h}^{*}-Q_{h}^{\pi_{k}}\right)\left(s_{h^{\prime}}^{k} a_{h}^{k}\right) \tag{С.1.10}
\end{align*}
$$

The third inequality follows from lemma 5 and the last inequality follows from the fact that $\operatorname{rad}\left(B_{h}^{k}\right)=\operatorname{rad}\left(B_{\text {init }}\right)=1$

Now, we can unify the bound (C.1.9) obtained in the first case where the algorithm selects a ball other than the initial ball and the bound (C.1.10) in second case where the initial ball is selected. To do that, we consider, by abuse of notation, that the initial ball is parent of itself i.e when $B_{h}^{k}=B_{\text {init }}$ we have $B_{h}^{k, p a}=B_{\text {init }}$ and we take the maximum over the two bounds

$$
\begin{align*}
\delta_{h}^{k} & \leq \alpha_{n_{h}^{k}\left(B_{h}^{k}\right)}^{0} H \mathbb{I}_{\left\{B_{h}^{k}=B_{\text {init }}\right\}}+\left(11 L+32 \sqrt{H^{3}} \imath\right) \operatorname{rad}\left(B_{h}^{k}\right) \\
& +\sum_{i=1}^{n_{h}^{k}\left(B_{h}^{k, \mathrm{pa}}\right)} \alpha_{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)}^{i} \phi_{h+1}^{k_{i}\left(B_{h}^{k, \mathrm{pa}}\right)}+\left(Q_{h}^{*}-Q_{h}^{\pi_{k}}\right)\left(s_{h}^{k} a_{h}^{k}\right) \tag{С.1.11}
\end{align*}
$$

we obtain the desired result but noting that

$$
\left.\begin{array}{l}
\left(Q_{h}^{\star}-Q_{h}^{\pi_{k}}\right)\left(s_{h}^{k}\right. \\
\left.a_{h}^{k}\right) \\
=\left[\mathbb{P}_{h}\left(V_{h}^{\star}-V_{h}^{\pi_{k}}\right)\right]\left(s_{h}{ }^{\prime}\right. \\
a_{h}^{k}
\end{array}\right) .
$$

## C.1.5 Bounding $\sum_{h=1}^{H} \sum_{k=1}^{K} \xi_{h+1}^{k}$

Lemma 12. With probability at least $1-p / 2$, we have

$$
\sum_{h=1}^{H} \sum_{k=1}^{K} \xi_{h+1}^{k} \leq 4 \sqrt{2 H^{3} K \imath}
$$

Let $\mathcal{F}_{k, h}$ be the $\sigma$-field generated by all the random variables until episode $k$, step $h$. Then, $\xi_{h+1}^{k}=\left[\left(P_{h}-\widehat{P}_{h}\right)\left(V_{h+1}^{*}-V_{h+1}^{\tau_{k}}\right)\right]\left(s_{h}^{k}, a_{h}^{k}\right)$ is a martingale difference sequence w.r.t the filtration $\left\{\mathcal{F}_{k, h}\right\}_{k, h \geq 0}$ bounded by $4 H$. By Azuma-Hoeffding (lemma 15), we have $\forall t>0, \operatorname{Pr}\left[\left|\sum_{h=1}^{H} \sum_{k=1}^{K} \xi_{h+1}^{k}\right| \geq t\right] \leq 2 \exp \left(\frac{-t^{2}}{32 H^{3} K}\right)$ Therefore,

$$
\begin{aligned}
\operatorname{Pr}\left[\left|\sum_{h=1}^{H} \sum_{k=1}^{K} \xi_{h+1}^{k}\right| \geq 4 \sqrt{2 H^{3} K ı}\right] & \leq 2 \exp \left(\frac{-32 H^{3} K \imath}{32 H^{3} K}\right) \\
& =2 \frac{p}{4 H^{2} K^{2}} \leq p / 2
\end{aligned}
$$

Hence, with probability at least $1-p / 2$, we have

$$
\sum_{h=1}^{H} \sum_{k=1}^{K} \xi_{h+1}^{k} \leq 4 \sqrt{2 H^{3} K^{2}}
$$

## C. 2 Misspecified Setting: Approximately Lipschtiz CASE

The proof structure is similar to the structure in Appendix B. We will particularly focus on the parts that require different treatments in the misspecified setting.

## C.2.1 Recursive Formula of $\widehat{\mathbf{Q}}_{h}^{k}(B)-Q_{h}^{\star}(s, a)$

Let $B \in \mathcal{B}_{h}^{k}$ and $(s, a) \in \operatorname{dom}_{h}^{k}(B)$.
Since $\sum_{i=0}^{t} \alpha_{t}^{i}=1$, we have that $Q_{h}^{\star}(s, a)=\alpha_{t}^{0} Q_{h}^{\star}(s, a)+\sum_{i=1}^{t} \alpha_{t}^{i} Q_{h}^{*}(s, a)$.
By the $\epsilon$-approximately Lipschitz assumption 5 and the fact $\forall i \in$ $[t],\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \in B$ and $(s, a) \in B$, we have:

$$
\left|Q_{h}^{\star}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)-Q_{h}^{\star}(s, a)\right| \leq L \cdot \operatorname{dist}\left(\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right),(x, a)\right)+2 \epsilon \leq 2 L \cdot \operatorname{rad}(B)+2 \epsilon .
$$

Then we have

$$
\begin{align*}
& Q_{h}^{\star}(s, a) \geq \alpha_{t}^{0} Q_{h}^{\star}(s, a)+\sum_{i=1}^{t} \alpha_{t}^{i}\left(Q_{h}^{\star}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)-2 L \cdot \operatorname{rad}(B)-2 \epsilon\right)  \tag{C.2.1}\\
& Q_{h}^{\star}(s, a) \leq \alpha_{t}^{0} Q_{h}^{\star}(s, a)+\sum_{i=1}^{t} \alpha_{t}^{i}\left(Q_{h}^{\star}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+2 L \cdot \operatorname{rad}(B)+2 \epsilon\right) . \tag{C.2.2}
\end{align*}
$$

By Bellman equation, we have $Q_{h}^{\star}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)=r_{h}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\left[P_{h} V_{h+1}^{\star}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)$. Recall $\left[\hat{P}_{h}^{k_{i}} V_{h+1}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)=V_{h+1}\left(s_{h+1}^{k_{i}}\right)$, we have:

$$
Q_{h}^{\star}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)=r_{h}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\left[\left(P_{h}-\widehat{P}_{h}^{k_{i}}\right) V_{h+1}^{\star}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+V_{h+1}^{\star}\left(s_{h+1}^{k_{i}}\right) .
$$

Substitute the above equality into Eq. (C.2.1) and (C.2.2), we have:

$$
\begin{aligned}
Q_{h}^{\star}(s, a) & \left.\geq \alpha_{t}^{0} Q_{h}^{\star}(s, a)\right)+\sum_{i=1}^{t} \alpha_{t}^{i}\left(r_{h}\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\left[\left(P_{h}-\widehat{P}_{h}^{k_{i}}\right) V_{h+1}^{\star}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right. \\
& \left.+V_{h+1}^{\star}\left(x_{h+1}^{k_{i}}\right)-2 L \cdot \operatorname{rad}(B)-2 \epsilon\right)
\end{aligned}
$$

$$
\begin{aligned}
Q_{h}^{\star}(s, a) \leq & \left.\alpha_{t}^{0} Q_{h}^{\star}(s, a)\right)+\sum_{i=1}^{t} \alpha_{t}^{i}\left(r_{h}\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+\left[\left(P_{h}-\widehat{P}_{h}^{k_{i}}\right) V_{h+1}^{\star}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right. \\
& \left.+V_{h+1}^{\star}\left(x_{h+1}^{k_{i}}\right)+2 L \cdot \operatorname{rad}(B)+2 \epsilon\right)
\end{aligned}
$$

Subtracting the formula in Lemma 4 from the two above inequalities, we have:

$$
\begin{align*}
\widehat{\mathrm{Q}}_{h}^{k}(B)-Q_{h}^{\star}(s, a) & \geq \sum_{i=1}^{t} \alpha_{t}^{i}\left(\left(\widehat{\mathrm{~V}}_{h+1}^{k_{i}}-V_{h+1}^{\star}\right)\left(x_{h+1}^{k_{i}}\right)+\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right. \\
& \left.+u_{i}-2 \epsilon\right) \tag{C.2.3}
\end{align*}
$$

$$
\begin{align*}
\widehat{\mathrm{Q}}_{h}^{k}(B)-Q_{h}^{\star}(s, a) & \leq \alpha_{t}^{0} H+\sum_{i=1}^{t} \alpha_{t}^{i}\left(\left(\widehat{\mathrm{~V}}_{h+1}^{k_{i}}-V_{h+1}^{\star}\right)\left(x_{h+1}^{k_{i}}\right)+\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right. \\
& \left.+u_{i}+4 L \cdot \operatorname{rad}(B)+2 \epsilon\right) \tag{C.2.4}
\end{align*}
$$

## C.2.2 Bounding of $Q_{h}^{k}(B)-Q_{h}^{*}(s, a)$

Lemma 13. Suppose Assumption 5 holds. For any $p \in(0,1)$, we have $\beta_{t}=$ $2 \sum_{i=1}^{t} \alpha_{t}^{i} u_{i} \leq 16 \sqrt{\frac{H^{3}}{t}}$ and, with probability at least $1-p / 2$, we have that for all $(s, a, h, k) \in \mathcal{S} \times \mathcal{A} \times[H] \times[K]$ and any ball $B$ such that $(s, a) \in \operatorname{dom}_{h}^{k}(B)$ :
(a) $\widehat{\mathrm{Q}}_{h}^{k}(B)-Q_{h}^{\star}(s, a) \geq-4(H-h+1) \epsilon$
(b) $\widehat{\mathrm{Q}}_{h}^{k}(B)-Q_{h}^{\star}(x, a) \leq \alpha_{t}^{0} \cdot H+\beta_{t}+4 L \cdot \operatorname{rad}(B)+2 \epsilon+\sum_{i=1}^{t} \alpha_{t}^{i}$. $\left(\widehat{\mathrm{V}}_{h+1}^{k_{i}}-V_{h+1}^{*}\right)\left(s_{h+1}^{k_{i}}\right)$
where $t=n_{h}^{k}(B)$ and $k_{1}, \cdots, k_{t}<k$ are the episodes where $B$ was selected at step $h$.

## C.2.3 High Probability Bound On The Sampling Noise

The same reasoning as in the subsection C.1.3 in the exact lipschitz case gives: with probability at least $1-p / 2: \forall(h, k) \in[H] \times[K]$ and for all ball $B \in \mathcal{B}_{h^{\prime}}^{k}$

$$
\begin{equation*}
\left|\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right| \leq \beta_{t} / 2, \text { where } t=n_{h}^{k}(B) \tag{C.2.5}
\end{equation*}
$$

## C.2.4 Approximate Optimism Of $Q$-values

We proceed by induction. By definition, we have $\widehat{Q}_{H+1}^{k}=Q_{H+1}^{*}=0$ which implies $Q_{H+1}^{k}(B)-Q_{H+1}^{*}(s, a)=-4(H-(H+1)+1) \epsilon$. Assume that $Q_{h+1}^{k}(B)-Q_{h+1}^{*}(s, a) \geq-4(H-(h+1)+1) \epsilon=-4(H-h) \epsilon$.

Let $i \in[1, t]$, We have $V_{h+1}^{\star}\left(s_{h+1}^{k_{i}}\right)=Q_{h+1}^{\star}\left(s_{h+1}^{k_{i}}, \pi_{h+1}^{\star}\left(s_{h+1}^{k_{i}}\right)\right)$. As the set of domains of active balls covers the entire space, there exists $B^{\star} \in \mathcal{B}_{h+1}^{k}$ such that $\left(s_{h+1}^{k_{i}}, \pi_{h+1}^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \in \operatorname{dom}_{h+1}^{k}\left(B^{*}\right)$. By the definition of index, we have $\operatorname{index}{ }_{h+1}^{k}\left(B^{\star}\right)=L \cdot \operatorname{rad}\left(B^{\star}\right)+\widehat{\mathrm{Q}}_{h+1}^{k}\left(\tilde{B}^{\star}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)$ for some ball $\tilde{B}^{\star}$.

We have

$$
\begin{aligned}
& \widehat{\mathrm{V}}_{h+1}^{k}\left(s_{h+1}^{k_{i}}\right)-V_{h+1}^{\star}\left(s_{h+1}^{k_{i}}\right) \\
& =\min \left\{H, \max _{B \in \operatorname{re1}_{h+1}^{k}\left(s_{h+1}^{k_{i}}\right)} \operatorname{index}_{h+1}^{k}(B)\right\}-Q_{h+1}^{\star}\left(s_{h+1^{\prime}}^{k_{i}} \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& \geq \max _{B \in \operatorname{rel}_{h+1}^{k}\left(s_{h+1}^{k_{i}}\right)} \operatorname{index}_{h+1}^{k}(B)-Q_{h+1}^{\star}\left(s_{h+1^{\prime}}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& \geq \operatorname{index}_{h+1}^{k}\left(B^{\star}\right)-Q_{h+1}^{\star}\left(s_{h+1}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& =L \cdot \operatorname{rad}\left(B^{\star}\right)+\widehat{Q}_{h+1}^{k}\left(\tilde{B}^{\star}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)-Q_{h+1}^{\star}\left(s_{h+1}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& \geq L \cdot \operatorname{rad}\left(B^{\star}\right)-4(H-h) \epsilon+Q_{h+1}^{\star}\left(s_{\tilde{B}^{\star}}, a_{\tilde{B}^{\star}}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)-Q_{h+1}^{\star}\left(s_{h+1}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& \geq L \cdot \operatorname{rad}\left(B^{\star}\right)-4(H-h) \epsilon+Q_{h+1}^{\star}\left(s_{B^{\star}}, a_{B^{\star}}\right)-2 \epsilon-Q_{h+1}^{\star}\left(s_{h+1}^{k_{i}}, \pi^{\star}\left(s_{h+1}^{k_{i}}\right)\right) \\
& \geq-(4(H-h)+2) \epsilon
\end{aligned}
$$

Where $\left(s_{B^{\star}}, a_{B^{\star}}\right)$ and $\left(s_{\tilde{B}^{\star}}, a_{\tilde{B}^{\star}}\right)$ denote respectively the centers of balls $B^{\star}$ and $\tilde{B}^{\star}$. The first inequality follows from $Q_{h+1}^{\star}(s, a) \leq H$ for any state-action pair $(s, a)$. The third inequality follows from the induction hypothesis. The fourth and the last inequalities follow from the assumption 5

Therefore, we have

$$
\begin{aligned}
Q_{h}^{k}(B)-Q_{h}^{*}(s, a) & \geq \sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(\left(\widehat{\mathrm{~V}}_{h+1}^{k_{i}}-V_{h+1}^{\star}\right)\left(s_{h+1}^{k_{i}}\right)+\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{\star}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+u_{i}-2 \epsilon\right) \\
& \geq-(4(H-h)+2) \epsilon-\beta_{t} / 2+\beta_{t} / 2-2 \epsilon=-4(H-h+1) \epsilon
\end{aligned}
$$

## C.2.5 Upper Bound of $\widehat{Q}_{h}^{k}(B)-Q_{h}^{\star}(s, a)$

We have:

$$
\begin{aligned}
& \widehat{\mathrm{Q}}_{h}^{k}(B)-Q_{h}^{\star}(s, a) \\
& \leq \alpha_{t}^{0} \cdot H+\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(\left(V_{h+1}^{k_{i}}-V_{h+1}^{*}\right)\left(x_{h+1}^{k_{i}}\right)+\left[\left(\widehat{P}_{h}^{k_{i}}-P_{h}\right) V_{h+1}^{*}\right]\left(x_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+u_{i}+4 \operatorname{rad}(B)+2 \epsilon\right) \\
& \leq \alpha_{t}^{0} H+\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(\widehat{\mathrm{~V}}_{h+1}^{k_{i}}-V_{h+1}^{*}\right)\left(x_{h+1}^{k_{i}}\right)+\beta_{t} / 2+\sum_{i=1}^{t} \alpha_{t}^{i} u_{i}+\sum_{i=1}^{t} \alpha_{t}^{i}(4 L \cdot \operatorname{rad}(B)+2 \epsilon) \\
& \leq \alpha_{t}^{0} H+\beta_{t}+4 L \cdot \operatorname{rad}(B)+2 \epsilon+\sum_{i=1}^{t} \alpha_{t}^{i} \cdot\left(\widehat{\mathrm{~V}}_{h+1}^{k_{i}}-V_{h+1}^{*}\right)\left(x_{h+1}^{k_{i}}\right)
\end{aligned}
$$

where the second inequality follows from the inequality C.2.5. The third inequality follows from $\sum_{i=1}^{t} \alpha_{t}^{i} \leq 1$
Lemma 14 (Approximate Optimism). Following the same setting as in Lemma 13, for any $(h, k)$, with probability at least $1-p / 2$, we have for any $s \in \mathcal{S}$ :

$$
\widehat{\mathrm{V}}_{h}^{k}(s) \geq V_{h}^{\star}(s)-(4(H-h)+5) \epsilon
$$

Proof. Let $s \in \mathcal{S}$, We have $V_{h}^{\star}(s)=Q_{h}^{\star}\left(s, \pi_{h}^{\star}(s)\right)$. As the set of domains of active balls covers the entire space, there exists $B^{\star} \in \mathcal{B}_{h+1}^{k}$ such that $\left(s, \pi_{h}^{\star}(s)\right) \in$ $\operatorname{dom}_{h}^{k}\left(B^{*}\right)$. By the definition of index, we have index ${ }_{h}^{k}\left(B^{\star}\right)=L \cdot \operatorname{rad}\left(B^{\star}\right)+$ $\widehat{\mathrm{Q}}_{h}^{k}\left(\tilde{B}^{\star}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)$ for some active ball $\tilde{B}^{\star}$.

We have

$$
\begin{aligned}
& \widehat{\mathrm{V}}_{h}^{k}(s)-V_{h}^{\star}(s) \\
& =\min \left\{H, \max _{B \in \operatorname{rel}_{h}^{k}(s)} \operatorname{index}\right. \\
& \geq \max _{B \in \operatorname{rel}_{h}^{k}(s)} \operatorname{index}_{h}^{k}(B)-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& \geq \operatorname{index}_{h}^{k}\left(B^{\star}\right)-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& =L \cdot \operatorname{rad}\left(B^{\star}\right)+\widehat{Q}_{h}^{k}\left(\tilde{B}^{\star}\right)+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& \geq L \cdot \operatorname{rad}\left(B^{\star}\right)+Q_{h}^{\star}\left(s_{\tilde{B}^{\star}}, a_{\tilde{B}^{\star}}\right)-4(H-h+1) \epsilon+L \cdot \operatorname{dist}\left(\tilde{B}^{\star}, B^{\star}\right)-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& \geq L \cdot \operatorname{rad}\left(B^{\star}\right)+Q_{h}^{\star}\left(s_{B^{\star}}, a_{B^{\star}}\right)-2 \epsilon-4(H-h+1) \epsilon-Q_{h}^{\star}\left(s, \pi^{\star}(s)\right) \\
& \geq-4 \epsilon-4(H-h+1) \epsilon=-(4(H-h)+5) \epsilon
\end{aligned}
$$

Where $\left(s_{B^{\star}}, a_{B^{\star}}\right)$ and ( $\left.s_{\tilde{B}^{\star}}, a_{\tilde{B}^{\star}}\right)$ denote respectively the centers of balls $B^{\star}$ and $\tilde{B}^{\star}$. The first inequality follows from $Q_{h}^{\star}(s, a) \leq H$ for any state-action pair $(s, a)$. The third inequality follows from lemma 13. The fourth and the last inequalities follow from assumption 5

## C.2.6 Regret Analysis

By the approximate optimism of our estimates with respect to the true value function (see lemma 14), we have with probability at least $1-p / 2$

$$
\begin{aligned}
\operatorname{REGRET}(K) & =\sum_{k=1}^{K}\left(V_{1}^{*}-V_{1}^{\pi_{k}}\right)\left(x_{1}^{k}\right) \\
& \leq \sum_{k=1}^{K}\left(\widehat{\mathrm{~V}}_{1}^{k}-V_{1}^{\pi_{k}}\right)\left(x_{1}^{k}\right)+K(4 H+1) \epsilon \\
& =\sum_{k=1}^{K} \delta_{1}^{k}+K(4 H+1) \epsilon
\end{aligned}
$$

Similarly to the Lemma 7, we can show that using Lemma 13 applied on $B_{h}^{k, p a}$ the parent of the selected ball at step $h$ of the episode $k$.

$$
\begin{aligned}
\delta_{h}^{k} & \leq \alpha_{n_{h}^{k}\left(B_{h}^{k}\right)}^{0} H \mathbb{I}_{\left\{B_{h}^{k}=B_{\text {init }}\right\}}+\left(11 L+32 \sqrt{H^{3}}\right) \\
& \operatorname{rad}\left(B_{h}^{k}\right) \\
& +\sum_{i=1}^{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)} \alpha_{n_{h}^{k}\left(B_{h}^{k, \text { pa }}\right)}^{i} \phi_{h+1}^{k_{i}\left(B_{h}^{k \text { pa }}\right)}-\phi_{h+1}^{k}+\delta_{h+1}^{k}+\xi_{h+1}^{k}+2 \epsilon
\end{aligned}
$$

Following the same steps of the Section 5.7.1 in the exact lipschitz setting, we obtain
$\left.\sum_{k=1}^{K} \delta_{h}^{k} \leq H+\left(11 L+32 \sqrt{H^{3}}\right)\right) \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right)+\left(1+\frac{1}{H}\right) \sum_{k=1}^{K} \delta_{h+1}^{k}+\sum_{k=1}^{K} \xi_{h+1}^{k}+2 K \epsilon$
By unrolling the last inequality for $h \in[H]$ and using the fact $\delta_{H+1}^{k}=0 \quad \forall k \in$ [K], we obtain

$$
\begin{align*}
\sum_{k=1}^{K} \delta_{1}^{k} & \leq \sum_{h=1}^{H}\left(1+\frac{1}{H}\right)^{h-1}\left(H+\left(11 L+32 \sqrt{H^{3} \imath}\right) \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right)+\sum_{k=1}^{K} \xi_{h+1}^{k}+2 K \epsilon\right) \\
& \leq 3 H^{2}+3\left(11 L+32 \sqrt{H^{3}} \imath\right) \sum_{h=1}^{H} \sum_{k=1}^{K} \operatorname{rad}\left(B_{h}^{k}\right)+3 \sum_{h=1}^{H} \sum_{k=1}^{K} \xi_{h+1}^{k}+6 H K \epsilon \tag{C.2.6}
\end{align*}
$$

Plugging bounds (5.7.2) and (5.7.3) from Section 5.7.1 in (C.2.6) and using union bound, we have with probability $1-p$

$$
\sum_{k=1}^{K} \delta_{1}^{k} \leq O\left(H^{2}+\sqrt{H^{3} K \imath}+\left(L+\sqrt{H^{5} \imath}\right) \min _{r_{0} \in(0,1)}\left\{K r_{0}+\sum_{\substack{r=2^{-i} \\ r \geq r_{0}}} \frac{M(r)}{r}\right\}+H K \epsilon\right)
$$

We obtain the regret bound in theorem 3 by noting that $\operatorname{REGRET}(K) \leq \sum_{k=1}^{K} \delta_{1}^{k}+$ $O(H K \epsilon)$.

## C. 3 TECHNICAL LEMMAS

Lemma 15 (Azuma-Hoeffding inequality). Suppose $\left\{X_{k}: k=0,1,2,3, \cdots\right\}$ is a martingale and $\left|X_{k}-X_{k-1}\right|<c_{k}$, almost surely. Then for all positive integers $N$ and all positive reals $t$,

$$
\operatorname{Pr}\left[\left|X_{N}-X_{0}\right| \geq t\right] \leq 2 \exp \left(\frac{-t^{2}}{2 \sum_{k=1}^{N} c_{k}^{2}}\right)
$$

Lemma 16 (Lemma 4.1 in Jin et al. [2018]). The following properties hold for $\alpha_{t}^{i}$ :
(a) $\frac{1}{\sqrt{t}} \leq \sum_{i=1}^{t} \frac{\alpha_{t}^{i}}{\sqrt{i}} \leq \frac{2}{\sqrt{t}}$ for every $t \geq 1$.
(b) $\max _{i \in[t]} \alpha_{t}^{i} \leq \frac{2 H}{t}$ and $\sum_{i=1}^{t}\left(\alpha_{t}^{i}\right)^{2} \leq \frac{2 H}{t}$ for every $t \geq 1$.
(c) $\sum_{t=i}^{\infty} \alpha_{t}^{i}=1+\frac{1}{H}$ for every $i \geq 1$.

## C.3.1 Few Reminders on Probability Theory

We consider a probability space $(\Omega, \mathcal{F}, P)$. We borrow notation from Qian et al. [2018]. We call filtration any increasing (for the inclusion) sequence of sub-$\sigma$-algebras of $\mathcal{F}$ i.e., $\left(\mathcal{F}_{n}\right)_{n \in \mathbb{N}}$ where $\forall n \in \mathbb{N}, \mathcal{F}_{n} \subset \mathcal{F}_{n+1} \subset \mathcal{F}$. We denote by $\mathcal{F}_{\infty}=\cup_{n \in \mathbb{N}} \mathcal{F}_{n}$.

Definition 7 (Stopping time). A random variable $\tau: \Omega: \rightarrow \mathbb{N} \cup\{+\infty\}$ is called stopping time with respect to a filtration $\left(\mathcal{F}_{n}\right)_{n \in \mathbb{N}}$ if for all $n \in \mathbb{N},\{\tau=n\} \in \mathcal{F}_{n}$.

Definition 8 ( $\sigma$-algebra at stopping time). let $\tau$ be a stopping time. An event prior to $\tau$ is any event $A \in \mathcal{F}_{\infty}$ s.t $A \cap \tau=n \in \mathcal{F}_{n}$ for all $n \in \mathbb{N}$. The set of events prior to $\tau$ is a $\sigma$-algebra denoted $\mathcal{F}_{r}$ and called $\sigma$-algebra at time $\tau$ :

$$
\mathcal{F}_{r}=\left\{A \in \mathcal{F}_{\infty}, \forall \in \mathbb{N}, A \cap \tau=n \in \mathcal{F}_{n}\right\}
$$

Lemma 17. let $\tau_{1}$ and $\tau_{2}$ be two stopping times with respect to the same filtration $\left(\mathcal{F}_{n}\right)_{n \in \mathbb{N}}$ s.t $\tau_{1} \leq \tau_{2}$ almost surely. Then $\mathcal{F}_{\tau_{1}} \subset \mathcal{F}_{\tau_{2}}$

## Appendix for Chapter 6

## D. 1 Technical Gaps in Published Bandit Papers

In this section, we highlight the technical error made by [Cheung et al., 2019] when controlling the bias term due to the non-stationarity of the reward function. Let us first recall the non-stationary linear bandit model

Definition 9 (Non-stationary linear bandit). At iteration $t$, the player makes a decision $A_{t}$ from a feasible set $\mathcal{A} \subset \mathbb{R}^{d}$, then observes the reward $r_{t}$ satisfying:

$$
\begin{equation*}
r_{t}=A_{t}^{\top} \boldsymbol{\theta}_{t}+z_{t} \tag{D.1.1}
\end{equation*}
$$

where $\boldsymbol{\theta}_{t}$ is the unknown regression parameter at iteration $t$ and $z_{t}$ is conditionally $\sigma$ subgaussian noise. We assume further that $\|A\| \leq 1, \forall A \in \mathcal{A}$ and $\left\|\boldsymbol{\theta}_{t}\right\| \leq S, \forall t$.

Cheung et al. [2019] propose the SW-UCB algorithm based on a sliding window approach of size $W$. At time $t$, actions are selected as follows:

$$
\begin{equation*}
A_{t}=\arg \max _{a \in \mathcal{A}} a^{\top} \hat{\boldsymbol{\theta}}_{t}+\beta\|a\|_{V_{t}^{-1}} \tag{D.1.2}
\end{equation*}
$$

where $\hat{\boldsymbol{\theta}}_{t}$ is the solution of the sliding window least squares problem:
$\hat{\boldsymbol{\theta}}_{t}=V_{t}^{-1} \sum_{\tau=\max \{1, t-W\}}^{t-1} A_{\tau} r_{\tau} \quad$, where $\quad V_{t}=\sum_{\tau=\max \{1, t-W\}}^{t-1} A_{\tau} A_{\tau}^{\top}+\lambda \cdot \mathbf{I}$ is the Gram matrix.
In the proof of lemma 1 in Cheung et al. [2019], the authors consider matrix $M=V_{t}^{-1} X$ where $X=\sum_{\tau=t-W}^{p} A_{\tau} A_{\tau}^{\top}$ for any integer $p \in\{t-W, \ldots, t-1\}$. They attempt to show that $M$ is positive semi-definite (PSD) (i.e $y^{\top} M y \geq 0, \forall y \in$ $\mathbb{R}^{d}$ ) as follows: they first prove that $M$ shares the same characteristic polynomial as the matrix $V_{t}^{-1 / 2} X V^{-1 / 2}$, then assert that since $V_{t}^{-1 / 2} X V^{-1 / 2}$ is PSD, $M$ is PSD as well.

Unfortunately, this last assertion does not hold in general. As a counterexample, let us consider the 2 dimensional identity matrix $\mathbf{I}$ and $B=\left((1,0)^{\top},(-10,1)^{\top}\right) . \quad \mathbf{I}$ and $B$ share the same characteristic polynomial
$p(x)=(x-1)^{2}$, $\mathbf{I}$ is obviously PSD but $B$ is not, as for $y=(1,1)^{\top}$, we have $y^{\top} B y=-8<0$.

Moreover, in general, a matrix of the form $M=V_{t}^{-1} X$ is not guaranteed to be PSD. If one sets $d=2, t=3, \lambda=1, A_{1}=(1,0)^{\top}$ and $A_{2}=(1,1)^{\top}$. with $A_{1}=$ $(1,0)^{\top}$ and $A_{2}=(1,1)^{\top}$, we have $M=V_{t}^{-1} A_{1} A_{1}^{\top}=\left((0.4,-0.2)^{\top},(0,0)^{\top}\right)$. If we consider $y=(1,5)^{\top}$, we have $y^{\top} M y=-0.6<0$.

## D. 2 Regret Reanalysis of D-LINUCB

Russac et al. [2019] propose the D-LINUCB algorithm, based on sequential weighted least squares regression. At time $t$, actions are selected as follows:

$$
\begin{equation*}
x_{t}=\arg \max _{x \in \mathcal{X}_{t}} x^{\top} \hat{\boldsymbol{\theta}}_{t}+\beta\|x\|_{V_{t}^{-1} \tilde{V}_{t} V_{t}^{-1}} \tag{D.2.1}
\end{equation*}
$$

where $V_{t}=\sum_{\tau=1}^{t-1} \eta^{-\tau} x_{\tau} x_{\tau}^{\top}+\lambda \eta^{-(t-1)} \cdot \mathbf{I}$ is the Gram matrix, $\tilde{V}_{t}=$ $\sum_{\tau=1}^{t-1} \eta^{-2 \tau} x_{\tau} x_{\tau}^{\top}+\lambda \eta^{-2(t-1)} \cdot \mathbf{I}$ and $\hat{\boldsymbol{\theta}}_{t}$ is the solution the weighted least squares problem:

$$
\begin{equation*}
\hat{\boldsymbol{\theta}}_{t}=V_{t}^{-1} \sum_{\tau=1}^{t-1} \eta^{-\tau} x_{\tau} r_{\tau} \tag{D.2.2}
\end{equation*}
$$

As our analysis follows the same proof steps as in Russac et al. [2019], we will only highlight our proposed fix to their technical error and the changes that it induces.

Non-stationarity bias Let $\overline{\boldsymbol{\theta}}_{t} \triangleq V_{t}^{-1} \sum_{\tau=1}^{t-1} \eta^{-\tau} x_{\tau} x_{\tau}^{\top} \boldsymbol{\theta}_{\tau}+\lambda \eta^{-(t-1)} \boldsymbol{\theta}_{t}$ the weighted average of the true regression parameters. To characterize the bias, Russac et al. [2019] attempt to control directly $\left\|\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right\|$. Instead, we propose to control $\left|x^{\top}\left(\boldsymbol{\theta}_{\underline{t}}-\overline{\boldsymbol{\theta}}_{t}\right)\right|$ for any $x \in \mathcal{X}$ and then use the fact that $\left\|\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right\|=\max _{x:\|x\|=1}\left|x^{\top}\left(\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right)\right|$

$$
\begin{aligned}
\left|x^{\top}\left(\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right)\right| & =\left|x^{\top} V_{t}^{-1} \sum_{\tau=1}^{t-1} \eta^{-\tau} x_{\tau} x_{\tau}^{\top}\left(\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right)\right| \\
& \leq \underbrace{\left|x^{\top} V_{t}^{-1} \sum_{\tau=t-W}^{t-1} \eta^{-\tau} x_{\tau} x_{\tau}^{\top}\left(\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right)\right|}_{(\star)}+\underbrace{\left|x^{\top} V_{t}^{-1} \sum_{\tau=1}^{t-W-1} \eta^{-\tau} x_{\tau} x_{\tau}^{\top}\left(\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right)\right|}_{(\star \star)}
\end{aligned}
$$

Bound on ( $\star$ ):

$$
\begin{aligned}
& \left|x^{\top} V_{t}^{-1} \sum_{\tau=t-W}^{t-1} \eta^{-\tau} x_{\tau} x_{\tau}^{\top}\left(\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right)\right| \\
& \leq \sum_{\tau=t-W}^{t-1} \eta^{-\tau}\left|x^{\top} V_{t}^{-1} x_{\tau}\right| \cdot\left|x_{\tau}^{\top}\left(\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right)\right| \quad \quad \text { (triangle inequality) } \\
& =\sum_{\tau=t-W}^{t-1} \eta^{-\tau}\left|x^{\top} V_{t}^{-1} x_{\tau}\right| \cdot\left|x_{\tau}^{\top}\left(\sum_{s=\tau}^{t-1}\left(\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right)\right)\right| \\
& \leq \sum_{\tau=t-W}^{t-1} \eta^{-\tau}\left|x^{\top} V_{t}^{-1} x_{\tau}\right| \cdot\left\|x_{\tau}\right\| \cdot\left\|\sum_{s=\tau}^{t-1}\left(\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right)\right\| \quad \quad \text { (Cauchy-Schwarz) } \\
& \leq \sum_{\tau=t-W}^{t-1} \eta^{-\tau}\left|x^{\top} V_{t}^{-1} x_{\tau}\right| \cdot \sum_{s=\tau}^{t-1}\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\| \\
& \leq \sum_{s=t-W}^{t-1} \sum_{\tau=t-W}^{s} \eta^{-\tau}\left|x^{\top} V_{t}^{-1} x_{\tau}\right| \cdot\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\| \\
& \left.\leq \sum_{s=t-W}^{t-1} \sqrt{\left[\sum_{\tau=t-W}^{s} \eta^{-\tau} x^{\top} V_{t}^{-1} x\right] \cdot\left[\sum_{\tau} \| \leq 1\right)} \sum_{\tau=t-W}^{s} \eta^{-\tau} x_{\tau}^{\top} V_{t}^{-1} x_{\tau}\right]
\end{aligned}\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\| l
$$

(Cauchy-Schwarz)

$$
\leq \sum_{s=t-W}^{t-1} \sqrt{\left[\sum_{\tau=t-W}^{s} \eta^{-\tau} x^{\top} V_{t}^{-1} x\right] \cdot d} \cdot\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\|
$$

$$
\leq\|x\| \sqrt{d} \sum_{s=t-W}^{t-1} \sqrt{\frac{\sum_{\tau=t-W}^{t-1} \eta^{-\tau}}{\lambda \eta^{-(t-1)}}} \cdot\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\| \quad\left(\lambda_{\max }\left(V_{t}^{-1}\right) \leq \frac{1}{\lambda \eta^{-(t-1)}}\right)
$$

$$
\leq\|x\| \sqrt{\frac{d}{\lambda(1-\eta)}} \sum_{s=t-W}^{t-1}\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\|
$$

Bound on ( $\star \star$ ):

$$
\begin{aligned}
&\left|x^{\top} V_{t}^{-1} \sum_{\tau=1}^{t-W-1} \eta^{-\tau} x_{\tau} x_{\tau}^{\top}\left(\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right)\right| \leq\|x\|\left\|V_{t}^{-1} \sum_{\tau=1}^{t-W-1} \eta^{-\tau} x_{\tau} x_{\tau}^{\top}\left(\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right)\right\| \\
& \leq\|x\| \frac{1}{\lambda \eta^{-(t-1)}}\left\|\sum_{\tau=1}^{t-W-1} \eta^{-\tau} x_{\tau} x_{\tau}^{\top}\left(\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right)\right\| \\
&\left(\left\|V_{t}^{-1}\right\|=\lambda_{\max }\left(V_{t}^{-1}\right) \leq \frac{1}{\lambda \eta^{-(t-1)}}\right) \\
& \leq\|x\| \frac{1}{\lambda} \sum_{\tau=1}^{t-W-1} \eta^{(t-1-\tau)}\left\|x_{\tau}\right\|^{2}\left\|\boldsymbol{\theta}_{\tau}-\boldsymbol{\theta}_{t}\right\| \\
& \leq\|x\| \frac{2 S}{\lambda} \frac{\eta^{W}}{1-\eta} \quad\left(\left\|\boldsymbol{\theta}_{t}\right\| \leq S \text { and }\left\|x_{t}\right\| \leq 1\right)
\end{aligned}
$$

We conclude for any $x \in \mathbb{R}^{d}$

$$
\left|x^{\top}\left(\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right)\right| \leq\|x\|\left(\sqrt{\frac{d}{\lambda(1-\eta)}} \sum_{s=t-W}^{t-1}\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\|+\frac{2 S}{\lambda} \frac{\eta^{W}}{1-\eta}\right)
$$

which proves that

$$
\left\|\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right\| \leq \sqrt{\frac{d}{\lambda(1-\eta)}} \sum_{s=t-W}^{t-1}\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\|+\frac{2 S}{\lambda} \frac{\eta^{W}}{1-\eta}
$$

Comparing to the bound on $\left\|\boldsymbol{\theta}_{t}-\overline{\boldsymbol{\theta}}_{t}\right\|$ in the proof of Russac et al. [2019], there is an extra factor $\sqrt{\frac{d}{\lambda(1-\eta)}}$ that multiplies the local non-stationarity term $\sum_{s=t-W}^{t-1}\left\|\boldsymbol{\theta}_{s}-\boldsymbol{\theta}_{s+1}\right\|$. This extra factor will consequently multiply the variation budget term in the final regret as stated in the following proposition:

Proposition 10. Under the assumption that $\sum_{t=1}^{K-1}\left\|\boldsymbol{\theta}_{t}-\boldsymbol{\theta}_{t+1}\right\| \leq \Delta$, for any $\delta \in$ $(0,1)$, if we set $\beta=\sqrt{\lambda} S+\sigma \sqrt{2 \log (1 / \delta)+d \log \left(1+\frac{1}{\lambda d(1-\gamma)}\right)}$ in the algorithm 1 D-LinUCB of Russac et al. [2019], then with probability $1-\delta$, for any $W>0$ the dynamic regret of $\mathrm{D}-\mathrm{LiNUCB}$ is at most

$$
\mathcal{O}\left(\sqrt{\frac{d}{\lambda(1-\eta)}} \Delta W+\frac{S}{\lambda} \frac{\eta^{W}}{1-\eta} K+\beta \sqrt{d K} \sqrt{K \log (1 / \eta)+\log \left(1+\frac{1}{d \lambda(1-\eta)}\right)}\right)
$$

Proposition 11. Under the same assumption as 10 If we set $\log (1 / \eta)=$ $d^{-1 / 4} \Delta^{1 / 2} K^{-1 / 2}, W=\frac{\log (K /(1-\eta))}{\log (1 / \eta)}$ and $\lambda=1$; for any $\delta \in(0,1)$; we have that with probability $1-\delta$, the dynamic regret of D-LINUCB is at most $\widetilde{\mathcal{O}}\left(d^{7 / 8} \Delta^{1 / 4} K^{3 / 4}\right)$.

Proof. With the choice $\log (1 / \eta)=d^{-1 / 4} \Delta^{1 / 2} K^{-1 / 2}$ and $W=\frac{\log (K /(1-\eta))}{\log (1 / \eta)}$; we have $\frac{\eta^{W}}{1-\eta} K=1, \eta=\exp \left(-\left(\frac{\Delta}{K}\right)^{1 / 2}\right)_{K \rightarrow \infty}^{\sim} 1-d^{-1 / 4} \Delta^{1 / 2} K^{-1 / 2}$ so that $\sqrt{\frac{d}{\lambda(1-\eta)}} \Delta W \sim \sqrt{d} \Delta \log (K /(1-\eta))\left(d^{1 / 4} \Delta^{-1 / 2} K^{1 / 2}\right)^{3 / 2}=\tilde{\mathcal{O}}\left(d^{7 / 8} \Delta^{1 / 4} K^{3 / 4}\right)$ and $\beta \sqrt{d K} \sqrt{K \log (1 / \eta)+\log \left(1+\frac{1}{d \lambda(1-\eta)}\right)}=\tilde{\mathcal{O}}\left(d K\left(d^{-1 / 4} \Delta^{1 / 2} K^{-1 / 2}\right)^{1 / 2}\right)=$ $\tilde{\mathcal{O}}\left(d^{7 / 8} \Delta^{1 / 4} K^{3 / 4}\right)$.

## D. 3 Regret Analysis of opt-WlSvi and Proof Outline

## D.3.1 Single Step Error Decomposition

In this section, we analyse the one-step error decomposition of the difference between the estimates $Q_{t, h}$ and $Q_{t, h}^{\pi}$ of a given policy $\pi$. To do that, we use the weighted $\operatorname{MDP}(\mathcal{S}, \mathcal{A}, \bar{P}, \bar{r})$ to isolate the bias term. The decomposition contains four parts: the reward bias and variance, the transition bias and variance, and the difference in value functions at step $h+1$. It can be written as:

$$
\begin{aligned}
\boldsymbol{\phi}(s, a)^{\top} \mathbf{w}_{t, h}-Q_{t, h}^{\pi}(s, a)= & \underbrace{\left(\bar{r}_{t, h}-r_{t, h}\right)(s, a)}_{\text {reward bias }}+\underbrace{\left(\widehat{r}_{t, h}-\bar{r}_{t, h}\right)(s, a)}_{\text {reward variance }} \\
& +\underbrace{\left[\left(\bar{P}_{t, h}-P_{t, h}\right) V_{t, h+1}^{\pi}\right](s, a)}_{\text {transition bias }}+\underbrace{\left[\left(\widehat{P}_{t, h}-\bar{P}_{t, h}\right) V_{t, h}\right](s, a)}_{\text {transition variance }} \\
& +\underbrace{\left[\bar{P}_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)}_{\text {difference in value functions of next step }} .
\end{aligned}
$$

The reward and transition bias terms are controlled by Lemma 9 using the fact that $\left\|V_{t, h}^{\pi}\right\|_{\infty} \leq H$. The difference in value-functions at step $h+1$ can be rewritten as $\left[P_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)+\left[\left(\bar{P}_{t, h}-P_{t, h}\right)\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)$. We control the second term by applying again Lemma 9 since $\left\|V_{t, h+1}-V_{t, h+1}^{\pi}\right\|_{\infty} \leq H$.

It remains now the two variance terms. The reward variance is easy to control and it reduces simply to the bias due to the regularization as we assume that $r$ is a deterministic function. Note that the assumption of deterministic reward is not a limiting assumption since the contribution of a stochastic reward in the final regret has lower order term than the contribution of a stochastic transition. We have, using the Cauchy-Shwartz inequality and $\left\|\widetilde{\Sigma}_{t, h}^{-1}\right\| \leq \frac{1}{\lambda \eta^{-2(t-1)}}$ :

$$
\begin{aligned}
\left|\left(\widehat{r}_{t, h}-\bar{r}_{t, h}\right)(s, a)\right| & =\lambda \eta^{-(t-1)}\left|\boldsymbol{\phi}(s, a)^{\top} \Sigma_{t, h}^{-1} \boldsymbol{\theta}_{t, h}\right| \\
& \leq \sqrt{d \lambda}\|\boldsymbol{\phi}(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}
\end{aligned}
$$

Controlloing the transition variance is more involved, and we differ the analysis to the next section. If we define bias $\triangleq$ bias $_{r}+$ bias $_{P}$ the total non-stationarity bias of the MDP, we can summarize the one-step analysis as follows:

$$
\begin{align*}
\phi(s, a)^{\top} \mathbf{w}_{t, h}-Q_{t, h}^{\pi}(s, a) \leq & 2 H \operatorname{Hias}(t, h)+\left[P_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a) \quad \text { (D.3.1 }  \tag{D.3.1}\\
& +\sqrt{d \lambda}\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}+\left[\left(\widehat{P}_{t, h}-\bar{P}_{t, h}\right) V_{t, h}\right](s, a)
\end{align*}
$$

## D.3.2 High Probability Bound on the Transition Variance

In this section, we will establish a high probability bound on the term ( $\widehat{P}_{t, h}-$ $\left.\bar{P}_{t, h}\right) V_{t, h}$. From the definitions of $\widehat{P}$ and $\bar{P}$ and the Cauchy-Schwartz inequality, we have

$$
\left[\left(\widehat{P}_{t, h}-\bar{P}_{t, h}\right) V_{t, h}\right](s, a) \leq\left(\left\|\sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau, h} \epsilon_{\tau, h}\right\|_{\tilde{\Sigma}_{t, h}^{-1}}+H \sqrt{d \lambda}\right)\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}
$$

where $\epsilon_{\tau, h}=V_{t, h+1}\left(s_{\tau, h+1}\right)-\left[P_{t, h} V_{t, h+1}\right]\left(s_{\tau, h}, a_{\tau, h}\right)$. If $V_{t, h+1}$ was a fixed function, $\epsilon_{\tau, h}$ would be zero-mean conditioned on the history of transitions up to step $h$ at episode $\tau$ and we would use the concentration of weighted selfnormalized processes [Russac et al., 2019] to get a high probability bound on $\left\|\sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau, h} \epsilon_{\tau, h}\right\|_{\tilde{\Sigma}_{t, h}^{-1}}$. However, as $V_{t, h+1}$ is estimated from past transitions and thus depends on the latter in a non-trivial way, we will show a concentration bound that holds uniformly for all possible value functions generated by the algorithm. We proceed first by establishing the boundness of iterates in the next Lemma.

Lemma 18 (Boundness of iterates). For any $(t, h) \in[K] \times[H]$, the weight $\mathbf{w}_{t, h}$ and the matrix $\Sigma_{t}^{-1} \widetilde{\Sigma}_{t} \Sigma_{t}^{-1}$ in Algorithm 6 satisfies:

$$
\left\|\mathbf{w}_{t, h}\right\| \leq 2 H \sqrt{\frac{d\left(1-\eta^{t-1}\right)}{\lambda(1-\eta)}} \text { and }\left\|\Sigma_{t}^{-1} \widetilde{\Sigma}_{t} \Sigma_{t}^{-1}\right\| \leq \frac{1}{\lambda}
$$

Any value function estimate produced by Algorithm 6 could be written in the following form

$$
V^{\mathbf{w}, \mathbf{A}}(\cdot)=\min \left\{\max _{a \in \mathcal{A}}\left\{\mathbf{w}^{\top} \boldsymbol{\phi}(\cdot, a)+\sqrt{\boldsymbol{\phi}(\cdot, a)^{\top} \mathbf{A} \boldsymbol{\phi}(\cdot, a)}\right\}, H\right\}
$$

where $\mathbf{w} \in \mathbb{R}^{d}$ and $\mathbf{A} \in \mathbb{R}^{d \times d}$ is a symmetric definite positive matrix that are in

$$
\mathcal{G}=\left\{\mathbf{w}, \mathbf{A}:\|\mathbf{w}\| \leq 2 H \sqrt{\frac{d}{\lambda(1-\eta)}} \text { and }\|\mathbf{A}\|_{\mathrm{F}} \leq \frac{\sqrt{d} \beta^{2}}{\lambda}\right\}
$$

The $\epsilon$-covering number of $\mathcal{G}$, identified as Euclidean ball in $\mathbb{R}^{d+d^{2}}$ of radius $2 H \sqrt{\frac{d}{\lambda(1-\eta)}}+\frac{\beta^{2} \sqrt{d}}{\lambda}$, is bounded by $\left(3\left(2 H \sqrt{\frac{d}{\lambda(1-\eta)}}+\frac{\beta^{2} \sqrt{d}}{\lambda}\right) / \epsilon\right)^{d+d^{2}}$. The latter number is exponential in the dimension $d$ but only the square root of its logarithm, which is linear in $d$, will contribute to the bound as we will show next.

By applying the concentration of weighted self-normalized processes [Russac et al., 2019] and using a union bound argument over an $\epsilon$-net of $\mathcal{G}$ with an appropriate value of $\epsilon$, we obtain the desired high probability bound stated in the following Lemma

Lemma 19. For any $\delta \in(0,1)$, with probability at least $1-\delta / 2$, we have for all $(t, h) \in[K] \times[H]$,

$$
\left\|\sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau, h} \epsilon_{\tau, h}\right\|_{\widetilde{\Sigma}_{t, h}^{-1}} \leq C d H \sqrt{\log \left(\frac{d H \beta}{\lambda(1-\eta)} \cdot \frac{2}{\delta}\right)}
$$

where $C>0$ is an absolute constant.
Finally, by combining the single error decomposition in Equation (D.3.1) and the transition concentration in Lemma 19 with an appropriate choice of $\beta$, we obtain the following high probability single-step bound.

Lemma 20 (Key lemma). There exists an absolute value c such that $\beta=c d H \sqrt{1}$ where $\tau=\log \left(\frac{2 d H}{(1-\eta) \delta}\right), \lambda=1$ and for any fixed policy $\pi$, we have with probability at least $1-\delta / 2$ for all $(s, a, h, t) \in \mathcal{S} \times \mathcal{A} \times[H] \times[K]$,

$$
\begin{aligned}
& \left|\phi(s, a)^{\top} \mathbf{w}_{t, h}-Q_{t, h}^{\pi}(s, a)-\left[P_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)\right| \\
& \quad \leq 2 \operatorname{Hbias}(t, h)+\beta\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}} .
\end{aligned}
$$

## D.3.3 Optimism

Now, we show that the true value functions can be upper bounded by the value functions computed by OPT-WLSVI plus a bias term. In fact, unlike the stationary case, we act optimistically with respect to the weighted average MDP. To prove this, we use the key Lemma 20 in the previous section and we proceed by induction argument over steps $h \in[H]$.

Lemma 21 (Optimism). For all $(s, a, t, h) \in \mathcal{S} \times \mathcal{A} \times[K] \times[H]$, we have with probability at least $1-\delta / 2$

$$
\begin{equation*}
Q_{t, h}(s, a)+2 H \sum_{h^{\prime}=h}^{H} b i a s(t, h) \geq Q_{t, h}^{\star}(s, a) \tag{D.3.2}
\end{equation*}
$$

## D.3.4 Final Regret Analysis

Now, having the results provided in previous sections at hand, we turn to proving the regret bound of our algorithm. Let $\pi_{t}$ the policy executed by the algorithm in step $h$ for $H$ steps to reach the end of the episode. If we define $\delta_{t, h} \triangleq V_{t, h}\left(s_{t, 1}\right)-V_{t, h}^{\pi_{t}}\left(s_{t, h}\right)$, a straightforward application of Lemma 21 is that the regret is upper bounded by the sum of $\delta_{t, h}$ and bias terms with probability at least $1-\delta / 2$ i.e

$$
\begin{equation*}
\operatorname{REGRET}(K) \underbrace{\leq}_{\text {by optimism }} \sum_{t=1}^{K} \delta_{t, h}+2 H \sum_{t=1}^{K} \sum_{h=1}^{H} \operatorname{bias}(t, h) . \tag{D.3.3}
\end{equation*}
$$

The policy $\pi_{t}$ is the greedy policy with respect to $Q_{t, h}$, and $a_{t, h}=\pi_{t}\left(s_{t}, h\right)=$ $\arg \max _{a \in \mathcal{A}} Q_{t, h}\left(s_{t, h}, a\right)$. Therefore, we have $\delta_{t, h}=Q_{t, h}\left(s_{t, h}, a_{t, h}\right)-Q_{t, h}^{\pi_{t}}\left(s_{t, h}, a_{t, h}\right)$. Using the definition of $Q_{t, h}$ and the key Lemma 20, we obtain with probability at least $1-\delta / 2$,

$$
\begin{aligned}
\delta_{t, h} & \leq\left[P_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi_{t}}\right)\right]\left(s_{t, h}, a_{t, h}\right)+2 \beta\left\|\phi_{t, h}\right\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}+2 \operatorname{Hbias}(t, h) \\
& =\delta_{t, h+1}+\xi_{t, h+1}+2 \beta\left\|\phi_{t, h}\right\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}+2 \operatorname{Hbias}(t, h)
\end{aligned}
$$

where we define $\xi_{t, h+1}=\left[P_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi_{t}}\right)\right]\left(s_{t, h,} a_{t, h}\right)-\left(V_{t, h+1}-\right.$ $\left.V_{t, h+1}^{\pi_{t}}\right)\left(s_{t, h+1}\right)$. Unrolling the last inequality $H$ times, we obtain

$$
\begin{equation*}
\delta_{t, 1} \leq \sum_{h=1}^{H} \widetilde{\xi}_{t, h}+2 \beta \sum_{h=1}^{H}\left\|\phi_{t, h}\right\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}+2 H \sum_{h=1}^{H} \operatorname{bias}(t, h) \tag{D.3.4}
\end{equation*}
$$

Hence, by combining Equations (D.3.3) and (D.3.4), we obtain with probability at least $1-\delta / 2$,

$$
\begin{align*}
\operatorname{REGRET}(K) \leq & \underbrace{\sum_{t=1}^{K} \sum_{h=1}^{H} \xi_{t, h}}_{(A)}+2 \beta \underbrace{\sum_{t=1}^{K} \sum_{h=1}^{H}\left\|\phi_{t, h}\right\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}}_{(B)} \\
& +4 H \underbrace{\sum_{t=1}^{K} \sum_{h=1}^{H} \operatorname{bias}(t, h)}_{(C)} . \tag{D.3.5}
\end{align*}
$$

Now, we proceed to upper bound the different terms in the RHS of Equation (D.3.5).

Term (A): The computation of $V_{t, h}$ is independent from $\left(s_{t, h}, a_{t, h}\right)$, therefore, $\left\{\xi_{t, h}\right\}$ is $2 H$-bounded martingale difference sequence. Therefore, by AzumaHoeffding, we have for all $t>0, P\left(\left|\sum_{t=1}^{K} \sum_{h=1}^{H} \xi_{t, h}\right| \geq t\right) \leq 2 \exp \left(\frac{-t^{2}}{16 H^{3} K}\right)$. Then, $P\left(\left|\sum_{t=1}^{K} \sum_{h=1}^{H} \xi_{t, h}\right| \geq 4 \sqrt{H^{3} K \log (4 / \delta)}\right) \leq \delta / 2$. Therefore, with probability at least $1-\delta / 2$, we have

$$
\begin{equation*}
\left|\sum_{t=1}^{K} \sum_{h=1}^{H} \xi_{t, h}\right| \leq \mathcal{O}\left(H^{3 / 2} \sqrt{K \imath}\right) \tag{D.3.6}
\end{equation*}
$$

Term (B): By application of Cauchy-Schwartz, we obtain

$$
\sum_{t=1}^{K} \sum_{h=1}^{H}\left\|\phi_{t, h}\right\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}} \leq \sqrt{K} \sum_{h=1}^{H} \sqrt{\sum_{t=1}^{K}\left\|\phi_{t, h}\right\|_{\Sigma_{t, h}^{-1} \tilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}^{2}}
$$

From Lemma 18, we have $\left\|\Sigma_{t}^{-1} \widetilde{\Sigma}_{t} \Sigma_{t}^{-1}\right\| \leq \frac{1}{\lambda}$, then, $\left\|\phi_{t, h}\right\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}} \leq$ $\frac{1}{\sqrt{\lambda}}\left\|\phi_{t, h}\right\|=\left\|\phi_{t, h}\right\| \leq 1$. So, we can use the bound on the sum of the squared norm of the features provided in proposition 4 of Russac et al. [2019] to obtain

$$
\begin{equation*}
\sum_{t=1}^{K} \sum_{h=1}^{H}\left\|\phi_{t, h}\right\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}} \leq H \sqrt{K} \sqrt{2 d K \log (1 / \eta)+2 d \log \left(1+\frac{1}{d \lambda(1-\eta)}\right)} \tag{D.3.7}
\end{equation*}
$$

Term (C): We control the bias term using the MDP variation budget as follows.

$$
\begin{align*}
& \sum_{t=1}^{K} \sum_{h=1}^{H} \operatorname{bias}(t, h) \leq \frac{4 H K \sqrt{d}}{\lambda} \frac{\eta^{W}}{1-\eta}+\sqrt{\frac{d}{\lambda(1-\eta)}} \\
& \sum_{t=1}^{K} \sum_{h=1}^{H} \sum_{s=t-W}^{t-1}\left\|\boldsymbol{\theta}_{s, h}-\boldsymbol{\theta}_{s+1, h}\right\|+\left\|\boldsymbol{\mu}_{s, h}(\mathcal{S})-\boldsymbol{\mu}_{s+1, h}(\mathcal{S})\right\| \\
& \leq \frac{4 H K \sqrt{d}}{\lambda} \frac{\eta^{W}}{1-\eta}+\sqrt{\frac{d}{\lambda(1-\eta)}} W \Delta \tag{D.3.8}
\end{align*}
$$

Finally, the desired regret bound in Theorem 4 is obtained by combining Equations (D.3.5), (D.3.6), (D.3.7) and (D.3.8).

## D. 4 Missing Proofs of Regret Analysis of OPT-WLSVI

## D.4.1 Linearity of $Q$-values: Lemma 8

Proof. The definition of non-stationary linear MDP from Assumption 6 together with the Bellman equation gives:

$$
\begin{aligned}
Q_{t, h}^{\pi} & =r_{t, h}(s, a)+\left[\mathbb{P}_{t, h} V_{t, h+1}^{\pi}\right](s, a) \\
& =\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\theta}_{t, h}+\int_{s^{\prime}} \boldsymbol{\phi}(s, a)^{\top} V_{t, h+1}^{\pi}\left(s^{\prime}\right) d \boldsymbol{\mu}_{t, h}\left(s^{\prime}\right) \\
& =\boldsymbol{\phi}(s, a)^{\top}\left(\boldsymbol{\theta}_{t, h}+\int_{s^{\prime}} V_{t, h+1}^{\pi}\left(s^{\prime}\right) d \boldsymbol{\mu}_{t, h}\left(s^{\prime}\right)\right)
\end{aligned}
$$

We define $\mathbf{w}_{t, h}^{\pi}$ to be the term inside the parentheses.

## D.4.2 Non-Stationarity Bias

## Proof of Lemma 9

## Reward Bias:

$$
\begin{aligned}
& \left|r_{t, h}(s, a)-\bar{r}_{t, h}(s, a)\right| \\
& \leq\left|\boldsymbol{\phi}(s, a)^{\top}\left(\boldsymbol{\theta}_{t, h}-\Sigma_{t, h}^{-1}\left(\sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top} \boldsymbol{\theta}_{\tau, h}+\lambda \eta^{-(t-1)} \boldsymbol{\theta}_{t, h}\right)\right)\right| \\
& =\left|\phi(s, a)^{\top} \sum_{\tau=1}^{t-1} \Sigma_{t, h}^{-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top}\left(\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right)\right| \\
& \leq \underbrace{\left|\phi(s, a)^{\top} \sum_{\tau=t-W}^{t-1} \Sigma_{t, h}^{-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top}\left(\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right)\right|}_{(\star)}+\underbrace{\left|\phi(s, a)^{\top} \sum_{\tau=1}^{t-W-1} \Sigma_{t, h}^{-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top}\left(\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right)\right|}_{(\star \star)}
\end{aligned}
$$

Bound on ( $\star$ ):

$$
\begin{aligned}
& \left|\boldsymbol{\phi}(s, a)^{\top} \sum_{\tau=t-W}^{t-1} \Sigma_{t, h}^{-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top}\left(\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right)\right| \\
& =\left|\sum_{\tau=t-W}^{t-1} \eta^{-\tau} \phi(s, a)^{\top} \Sigma_{t, h}^{-1} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top}\left(\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right)\right| \\
& \leq \sum_{\tau=t-W}^{t-1} \eta^{-\tau}\left|\phi(s, a)^{\top} \Sigma_{t, h}^{-1} \boldsymbol{\phi}_{\tau, h}\right| \cdot\left|\boldsymbol{\phi}_{\tau, h}^{\top}\left(\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right)\right| \\
& \leq \sum_{\tau=t-W}^{t-1} \eta^{-\tau}\left|\phi(s, a)^{\top} \Sigma_{t, h}^{-1} \boldsymbol{\phi}_{\tau, h}\right|\left\|\boldsymbol{\phi}_{\tau, h}\right\|\left\|\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right\| \\
& \leq \sum_{\tau=t-W}^{t-1} \eta^{-\tau}\left|\phi(s, a)^{\top} \Sigma_{t, h}^{-1} \boldsymbol{\phi}_{\tau, h}\right|\left\|\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right\| \\
& =\sum_{\tau=t-W}^{t-1} \eta^{-\tau}\left|\phi(s, a)^{\top} \Sigma_{t, h}^{-1} \boldsymbol{\phi}_{\tau, h}\right|\left\|\sum_{s=\tau}^{t-1} \boldsymbol{\theta}_{s, h}-\boldsymbol{\theta}_{s+1, h}\right\| \\
& \leq \sum_{\tau=t-W}^{t-1} \eta^{-\tau}\left|\phi(s, a)^{\top} \Sigma_{t, h}^{-1} \boldsymbol{\phi}_{\tau, h}\right| \sum_{s=\tau}^{t-1}\left\|\boldsymbol{\theta}_{s, h}-\boldsymbol{\theta}_{s+1, h}\right\| \\
& \left.\leq \sum_{s=t-W}^{t-1} \sum_{\tau=t-W}^{s} \eta^{-\tau}\left|\phi(s, a)^{\top} \Sigma_{t, h}^{-1} \boldsymbol{\phi}_{\tau, h}\right|\left\|\boldsymbol{\theta}_{s, h}-\boldsymbol{\theta}_{s+1, h}\right\| \leq 1\right) \\
& \leq \sum_{s=t-W}^{t-1} \sqrt{\left[\sum_{\tau=t-W}^{s} \eta^{-\tau} \phi(s, a)^{\top} \Sigma_{t, h}^{-1} \phi(s, a)\right] \cdot\left[\sum_{\tau=t-W}^{s-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h}^{\top} \Sigma_{t, h}^{-1} \boldsymbol{\phi}_{\tau, h}\right]} \\
& \leq \boldsymbol{\theta}_{s, h}-\boldsymbol{\theta}_{s+1, h} \| \\
& \leq \sum_{s=t-W}^{t-1} \sqrt{\left[\sum_{\tau=t-W}^{s} \eta^{-\tau} \phi(s, a)^{\top} \Sigma_{t, h}^{-1} \phi(s, a)\right] \cdot \sqrt{\boldsymbol{d}} \cdot\left\|\boldsymbol{\theta}_{s, h}-\boldsymbol{\theta}_{s+1, h}\right\|}
\end{aligned}
$$

(by Lemma 22)

$$
\leq \sum_{s=t-W}^{t-1} \sqrt{d} \sum_{s=t-W}^{t-1} \sqrt{\frac{\sum_{\tau=t-W}^{t-1} \eta^{-\tau}}{\lambda \eta^{-(t-1)}}}\left\|\boldsymbol{\theta}_{s, h}-\boldsymbol{\theta}_{s+1, h}\right\|
$$

$$
\left(\|\phi(s, a)\| \leq 1 \text { and } \lambda_{\max }\left(\Sigma_{t, h}^{-1}\right) \leq \frac{1}{\lambda \eta^{-(t-1)}}\right)
$$

$$
\leq \sqrt{\frac{d}{\lambda(1-\eta)}} \sum_{s=t-W}^{t-1}\left\|\boldsymbol{\theta}_{s, h}-\boldsymbol{\theta}_{s+1, h}\right\|
$$

Bound on ( $\star \star$ ):

$$
\begin{aligned}
& \left|\phi(s, a)^{\top} \sum_{\tau=1}^{t-W-1} \Sigma_{t, h}^{-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top}\left(\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right)\right| \\
& \leq \frac{1}{\lambda}\|\phi(s, a)\| \sum_{\tau=1}^{t-W-1} \eta^{t-\tau-1}\left\|\boldsymbol{\phi}_{\tau, h}\right\| \cdot\left|\boldsymbol{\phi}_{\tau, h}^{\top}\left(\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right)\right| \quad\left(\lambda_{\max }\left(\Sigma_{t, h}^{-1}\right) \leq \frac{1}{\lambda \eta^{-(t-1)}}\right) \\
& \leq \frac{1}{\lambda}\|\phi(s, a)\| \sum_{\tau=1}^{t-W-1} \eta^{t-\tau-1}\left\|\boldsymbol{\phi}_{\tau, h}\right\|^{2}\left\|\boldsymbol{\theta}_{t, h}-\boldsymbol{\theta}_{\tau, h}\right\| \\
& \leq \frac{2 \sqrt{d}}{\lambda} \frac{\eta^{W}}{1-\eta}
\end{aligned} \quad\left(\phi(s, a) \leq 1 \text { and }\left\|\boldsymbol{\theta}_{t, h}\right\| \leq \sqrt{d}\right)
$$

Transition Bias: $\forall f: \mathcal{S} \rightarrow \mathbb{R}$ such that $\|f\|_{\infty}<\infty$ (real-valued bounded function), similarly to what we have done for the reward function, we obtain

$$
\begin{aligned}
&\left|\left[\left(\mathbb{P}_{t, h}-\overline{\mathbb{P}}_{t, h}\right) f\right](s, a)\right| \leq\left\|\Sigma_{t, h}^{-1} \sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h} \boldsymbol{\phi}_{\tau, h}^{\top} \int f\left(s^{\prime}\right)\left(d \boldsymbol{\mu}_{t, h}\left(s^{\prime}\right)-d \boldsymbol{\mu}_{\tau, h}\left(s^{\prime}\right)\right)\right\| \\
&(\|\boldsymbol{\phi}(s, a)\| \leq 1) \\
& \leq \sqrt{\frac{d}{\lambda(1-\eta)} \sum_{s=t-W}^{t-1}\left\|\int f\left(s^{\prime}\right)\left(d \boldsymbol{\mu}_{s, h}\left(s^{\prime}\right)-d \boldsymbol{\mu}_{s+1, h}\left(s^{\prime}\right)\right)\right\|} \\
&+\frac{1}{\lambda} \sum_{\tau=1}^{t-W-1} \eta^{t-\tau-1}\left\|\int f\left(s^{\prime}\right)\left(d \boldsymbol{\mu}_{t, h}\left(s^{\prime}\right)-d \boldsymbol{\mu}_{\tau, h}\left(s^{\prime}\right)\right)\right\|
\end{aligned}
$$

Furthermore,

$$
\begin{aligned}
\left\|\int f\left(s^{\prime}\right)\left(d \boldsymbol{\mu}_{t, h}\left(s^{\prime}\right)-d \boldsymbol{\mu}_{\tau, h}\left(s^{\prime}\right)\right)\right\| & =\sqrt{\sum_{l=1}^{d}\left|\int f\left(s^{\prime}\right)\left(d \boldsymbol{\mu}_{t, h}^{(l)}\left(s^{\prime}\right)-d \boldsymbol{\mu}_{\tau, h}^{(l)}\left(s^{\prime}\right)\right)\right|^{2}} \\
& \leq\|f\|_{\infty} \sqrt{\left.\sum_{l=1}^{d} \mid \boldsymbol{\mu}_{t, h}^{(l)}(\mathcal{S})-\boldsymbol{\mu}_{\tau, h}^{(l)}(\mathcal{S})\right)\left.\right|^{2}} \\
& =\|f\|_{\infty}\left\|\boldsymbol{\mu}_{s, h}(\mathcal{S})-\boldsymbol{\mu}_{s+1, h}(\mathcal{S})\right\| \\
& \leq 2 \sqrt{d}\|f\|_{\infty}
\end{aligned}
$$

Therefore,

$$
\left|\left[\left(\mathbb{P}_{t, h}-\overline{\mathbb{P}}_{t, h}\right) f\right](s, a)\right| \leq\|f\|_{\infty}\left(\sqrt{\frac{d}{\lambda(1-\eta)}} \sum_{s=t-W}^{t-1}\left\|\left(\boldsymbol{\mu}_{s, h}(\mathcal{S})-\boldsymbol{\mu}_{s+1, h}(\mathcal{S})\right)\right\|+\frac{2 \sqrt{d}}{\lambda} \frac{\eta^{W}}{1-\eta}\right)
$$

## D.4.3 Single Step Error Decomposition

We provide here the full derivation of the single-error decomposition. We have for all $(t, h) \in[K] \times[H]$

$$
\begin{aligned}
\boldsymbol{\phi}(s, a)^{\top} \mathbf{w}_{t, h}-Q_{t, h}^{\pi}(s, a)= & \underbrace{\left(\bar{r}_{t, h}-r_{t, h}\right)(s, a)}_{\text {reward bias }}+\underbrace{\left(\widehat{r}_{t, h}-\bar{r}_{t, h}\right)(s, a)}_{\text {reward variance }}+ \\
& \underbrace{\left[\left(\bar{P}_{t, h}-P_{t, h}\right) V_{t, h+1}^{\pi}\right](s, a)}_{\text {transition bias }}+\underbrace{\left[\left(\widehat{P}_{t, h}-\bar{P}_{t, h}\right) V_{t, h+1}\right](s, a)}_{\text {transition variance }}+ \\
& \underbrace{\left[\bar{P}_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)}_{\text {difference in value functions of next step }} .
\end{aligned}
$$

Reward \& transition bias: Thanks to Lemma 9, we have

$$
\begin{aligned}
\left|\bar{r}_{t, h}(s, a)-r_{t, h}(s, a)\right| & \leq \operatorname{bias}_{r}(t, h), \\
\left|\left[\left(\bar{P}_{t, h}-P_{t, h}\right) V_{t, h+1}^{\pi}\right](s, a)\right| & \leq \operatorname{bias}_{P}(t, h) .
\end{aligned} \quad\left(\left\|V_{t, h+1}^{\pi}\right\|_{\infty} \leq H\right)
$$

## Difference in value functions of next step:

$$
\begin{aligned}
{\left[\bar{P}_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a) } & =\left[P_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)+\left[\left(\bar{P}_{t, h}-P_{t, h}\right)\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a) \\
\leq & {\left[P_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)+\operatorname{bias}_{P}(s, a) . } \\
& \quad\left(\left\|V_{t, h+1}-V_{t, h+1}^{\pi}\right\|_{\infty} \leq H\right)
\end{aligned}
$$

Reward variance: The reward variance here reduces simply to the bias due to the regularization as we assume that $r$ is a deterministic function.

$$
\begin{aligned}
\left|\left(\widehat{r}_{t, h}-\bar{r}_{t, h}\right)(s, a)\right| & =\lambda \eta^{-(t-1)}\left|\left\langle\boldsymbol{\phi}(s, a), \Sigma_{t, h}^{-1} \boldsymbol{\theta}_{t, h}\right\rangle\right| \\
& \leq \lambda \eta^{-(t-1)}\|\boldsymbol{\phi}(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}\left\|\Sigma_{t, h}^{-1} \boldsymbol{\theta}_{t, h}\right\|_{\Sigma_{t, h} \widetilde{\Sigma}_{t, h}^{-1} \Sigma_{t, h}} \\
& =\lambda \eta^{-(t-1)}\|\boldsymbol{\phi}(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}\left\|\boldsymbol{\theta}_{t, h}\right\|_{\widetilde{\Sigma}_{t, h}^{-1}} \\
& \leq \lambda \eta^{-(t-1)}\|\boldsymbol{\phi}(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}} \sqrt{\left\|\widetilde{\Sigma}_{t, h}^{-1}\right\|}\left\|\boldsymbol{\theta}_{t, h}\right\| \\
& \leq \sqrt{d \lambda}\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}
\end{aligned}
$$

The last step follows from $\left\|\boldsymbol{\theta}_{t, h}\right\| \leq \sqrt{d}$ (Assumption 6) and $\left\|\widetilde{\Sigma}_{t, h}^{-1}\right\| \leq \frac{1}{\lambda \eta^{-2(t-1)}}$.

If we define bias $\triangleq$ bias $_{r}+2 \cdot$ bias $_{P}$ the total non-stationarity bias of the MDP, we can summarize the one-step analysis as follows:

$$
\begin{aligned}
& \phi(s, a)^{\top} \mathbf{w}_{t, h}-Q_{t, h}^{\pi}(s, a) \leq \operatorname{bias}(t, h)+\left[\mathbb{P}_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)+ \\
& \sqrt{d \lambda}\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}+\left[\left(\widehat{P}_{t, h}-\bar{P}_{t, h}\right) V_{t, h}\right](s, a)
\end{aligned}
$$

## D.4.4 Boundness of iterates

We will start with the following elementary lemma:
Lemma 22. Let $\Sigma_{t}=\sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau} \boldsymbol{\phi}_{\tau}^{\top}+\lambda \eta^{-(t-1)} \mathbf{I}$ where $\phi_{\tau} \in \mathbb{R}^{d}$ and $\lambda>0, \eta \in$ $(0,1)$. Then:

$$
\sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau}^{\top} \Sigma_{t}^{-1} \phi_{\tau} \leq d
$$

Proof. We have $\sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau}^{\top} \Sigma_{t}^{-1} \boldsymbol{\phi}_{\tau}=\sum_{\tau=1}^{t-1} \operatorname{tr}\left(\eta^{-\tau} \boldsymbol{\phi}_{\tau}^{\top} \Sigma_{t}^{-1} \boldsymbol{\phi}_{\tau}\right)=$ $\operatorname{tr}\left(\Sigma_{t}^{-1} \sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau} \boldsymbol{\phi}_{\tau}^{\top}\right)$. Given the eigenvalue decomposition $\sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau} \boldsymbol{\phi}_{\tau}=$ $\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{d}\right)^{\top}$, we have $\Sigma_{t}=\operatorname{diag}\left(\lambda_{1}+\lambda \eta^{-(t-1)}, \ldots, \lambda_{d}+\lambda \eta^{-(t-1)}\right)^{\top}$, and $\operatorname{tr}\left(\Sigma_{t}^{-1} \sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau} \phi_{\tau}^{\top}\right)=\sum_{i=1}^{d} \frac{\lambda_{j}}{\lambda_{j}+\lambda \eta^{-(t-1)}} \leq d$

## Proof of Lemma 18

$\underline{\text { Bound on }\left\|\mathbf{w}_{t, h}\right\| \text { : For any vector } v \in \mathbb{R}^{d} \text {, we have }}$

$$
\begin{aligned}
\left|v^{\top} \mathbf{w}_{t, h}\right| & =\left|v^{\top} \Sigma_{t, h}^{-1} \sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau, h}\left[r_{\tau, h}+\max _{a} Q_{\tau, h+1}\left(s^{\tau, h+1}, a\right)\right]\right| \\
& \leq \sum_{\tau=1}^{t-1} \eta^{-\tau}\left|v^{\top} \Sigma_{t, h}^{-1} \phi_{\tau, h}\right| \cdot 2 H \leq \sqrt{\left[\sum_{\tau=1}^{t-1} \eta^{-\tau} v^{\top} \Sigma_{t, h}^{-1} v\right] \cdot\left[\sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau, h}^{\top} \Sigma_{t, h}^{-1} \phi_{\tau, h}\right] \cdot 2 H} \\
& \leq 2 H\|v\| \sqrt{\frac{\sum_{\tau=1}^{t-1} \eta^{-\tau}}{\lambda \eta^{-(t-1)}} \cdot d}=2 H\|v\| \sqrt{\frac{d\left(1-\eta^{t-1}\right)}{\lambda(1-\eta)}}
\end{aligned}
$$

where the third inequality is due to Lemma 22 and the fact that the eigenvalues of $\Sigma_{t, h}^{-1}$ are upper bounded by $\frac{1}{\lambda \eta^{-(t-1)}}$. The remainder of the proof follows from the fact that $\left\|\mathbf{w}_{t, h}\right\|=\max _{v:\|v\|=1}\left|v^{\top} \mathbf{w}_{t, h}\right|$.

Bound on $\left\|\Sigma_{t}^{-1} \widetilde{\Sigma}_{t} \Sigma_{t}^{-1}\right\|:$
$\widetilde{\Sigma}_{t}=\sum_{\tau=1}^{t-1} \eta^{-2 \tau} \boldsymbol{\phi}_{\tau} \boldsymbol{\phi}_{\tau}^{\top}+\lambda \eta^{-2(t-1)} \mathbf{I} \leq \eta^{-(t-1)} \sum_{\tau=1}^{t-1} \eta^{-\tau} \boldsymbol{\phi}_{\tau} \boldsymbol{\phi}_{\tau}^{\top}+\lambda \eta^{-2(t-1)} \mathbf{I}=\eta^{-(t-1)} \Sigma_{t}$

Hence,

$$
\Sigma_{t}^{-1} \widetilde{\Sigma}_{t} \Sigma_{t}^{-1} \leq \eta^{-(t-1)} \Sigma_{t}^{-1} \Sigma_{t} \Sigma_{t}^{-1}=\eta^{-(t-1)} \Sigma_{t}^{-1}
$$

and

$$
\left\|\Sigma_{t}^{-1} \widetilde{\Sigma}_{t} \Sigma_{t}^{-1}\right\| \leq \eta^{-(t-1)}\left\|\Sigma_{t}^{-1}\right\| \leq \eta^{-(t-1)} \frac{1}{\lambda \eta^{-(t-1)}}=\frac{1}{\lambda}
$$

## D.4.5 Transition Concentration

$$
\begin{align*}
&\left|\left[\left(\widehat{P}_{t, h}-\bar{P}_{t, h}\right) V_{t, h}\right](s, a)\right| \\
& \leq \mid \phi(s, a)^{\top}\left(\Sigma_{t, h}^{-1} \sum_{\tau=1}^{k-1} \eta^{-\tau} \phi_{\tau, h}\left(V_{t, h+1}\left(s_{\tau, h+1}\right)-\left[\mathbb{P}_{t, h} V_{t, h+1}\right]\left(s_{\tau, h} a_{\tau, h}\right)\right)\right. \\
&\left.-\lambda \eta^{-(t-1)} \Sigma_{t, h}^{-1} \boldsymbol{\mu}_{t, h} V_{t, h+1}\right) \mid \\
& \leq\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{t}_{t, h} \Sigma_{t, h}^{-1}}\left(\left\|\sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau, h}\left(V_{t, h+1}\left(s_{\tau, h+1}\right)-\left[\mathbb{P}_{t, h} V_{t, h+1}\right]\left(s_{\tau, h}, a_{\tau, h}\right)\right)\right\|_{\widetilde{\Sigma}_{t, h}^{-1}}\right. \\
&\left.+\lambda \eta^{-(t-1)} \sqrt{\left\|\widetilde{\Sigma}_{t, h}^{-1}\right\|}\left\|\boldsymbol{\mu}_{t, h}(\mathcal{S})\right\|\left\|V_{t, h+1}\right\|_{\infty}\right) \quad(\text { Cauchy-Schwarz }) \\
& \leq\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}\left(\left\|\sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau, h}\left(V_{t, h+1}\left(s_{\tau, h+1}\right)-\left[\mathbb{P}_{t, h} V_{t, h+1}\right]\left(s_{\tau, h}, a_{\tau, h}\right)\right)\right\|_{\widetilde{\Sigma}_{t, h}^{-1}}+H \sqrt{d \lambda}\right) \tag{D.4.2}
\end{align*}
$$

The last step follows from $\left\|\boldsymbol{\mu}_{t, h}(\mathcal{S})\right\| \leq \sqrt{d}$ (Assumption 6) and $\left\|\widetilde{\Sigma}_{t, h}^{-1}\right\| \leq$ $\frac{1}{\lambda \eta^{-2(t-1)}}$.

Let us now consider the following function form:

$$
\begin{equation*}
V^{\mathbf{w}, \mathbf{A}}(\cdot)=\min \left\{\max _{a \in \mathcal{A}}\left\{\mathbf{w}^{\top} \boldsymbol{\phi}(\cdot, a)+\sqrt{\boldsymbol{\phi}(\cdot, a)^{\top} \mathbf{A} \boldsymbol{\phi}(\cdot, a)}\right\}, H\right\} \tag{D.4.3}
\end{equation*}
$$

where $\mathbf{w} \in \mathbb{R}^{d}$ and $\mathbf{A} \in \mathbb{R}^{d \times d}$ is a symmetric definite positive matrix that are in

$$
\begin{equation*}
\mathcal{G}=\left\{\mathbf{w}, \mathbf{A}:\|\mathbf{w}\| \leq 2 H \sqrt{\frac{d}{\lambda(1-\eta)}} \text { and }\|\mathbf{A}\|_{\mathrm{F}} \leq \frac{\sqrt{d} \beta^{2}}{\lambda}\right\} \tag{D.4.4}
\end{equation*}
$$

In the technical Lemma 26, we prove a concentration bound that holds uniformly for any function on the form $V^{\mathbf{w}, \mathbf{A}}$ where $\mathbf{w}, \mathbf{A} \in \mathcal{G}$. The statement and full proof of this lemma is defered to section D. 5 of the appendix. As a corollary of Lemma 26, we can prove the concentration of the transition as follows.

For any $\tau>0, h \in[H]$, let $\mathcal{F}_{\tau, h}$ be the $\sigma$-field generated by all the random variables until episode $\tau$, step $h$. $\left\{s_{\tau, h}\right\}$ defines a stochastic process on state space $\mathcal{S}$ with corresponding filtration $\left\{\mathcal{F}_{\tau, h}\right\}$. We have

$$
\begin{aligned}
V_{\tau, h+1}(\cdot) & =\max _{a \in \mathcal{A}}\left\{\min \left\{\mathbf{w}_{t, h+1}^{\top} \boldsymbol{\phi}(\cdot, a)+\beta_{t}\left[\boldsymbol{\phi}(\cdot, a)^{\top} \Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1} \boldsymbol{\phi}(\cdot, a)\right]^{1 / 2}, H\right\}\right\} \\
& =V^{\mathbf{w}, \mathbf{A}}(\cdot)
\end{aligned}
$$

where $\mathbf{w}=\mathbf{w}_{t, h+1}$ and $\mathbf{A}=\beta_{t}^{2} \Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}$. We have $\|\mathbf{w}\| \leq 2 H \sqrt{\frac{d\left(1-\eta^{t-1}\right)}{\lambda(1-\eta)}}$ and $\|\mathbf{A}\|_{F} \leq \sqrt{d} \beta_{t}^{2}\|A\| \leq \frac{\sqrt{d} \beta_{t}^{2}}{\lambda}$ by Lemma 18. Therefore $(\mathbf{w}, \mathbf{A}) \in \mathcal{G}$ and we can apply Lemma 26: we have with probability at least $1-\delta / 2$

$$
\begin{align*}
& \left\|\sum_{\tau=1}^{t-1} \eta^{-\tau} \phi_{\tau, h}\left(V_{t, h+1}\left(s_{\tau, h+1}\right)-\left[\mathbb{P}_{t, h} V_{t, h+1}\right]\left(s_{\tau, h}, a_{\tau, h}\right)\right)\right\|_{\widetilde{\Sigma}_{t, h}^{-1}} \\
& \leq C d H \sqrt{\log \left(d H \beta \frac{1}{\lambda(1-\eta)}\right)+\log (2 / \delta)} \tag{D.4.5}
\end{align*}
$$

Combining Equation (D.4.2) and (D.4.5), we obtain that with probability at least $1-\delta / 2$ :

$$
\begin{aligned}
& \left|\left[\left(\widehat{P}_{t, h}-\bar{P}_{t, h}\right) V_{t, h}\right](s, a)\right| \\
& \leq\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}\left(C d H \sqrt{\log \left(d H \beta \frac{1}{\lambda(1-\eta)}\right)+\log (2 / \delta)}+H \sqrt{d \lambda}\right)
\end{aligned}
$$

## D.4.6 Single-Step High Probability Upper Bound

## Proof of Lemma 20

We have shown so far:

$$
\begin{aligned}
\phi(s, a)^{\top} \mathbf{w}_{t, h} & -Q_{t, h}^{\pi}(s, a) \leq \operatorname{bias}(t, h)+\left[\mathbb{P}_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a) \\
& +\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}\left(C d H \sqrt{\log \left(d H \beta \frac{1}{\lambda(1-\eta)}\right)+\log (2 / \delta)}+H \sqrt{d \lambda}+\sqrt{d \lambda}\right)
\end{aligned}
$$

so there exists an absolute constant $C^{\prime}>0$ such that

$$
\begin{aligned}
\phi(s, a)^{\top} \mathbf{w}_{t, h} & -Q_{t, h}^{\pi}(s, a) \leq \operatorname{bias}(t, h)+\left[\mathbb{P}_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a) \\
& +\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}\left(C^{\prime} d H \sqrt{\lambda} \sqrt{\log \left(d H \beta \frac{1}{\lambda(1-\eta)}\right)+\log (2 / \delta)}\right)
\end{aligned}
$$

The missing ingredient to prove our key lemma is the choice of the parameter $\beta$
Now, we would like to find an appropriate choice of $\beta$ such that

$$
\begin{equation*}
C^{\prime} d H \sqrt{\lambda} \sqrt{\log \left(d H \beta \frac{1}{\lambda(1-\eta)}\right)+\log (2 / \delta)} \leq \beta \tag{D.4.6}
\end{equation*}
$$

First we set $\lambda=1$. A good candidate for $\beta$ is in the form of $\beta=c_{\beta} d H \sqrt{2}$ where $c_{\beta}>1$ is an absolute constant and $\iota=\log \left(\frac{2 d H}{\delta(1-\eta)}\right)$. With this choice, we obtain:

$$
\begin{aligned}
C^{\prime} d H \sqrt{\lambda} \sqrt{\log \left(d H \beta \frac{1}{\lambda(1-\eta)}\right)+\log (2 / \delta)} & =C^{\prime} d H \sqrt{\log \left(c_{\beta} d H \sqrt{\imath}\right)+\imath} \\
& \leq C^{\prime \prime} d H \sqrt{\log \left(c_{\beta}\right)+\imath} \\
& \left(C^{\prime \prime}>0 \text { is an absolute constant }\right) \\
& \leq C^{\prime \prime} d H\left(\sqrt{\log \left(c_{\beta}\right)}+\sqrt{\imath}\right)
\end{aligned}
$$

Let $c_{\beta}>1$ such that $C^{\prime \prime}\left(\sqrt{\log \left(c_{\beta}\right)}+\sqrt{\log (2)}\right) \leq \frac{c_{\beta}}{\sqrt{2}} \sqrt{\log (2)}$. In particular we have necessarily $\frac{c_{\beta}}{\sqrt{2}} \geq C^{\prime \prime}$ Therefore, we have:

$$
\begin{aligned}
C^{\prime \prime}\left(\sqrt{\log \left(c_{\beta}\right)}+\sqrt{\imath}\right) & =C^{\prime \prime}\left(\sqrt{\log \left(c_{\beta}\right)}+\sqrt{\log (2)+(\imath-\log (2)}\right) \quad(1 \geq \log (2)) \\
& \leq C^{\prime \prime}\left(\sqrt{\log \left(c_{\beta}\right)}+\sqrt{\log (2)}+\sqrt{(\imath-\log (2)}\right) \\
& \leq \frac{c_{\beta}}{\sqrt{2}} \sqrt{\log (2)}+C^{\prime \prime} \sqrt{(1-\log (2)} \\
& \leq \frac{c_{\beta}}{\sqrt{2}}(\sqrt{\log (2)}+\sqrt{(1-\log (2)}) \\
& \leq \frac{c_{\beta}}{\sqrt{2}} \sqrt{2} \sqrt{\log (2)+\imath-\log (2)} \\
& \leq c_{\beta} \sqrt{\imath}
\end{aligned}
$$

Therefore, with this choice of $c_{\beta}$ and $\beta=c_{\beta} d H \sqrt{2}$, we obtain that

$$
\left|\left\langle\phi(s, a), \mathbf{w}_{t, h}\right\rangle-Q_{t, h}^{\pi}(s, a)-\left[\mathbb{P}_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\pi}\right)\right](s, a)\right| \leq \operatorname{bias}(t, h)+\beta\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{h}_{t, h} \Sigma_{t, h}^{-1}} .
$$

## D.4.7 Optimism

## Proof of Lemma 21

We proceed by induction. By definition, we have $Q_{t, H+1}=Q_{t, H+1}^{\star}=0$ and the desired statement trivially holds at step $H+1$. Now, assume that the
statement holds for $h+1$. Consider step $h$. By Lemma 20, we have

$$
\left|\phi(s, a)^{\top} \mathbf{w}_{t, h}-Q_{t, h}^{\star}(s, a)-\left[\mathbb{P}_{t, h}\left(V_{t, h+1}-V_{t, h+1}^{\star}\right)\right](s, a)\right| \leq \boldsymbol{\operatorname { b i a s }}(t, h)+\beta\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}
$$

Moreover, we have

$$
\begin{aligned}
V_{t, h+1}^{\star}(s)-V_{t, h+1}(s) & =\max _{a \in \mathcal{A}} Q_{t, h+1}^{\star}(s, a)-\max _{a \in \mathcal{A}} Q_{t, h+1}(s, a) \\
& \leq \max _{a \in \mathcal{A}}\left(Q_{t, h+1}^{\star}(s, a)-Q_{t, h+1}(s, a)\right) \\
& \leq \sum_{h^{\prime}=h+1}^{H} \operatorname{bias}(t, h) \quad \text { (by the induction hypothesis) }
\end{aligned}
$$

Therefore, we obtain

$$
\begin{aligned}
Q_{t, h}^{\star}(s, a) & \leq \phi(s, a)^{\top} \mathbf{w}_{t, h}+\beta\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}+\left[\mathbb{P}_{t, h}\left(V_{t, h+1}^{\star}-V_{t, h+1}\right)\right](s, a)+\boldsymbol{b i a s}(t, h) \\
& \leq \phi(s, a)^{\top} \mathbf{w}_{t, h}+\beta\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}+\sum_{h^{\prime}=h}^{H} \boldsymbol{\operatorname { b i a s }}(t, h)
\end{aligned}
$$

We have

$$
\begin{aligned}
Q_{t, h}^{\star}(s, a) & \left.\leq \boldsymbol{\phi}(s, a)^{\top} \mathbf{w}_{t, h}\right\rangle++\beta\|\phi(s, a)\|_{\Sigma_{t, h}^{-1} \widetilde{\Sigma}_{t, h} \Sigma_{t, h}^{-1}}+\sum_{h^{\prime}=h}^{H} \operatorname{bias}(t, h) \\
& =Q_{t, h}(s, a)+\sum_{h^{\prime}=h}^{H} \operatorname{bias}(t, h)
\end{aligned}
$$

## D. 5 Technical Lemmas

Lemma 23 (Concentration of weighted self-normalized processes [Russac et al., 2019]). Let $\left\{\epsilon_{t}\right\}_{t=1}^{\infty}$ be a real-valued stochastic process with corresponding filtration $\{\mathcal{F}\}_{t=1}^{\infty}$. Let $\epsilon_{t} \mid \mathcal{F}_{t-1}$ be zero-mean and $\sigma$-subGaussiasn; i.e $\mathbb{E}\left[\epsilon_{t} \mid \mathcal{F}_{t-1}\right]=0$ and

$$
\forall \lambda \in \mathbb{R}, \quad \mathbb{E}\left[e^{\lambda \epsilon_{t}} \mid \mathcal{F}_{t-1}\right] \leq e^{\lambda^{2} \sigma^{2} / 2}
$$

Let $\left\{\phi_{t}\right\}_{t=1}^{\infty}$ be a predictable $\mathbb{R}^{d}$-valued stochastic process (i.e $\phi_{t}$ is $\mathcal{F}_{t-1^{-}}$ measurable) and $\left\{\omega_{t}\right\}_{t=0}^{\infty}$ be a sequence of predictable and positive weights. Let $\widetilde{\Sigma}_{t}=\sum_{s=1}^{t} \omega_{s}^{2} \boldsymbol{\phi}_{s} \boldsymbol{\phi}_{s}^{\top}+\mu_{t} \cdot \mathbf{I}$ where $\left\{\mu_{t}\right\}_{t=1}^{\infty}$ a deterministic sequence of scalars. Then for any $\delta>0$, with probability at least $1-\delta$, we have for all $t \geq 0$ :

$$
\begin{equation*}
\left\|\sum_{s=1}^{t} \omega_{s} \phi_{s} \epsilon_{s}\right\|_{\tilde{\Sigma}_{t}^{-1}} \leq \sigma \sqrt{2 \log \left(\frac{1}{\delta}\right)+\log \left(\frac{\operatorname{det}\left(\widetilde{\Sigma}_{t}\right)}{\mu_{t}^{d}}\right)} \tag{D.5.1}
\end{equation*}
$$

Lemma 24 (Determinant inequality for the weighted Gram matrix Russac et al. [2019]). Let $\left\{\lambda_{t}\right\}_{t=0}^{\infty}$ and $\left\{\omega_{t}\right\}_{t=0}^{\infty}$ be a deterministic sequence of scalars. Let $\Sigma_{t}=$ $\sum_{t=1}^{t} w_{s} \boldsymbol{\phi}_{s} \boldsymbol{\phi}_{s}^{\top}+\lambda_{t} \cdot \mathbf{I}$ be the weighted Gram matrix. Under the assumption $\forall t,\left\|\boldsymbol{\phi}_{t}\right\| \leq$ 1 , the following holds

$$
\begin{equation*}
\operatorname{det}\left(\Sigma_{t}\right) \leq\left(\lambda_{t}+\frac{\sum_{s=1}^{t} \omega_{s}}{d}\right)^{d} \tag{D.5.2}
\end{equation*}
$$

Lemma 25 (Covering Number of Euclidean Ball [Pollard, 1990]). For any $\epsilon>0$, the $\epsilon$-covering number of the Euclidean ball in $\mathbb{R}^{d}$ with radius $R>0$ is upper bounded by $\left(\frac{3 R}{\epsilon}\right)^{d}$

Lemma 26 (Uniform concentration). Let $\left\{s_{t}\right\}_{t=1}^{\infty}$ be a stochastic process on state state $\mathcal{S}$ with corresponding filtration $\left\{\mathcal{F}_{t}\right\}_{t=0}^{\infty}$. Let $\left\{\phi_{t}\right\}_{t=1}^{\infty}$ be an $\mathbb{R}^{d}$-valued stochastic process where $\phi_{t}$ is $\mathcal{F}_{t-1}$-measurable, and $\|\phi\| \leq 1$. Let $\widetilde{\Sigma}_{t}=\sum_{\tau=1}^{t} \eta^{-2 \tau} \phi_{\tau} \phi_{\tau}^{\top}+\lambda \eta^{-2 t}$. I. Then for any $\epsilon \in(0,1)$ and $\delta>0$, with probability at least $1-\delta$, for all $t>0$ and for all $\mathbf{w}, \mathbf{A} \in \mathcal{G}$ defined in (D.4.4), we have

$$
\begin{align*}
& \left\|\sum_{\tau=1}^{t} \eta^{-\tau} \phi_{\tau}\left(V^{\mathbf{w}, \mathbf{A}}\left(s_{\tau}\right)-\mathbb{E}\left[V^{\mathbf{w}, \mathbf{A}}\left(s_{\tau}\right) \mid \mathcal{F}_{\tau-1}\right]\right)\right\|_{\tilde{\Sigma}_{t}^{-1}} \\
& \leq \mathcal{O}\left(d H \sqrt{\log \left(d H \beta_{t} \frac{1}{\lambda(1-\eta)}\right)+\log (1 / \delta)}\right) \tag{D.5.3}
\end{align*}
$$

where $V^{\mathbf{w}, \mathbf{A}}$ is defined in Equation (D.4.3).
Proof. Let $t>0$. For $\mathbf{w}, \mathbf{A} \in \mathcal{O}_{t}$ and $\tau \in[t]$ define

$$
\begin{equation*}
\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}=V^{\mathbf{w}, \mathbf{A}}\left(s_{\tau}\right)-\mathbb{E}\left[V^{\mathbf{w}, \mathbf{A}}\left(s_{\tau}\right) \mid \mathcal{F}_{\tau-1}\right] \tag{D.5.4}
\end{equation*}
$$

Then $\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}$ defines a martingale difference sequence with filtration $\mathcal{F}_{\tau}$. Moreover, by the definition of $V^{\mathbf{w}, \mathbf{A}}$, each $\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}$ is bounded in absolute value by $H$, so that each $\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}$ is $H$-subgaussian random variable.

So, by lemma 23 , the $\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}$ induce a self normalizing process so that for any $\delta>0$, with probability at least $1-\delta$, we have for all $t>0$ :

$$
\begin{align*}
\left\|\sum_{\tau=1}^{t} \eta^{-\tau} \phi_{s} \epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}\right\|_{\tilde{\Sigma}_{t}^{-1}} & \leq H \sqrt{2 \log \left(\frac{1}{\delta}\right)+\log \left(\frac{\operatorname{det}\left(\widetilde{\Sigma}_{t}\right)}{\left(\lambda \eta^{-2 t}\right)^{d}}\right)}  \tag{D.5.5}\\
& \leq H \sqrt{2 \log \left(\frac{1}{\delta}\right)+d \log \left(1+\frac{1-\eta^{2 t}}{\lambda d\left(1-\eta^{2}\right)}\right)} \tag{D.5.6}
\end{align*}
$$

The last step is due to $\operatorname{det}\left(\widetilde{\Sigma}_{t}\right) \leq\left(\lambda \eta^{-2 t}+\frac{\eta^{-2 t}-1}{d\left(1-\eta^{-d}\right)}\right)^{d}$ by lemma 24 .

Let $\mathcal{N}_{\epsilon}(\mathcal{G})$ be covering number of $\mathcal{G}$. So, by union bound, with probability $\delta$. For all $\tilde{\mathbf{w}}, \tilde{\mathbf{A}}$ in the $\epsilon$-covering of $\mathcal{G}$ that

$$
\begin{equation*}
\left\|\sum_{\tau=1}^{t} \eta^{-\tau} \phi_{s} \epsilon_{\tau}^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\right\|_{\tilde{\Sigma}_{t}^{-1}} \leq H \sqrt{2 \log \left(\frac{\mathcal{N}_{\epsilon}(\mathcal{G})}{\delta}\right)+d \log \left(1+\frac{1-\eta^{2 t}}{\lambda d\left(1-\eta^{2}\right)}\right)} \tag{D.5.7}
\end{equation*}
$$

For any $(\mathbf{w}, \mathbf{A}) \in \mathcal{G}$, we choose a specific $(\tilde{\mathbf{w}}, \tilde{\mathbf{A}})$ in the $\epsilon$-covering of $\mathcal{G}$ such that $\|\mathbf{w}-\tilde{\mathbf{w}}\| \leq \epsilon$ and $\|\mathbf{A}-\tilde{\mathbf{A}}\|_{\mathrm{F}} \leq \epsilon$.

$$
\begin{align*}
\left\|\sum_{\tau=1}^{t} \eta^{-\tau} \phi_{s} \epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}\right\|_{\tilde{\Sigma}_{t}^{-1}} & \leq\left\|\sum_{\tau=1}^{t} \eta^{-\tau} \boldsymbol{\phi}_{s} \epsilon_{\tau}^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\right\|_{\tilde{\Sigma}_{t}^{-1}}+\left\|\sum_{\tau=1}^{t} \eta^{-\tau} \boldsymbol{\phi}_{s}\left(\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}-\epsilon_{\tau}^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\right)\right\|_{\tilde{\Sigma}_{t}^{-1}} \\
& \leq H \sqrt{2 \log \left(\frac{\mathcal{N}_{\epsilon}(\mathcal{G})}{\delta}\right)+d \log \left(1+\frac{1-\eta^{2 t}}{\lambda d\left(1-\eta^{2}\right)}\right)} \\
& +\left\|\sum_{\tau=1}^{t} \eta^{-\tau} \boldsymbol{\phi}_{s}\left(\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}-\epsilon_{\tau}^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\right)\right\|_{\tilde{\Sigma}_{t}^{-1}} \tag{D.5.8}
\end{align*}
$$

We can bound

$$
\begin{aligned}
\left\|\sum_{\tau=1}^{t} \eta^{-\tau} \phi_{s}\left(\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}-\epsilon_{\tau}^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\right)\right\|_{\widetilde{\Sigma}_{t}^{-1}} & \leq \frac{1}{\sqrt{\lambda} \eta^{-t}}\left\|\sum_{\tau=1}^{t} \eta^{-\tau} \boldsymbol{\phi}_{s}\left(\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}-\epsilon_{\tau}^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\right)\right\| \\
& \leq \frac{1}{\sqrt{\lambda} \eta^{-t}} \cdot \frac{\eta^{-t}-1}{1-\eta} \sup _{\tau}\left|\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}-\epsilon_{\tau}^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\right| \\
& =\frac{1-\eta^{t}}{\sqrt{\lambda}(1-\eta)} \sup _{\tau}\left|\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}-\epsilon_{\tau}^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\right| \\
& \leq \frac{2\left(1-\eta^{t}\right)}{\sqrt{\lambda}(1-\eta)} \sup _{\tau}\left|V^{\mathbf{w}, \mathbf{A}}\left(s_{\tau}\right)-V^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\left(s_{\tau}\right)\right|
\end{aligned}
$$

By the definition of $V^{\mathbf{w}, \mathbf{A}}$, we have

$$
\begin{aligned}
& \sup _{\tau}\left|V^{\mathbf{w}, \mathbf{A}}\left(s_{\tau}\right)-V^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\left(s_{\tau}\right)\right| \\
& \leq \sup _{s, a}\left|\left(\mathbf{w}^{\top} \phi(s, a)+\sqrt{\phi(s, a)^{\top} \mathbf{A} \phi(s, a)}\right)-\left(\tilde{\mathbf{w}}^{\top} \phi(s, a)+\sqrt{\phi(s, a)^{\top} \tilde{\mathbf{A}} \phi(s, a)}\right)\right| \\
& \leq \sup _{\phi \in \mathbb{R}^{d}:\|\phi\| \leq 1}\left|\left(\mathbf{w}^{\top} \boldsymbol{\phi}+\sqrt{\boldsymbol{\phi}^{\top} \mathbf{A} \boldsymbol{\phi}}\right)-\left(\tilde{\mathbf{w}}^{\top} \boldsymbol{\phi}+\sqrt{\phi^{\top} \tilde{\mathbf{A}} \boldsymbol{\phi}}\right)\right| \\
& \leq \sup _{\phi \in \mathbb{R}^{d}:\|\phi\| \leq 1}\left|(\mathbf{w}-\tilde{\mathbf{w}})^{\top} \boldsymbol{\phi}\right|+\sup _{\phi \in \mathbb{R}^{d}:\|\phi\| \leq 1} \sqrt{\phi^{\top}(\mathbf{A}-\tilde{\mathbf{A}}) \boldsymbol{\phi}} \\
& =\|\mathbf{w}-\tilde{\mathbf{w}}\|+\sqrt{\|\mathbf{A}-\tilde{\mathbf{A}}\|} \\
& \leq\|\mathbf{w}-\tilde{\mathbf{w}}\|+\sqrt{\|\mathbf{A}-\tilde{\mathbf{A}}\|_{\mathrm{F}}} \leq \epsilon+\sqrt{\epsilon} \leq 2 \sqrt{\epsilon} .
\end{aligned}
$$

Therefore,

$$
\begin{equation*}
\left\|\sum_{\tau=1}^{t} \eta^{-\tau} \phi_{s}\left(\epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}-\epsilon_{\tau}^{\tilde{\mathbf{w}}, \tilde{\mathbf{A}}}\right)\right\|_{\widetilde{\Sigma}_{t}^{-1}} \leq \frac{4\left(1-\eta^{t}\right)}{\sqrt{\lambda}(1-\eta)} \sqrt{\epsilon} \tag{D.5.9}
\end{equation*}
$$

The $\epsilon$-covering number of $\mathcal{G}$ as Euclidean ball in $\mathbb{R}^{d+d^{2}}$ of radius $2 H \sqrt{\frac{d}{\lambda(1-\eta)}}+\frac{\beta^{2} \sqrt{d}}{\lambda}$ is bounded by Lemma 25 as $\left(3\left(2 H \sqrt{\frac{d}{\lambda(1-\eta)}}+\frac{\beta^{2} \sqrt{d}}{\lambda}\right) / \epsilon\right)^{d+d^{2}}$. Now, combining Equations (D.5.8) and (D.5.9) we obtain:

$$
\begin{aligned}
& \left\|\sum_{\tau=1}^{t} \eta^{-\tau} \boldsymbol{\phi}_{s} \epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}\right\|_{\widetilde{\Sigma}_{t}^{-1}} \\
& \leq 2 H \sqrt{2\left(d^{2}+d\right) \log \left(\frac{3}{\epsilon}\left(2 H \sqrt{\frac{d}{\lambda(1-\eta)}}+\frac{\beta^{2} \sqrt{d}}{\lambda}\right)\right)+2 \log \left(\frac{1}{\delta}\right)+d \log \left(1+\frac{1-\eta^{2 t}}{\lambda d\left(1-\eta^{2}\right)}\right)} \\
& +\frac{4 \sqrt{\epsilon}\left(1-\eta^{t}\right)}{\sqrt{\lambda}(1-\eta)}
\end{aligned}
$$

Finally by taking $\epsilon=\frac{\lambda(1-\eta)^{2}}{16}$ and keeping only dominant term for each parameter, we obtain.

$$
\left\|\sum_{\tau=1}^{t} \eta^{-\tau} \phi_{s} \epsilon_{\tau}^{\mathbf{w}, \mathbf{A}}\right\|_{\tilde{\Sigma}_{t}^{-1}} \leq \mathcal{O}\left(d H \sqrt{\log \left(d H \beta \frac{1}{\lambda(1-\eta)}\right)+\log (1 / \delta)}\right)
$$

## Appendix for Chapter 7

## E. 1 Proofs

The first proposition is a direct consequence of the definition (7.8.2) of successor states with a goal space $G$.
Proposition 12. Let $\phi: S \times A \rightarrow G$ be a map to some goal space $G$.
Let $\pi$ be some policy, and let $M^{\pi}$ be the successor state measure (7.8.2) of $\pi$ in goal space $G$. Let $m^{\pi}$ be the density of $M^{\pi}$ with respect to some positive probability measure $\rho$ on $G$.

Let $r: G \rightarrow \mathbb{R}$ be some function on $G$, and define the reward function $R(s, a) \triangleq$ $r(\phi(s, a))$ on $S \times A$.

Then the $Q$-function $Q^{\pi}$ of policy $\pi$ for reward $R$ is

$$
\begin{aligned}
Q^{\pi}(s, a) & =\int r(g) M^{\pi}(s, a, \mathrm{~d} g) \\
& =\mathbb{E}_{g \sim \rho}\left[r(g) m^{\pi}(s, a, g)\right]
\end{aligned}
$$

Proof of Proposition 12. For each time $t \geq 0$, let $P_{t}^{\pi}\left(s_{0}, a_{0}, \mathrm{~d} g\right)$ be the probability distribution of $g=\phi\left(s_{t}, a_{t}\right)$ over trajectories of the policy $\pi$ starting at $\left(s_{0}, a_{0}\right)$ in the MDP. Thus, by the definition (7.8.2),

$$
\begin{equation*}
M^{\pi}(s, a, \mathrm{~d} g)=\sum_{t \geq 0} \gamma^{t} P_{t}^{\pi}(s, a, \mathrm{~d} g) \tag{E.1.1}
\end{equation*}
$$

The $Q$-function of $\pi$ for the reward $R$ is by definition (the sums and integrals are finite since $R$ is bounded)

$$
\begin{aligned}
Q^{\pi}(s, a) & =\sum_{t \geq 0} \gamma^{t} \mathbb{E}\left[R\left(s_{t}, a_{t}\right) \mid s_{0}=s, a_{0}=a, \pi\right] \\
& =\sum_{t \geq 0} \gamma^{t} \mathbb{E}\left[r\left(\phi\left(s_{t}, a_{t}\right)\right) \mid s_{0}=s, a_{0}=a, \pi\right] \\
& =\sum_{t \geq 0} \gamma^{t} \int_{g} r(g) P_{t}^{\pi}(s, a, \mathrm{~d} g) \\
& =\int_{g} r(g) M^{\pi}(s, a, \mathrm{~d} g) \\
& =\int_{g} r(g) m^{\pi}(s, a, g) \rho(\mathrm{d} g) \\
& =\mathbb{E}_{g \sim \rho}\left[r(g) m^{\pi}(s, a, g)\right]
\end{aligned}
$$

by definition of the density $m^{\pi}$.
Proof of Theorems 1 and 2. Theorem 1 is a particular case of Theorem 2 ( $\phi=$ Id and $\bar{m}=0$ ), so we only prove the latter.

Let $R(s, a)=r(\phi(s, a))$ be a reward function as in the theorem.
The $Q$-function of $\pi$ for the reward $R$ is, by Proposition 12,

$$
\begin{equation*}
Q^{\pi}(s, a)=\mathbb{E}_{g \sim \rho}\left[r(g) m^{\pi}(s, a, g)\right] \tag{E.1.2}
\end{equation*}
$$

The assumptions state that for any $z \in Z, m^{\pi_{z}}(s, a, g)$ is equal to $F(s, a, z)^{\top} B(g)+\bar{m}(s, z, g)$. Therefore, for any $z \in Z$ we have

$$
\begin{aligned}
Q^{\pi_{z}}(s, a) & =\mathbb{E}_{g \sim \rho}\left[r(g) F(s, a, z)^{\top} B(g)+r(g) \bar{m}(s, z, g)\right] \\
& =F(s, a, z)^{\top} \mathbb{E}_{g \sim \rho}[r(g) B(g)]+\mathbb{E}_{g \sim \rho}[r(g) \bar{m}(s, z, g)] \\
& =F(s, a, z)^{\top} z_{R}+\bar{V}^{z}(s)
\end{aligned}
$$

by definition of $z_{R}$ and $\bar{V}$. This proves the claim (7.8.7) about $Q$-functions.
By definition, the policy $\pi_{z}$ selects the action $a$ that maximizes $F(s, a, z)^{\top} z$. Take $z=z_{R}$. Then

$$
\begin{align*}
\pi_{z_{R}} & =\underset{a}{\operatorname{argmax}} F\left(s, a, z_{R}\right)^{\top} z_{R}  \tag{E.1.3}\\
& =\underset{a}{\operatorname{argmax}}\left\{F\left(s, a, z_{R}\right)^{\top} z_{R}+\bar{V}^{z}(s)\right\} \tag{E.1.4}
\end{align*}
$$

since the last term does not depend on $a$.
This quantity is equal to $Q^{\pi_{z_{R}}}(s, a)$. Therefore,

$$
\begin{equation*}
\pi_{z_{R}}=\underset{a}{\operatorname{argmax}} Q^{\pi_{z_{R}}}(s, a) \tag{E.1.5}
\end{equation*}
$$

and by the above, $Q^{\pi_{z_{R}}}(s, a)$ is indeed equal to the $Q$-function of policy $\pi_{z_{R}}$ for the reward $R$. Therefore, $\pi_{z_{R}}$ and $Q^{\pi_{z_{R}}}$ constitute an optimal Bellman pair for reward $R$. Since $Q^{\pi_{z_{R}}}(s, a)$ is the $Q$-function of $\pi_{Z_{R}}$, it satisfies the Bellman equation

$$
\begin{align*}
Q^{\pi_{z_{R}}(s, a)} & =R(s, a)+\gamma \mathbb{E}_{s^{\prime} \mid(s, a)} Q^{\pi_{z_{R}}}\left(s^{\prime}, \pi_{z_{R}}\left(s^{\prime}\right)\right)  \tag{E.1.6}\\
& =R(s, a)+\gamma \mathbb{E}_{s^{\prime} \mid(s, a)} \max _{a^{\prime}} Q^{\pi_{z_{R}}}\left(s^{\prime}, a^{\prime}\right) \tag{E.1.7}
\end{align*}
$$

by (E.1.5). This is the optimal Bellman equation for $R$, and $\pi_{z_{R}}$ is the optimal policy for $R$.

We still have to prove the last statement of Theorem 2. Since $\pi_{z_{R}}$ is an optimal policy for $R$, for any other policy $\pi_{z}$ and state-action $(s, a)$ we have

$$
\begin{equation*}
Q^{\pi_{z_{R}}(s, a) \geq Q^{\pi_{z}}(s, a) .} \tag{E.1.8}
\end{equation*}
$$

Using the formulas above for $Q^{\pi}$, with $\bar{m}=0$, this rewrites as

$$
\begin{equation*}
F\left(s, a, z_{R}\right)^{\top} z_{R} \geq F(s, a, z)^{\top} z_{R} \tag{E.1.9}
\end{equation*}
$$

as needed. Thus $F\left(s, a, z_{R}\right)^{\top} z_{R} \geq \sup _{z \in Z} F(s, a, z)^{\top} z_{R}$, and equality occurs by taking $z=z_{R}$. This ends the proof of Theorem 2 .

Proof of Proposition 6. Assume $d=\# S \times \# A$; extra dimensions can just be ignored by setting the extra components of $F$ and $B$ to 0 .

With $d=\# S \times \# A$, we can index the components of $Z$ by pairs $(s, a)$.
First, let us set $B(s, a) \triangleq \mathbb{1}_{s, a}$.
Let $r: S \times A \rightarrow \mathbb{R}$ be any reward function. Let $z_{R} \in \mathbb{R}^{\# S \times \# A}$ be defined as in Theorem 1, namely, $z_{R}=\mathbb{E}_{(s, a) \sim \rho}[r(s, a) B(s, a)]$. With our choice of $B$, the components of $z_{R}$ are $\left(z_{R}\right)_{s, a}=r(s, a) \rho(s, a)$. Since $\rho>0$, the correspondence $r \leftrightarrow z_{R}$ is bijective.

Let us now define $F$. Take $z \in Z$. Since $r \leftrightarrow z_{R}$ is bijective, this $z$ is equal to $z_{R}$ for some reward function $r$. Let $\pi_{z}$ be an optimal policy for this reward $r$ in the MDP. Let $M^{\pi_{r}}$ be the successor state measure of policy $\pi_{z}$, namely:

$$
\begin{equation*}
M^{\pi_{z}}\left(s, a, s^{\prime}, a^{\prime}\right)=\sum_{t \geq 0} \gamma^{t} \operatorname{Pr}\left(\left(s_{t}, a_{t}\right)=\left(s^{\prime}, a^{\prime}\right) \mid\left(s_{0}, a_{0}\right)=(s, a), \pi_{z}\right) . \tag{E.1.10}
\end{equation*}
$$

Now define $F(s, a, z)$ by setting its $\left(s^{\prime}, a^{\prime}\right)$ component to $M^{\pi_{z}}\left(s, a, s^{\prime}, a^{\prime}\right) / \rho\left(s^{\prime}, a^{\prime}\right)$ for each $\left(s^{\prime}, a^{\prime}\right)$ :

$$
\begin{equation*}
F(s, a, z)_{s^{\prime}, a^{\prime}} \triangleq M^{\pi_{z}}\left(s, a, s^{\prime}, a^{\prime}\right) / \rho\left(s^{\prime}, a^{\prime}\right) \tag{E.1.11}
\end{equation*}
$$

Then we have

$$
\begin{align*}
& F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)=\sum_{s^{\prime \prime}, a^{\prime \prime}} F(s, a, z)_{s^{\prime \prime}, a^{\prime \prime}} B\left(s^{\prime}, a^{\prime}\right)_{s^{\prime \prime}, a^{\prime \prime}} \\
&=F(s, a, z)_{s^{\prime}, a^{\prime}}=M^{\pi_{z}}\left(s, a, s^{\prime}, a^{\prime}\right) / \rho\left(s^{\prime}, a^{\prime}\right) \tag{E.1.12}
\end{align*}
$$

because by our choice of $B, B\left(s^{\prime}, a^{\prime}\right)_{s^{\prime \prime} a^{\prime \prime}}=\mathbb{1}_{s^{\prime}=s^{\prime \prime}, a^{\prime}=a^{\prime \prime}}$.
Thus, $F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)$ is the density of the successor state measure $M^{\pi_{z}}$ of policy $\pi_{z}$ with respect to $\rho$, as needed.

We still have to check that $\pi_{z}$ satisfies $\pi_{z}(s)=\operatorname{argmax} F(s, a, z)^{\top} z$ (since this is not how it was defined). Since $\pi_{z}$ was defined as an optimal policy for the reward $r$ associated with $z$, it satisfies

$$
\begin{equation*}
\pi_{z}(s)=\underset{a}{\operatorname{argmax}} Q^{\pi_{z}}(s, a) \tag{E.1.13}
\end{equation*}
$$

with $Q^{\pi_{z}}(s, a)$ the $Q$-function of policy $\pi_{z}$ for the reward $r$. This $Q$-function is equal to the cumulated expected reward

$$
\begin{aligned}
Q^{\pi_{z}}(s, a) & =\sum_{t \geq 0} \gamma^{t} \mathbb{E}\left[r\left(s_{t}, a_{t}\right) \mid s_{0}=s, a_{0}=a, \pi_{z}\right] \\
& =\sum_{t \geq 0} \gamma^{t} \sum_{s^{\prime}, a^{\prime}} r\left(s^{\prime}, a^{\prime}\right) \operatorname{Pr}\left(\left(s_{t}, a_{t}\right)=\left(s^{\prime}, a^{\prime}\right) \mid s_{0}=s, a_{0}=a, \pi_{z}\right) \\
& =\sum_{s^{\prime}, a^{\prime}} r\left(s^{\prime}, a^{\prime}\right) \sum_{t \geq 0} \gamma^{t} \operatorname{Pr}\left(\left(s_{t}, a_{t}\right)=\left(s^{\prime}, a^{\prime}\right) \mid s_{0}=s, a_{0}=a, \pi_{z}\right) \\
& =\sum_{s^{\prime}, a^{\prime}} r\left(s^{\prime}, a^{\prime}\right) M^{\pi_{z}}\left(s, a, s^{\prime}, a^{\prime}\right) \\
& =\sum_{s^{\prime}, a^{\prime}} r\left(s^{\prime}, a^{\prime}\right) F(s, a, z)_{s^{\prime}, a^{\prime}} \rho\left(s^{\prime}, a^{\prime}\right) \\
& =F(s, a, z)^{\top}\left(\sum_{s^{\prime} a^{\prime}} r\left(s^{\prime}, a^{\prime}\right) \rho\left(s^{\prime}, a^{\prime}\right) \mathbb{1}_{s^{\prime} a^{\prime}}\right) \\
& =F(s, a, z)^{\top} z
\end{aligned}
$$

since $z$ is equal to $\mathbb{E}_{\left(s^{\prime}, a^{\prime}\right) \sim \rho}\left[r\left(s^{\prime}, a^{\prime}\right) B\left(s^{\prime}, a^{\prime}\right)\right]$. This proves that $\pi_{z}(s)=$ $\operatorname{argmax}_{a} Q^{\pi_{z}(s, a)}=\operatorname{argmax}_{a} F(s, a, z)^{\top} z$. So this choice of $F$ and $B$ satisfies all the properties claimed.

We will rely on the following two basic results in Q-learning.
Proposition 13 ( $r \mapsto Q^{\star}$ is Lipschitz in sup-norm). Let $r_{1}, r_{2}: S \times A \rightarrow \mathbb{R}$ be two bounded reward functions. Let $Q_{1}^{\star}$ and $Q_{2}^{\star}$ be the corresponding optimal $Q$-functions, and likewise for the $V$-functions. Then

$$
\begin{equation*}
\sup _{S \times A}\left|Q_{1}^{\star}-Q_{2}^{\star}\right| \leq \frac{1}{1-\gamma} \sup _{S \times A}\left|r_{1}-r_{2}\right| \text { and } \sup _{S}\left|V_{1}^{\star}-V_{2}^{\star}\right| \leq \frac{1}{1-\gamma} \sup _{S \times A}\left|r_{1}-r_{2}\right| . \tag{E.1.14}
\end{equation*}
$$

Moreover for any policy $\pi$, we have

$$
\begin{equation*}
\sup _{S \times A}\left|Q_{1}^{\pi}-Q_{2}^{\pi}\right| \leq \frac{1}{1-\gamma} \sup _{S \times A}\left|r_{1}-r_{2}\right| \text { and } \sup _{S}\left|V_{1}^{\pi}-V_{2}^{\pi}\right| \leq \frac{1}{1-\gamma} \sup _{S \times A}\left|r_{1}-r_{2}\right| . \tag{E.1.15}
\end{equation*}
$$

Proof. Assume sup ${ }_{S \times A}\left|r_{1}-r_{2}\right| \leq \epsilon$ for some $\epsilon \geq 0$.
For any policy $\pi$, let $Q_{1}^{\pi}$ be its $Q$-function for reward $r_{1}$, and likewise for $r_{2}$. Let $\pi_{1}$ and $\pi_{2}$ be optimal policies for $r_{1}$ and $r_{2}$, respectively. Then for any

$$
\begin{aligned}
& (s, a) \in S \times A, \\
& Q_{1}^{\star}(s, a)=Q_{1}^{\pi_{1}}(s, a) \\
& \geq Q_{1}^{\pi_{2}}(s, a) \\
& =\sum_{t \geq 0} \gamma^{t} \mathbb{E}\left[r_{1}\left(s_{t}, a_{t}\right) \mid \pi_{2},\left(s_{0}, a_{0}\right)=(s, a)\right] \\
& \geq \sum_{t \geq 0} \gamma^{t} \mathbb{E}\left[r_{2}\left(s_{t}, a_{t}\right)-\epsilon \mid \pi_{2},\left(s_{0}, a_{0}\right)=(s, a)\right] \\
& =\sum_{t \geq 0} \gamma^{t} \mathbb{E}\left[r_{2}\left(s_{t}, a_{t}\right) \mid \pi_{2},\left(s_{0}, a_{0}\right)=(s, a)\right]-\frac{\epsilon}{1-\gamma} \\
& =Q_{2}^{\star}(s, a)-\frac{\epsilon}{1-\gamma}
\end{aligned}
$$

and likewise in the other direction, which ends the proof for $Q$-functions. The case of $V$-functions follows by restricting to the optimal actions at each state $s$.

Now, let $\pi$ a policy. We have

$$
\begin{aligned}
\left|Q_{1}^{\pi}(s, a)-Q_{2}^{\pi}(s, a)\right|= & \mid \sum_{t \geq 0} \gamma^{t} \mathbb{E}\left[r_{1}\left(s_{t}, a_{t}\right) \mid \pi,\left(s_{0}, a_{0}\right)=(s, a)\right] \\
& -\sum_{t \geq 0} \gamma^{t} \mathbb{E}\left[r_{2}\left(s_{t}, a_{t}\right) \mid \pi,\left(s_{0}, a_{0}\right)=(s, a)\right] \mid \\
\leq & \sum_{t \geq 0} \gamma^{t} \mathbb{E}\left[\left|r_{2}\left(s_{t}, a_{t}\right)-r_{1}\left(s_{t}, a_{t}\right)\right| \mid \pi,\left(s_{0}, a_{0}\right)=(s, a)\right] \\
\leq & \frac{1}{1-\gamma} \sup _{S \times A}\left|r_{1}-r_{2}\right|
\end{aligned}
$$

The case of $V$-functions follows by taking the expectation over actions according to $\pi$.

Proposition 14. Let $f: S \times A \rightarrow \mathbb{R}$ be any function, and define a policy $\pi_{f}$ by $\pi_{f}(s) \triangleq$ $\operatorname{argmax}_{a} f(s, a)$. Let $r: S \times A \rightarrow \mathbb{R}$ be some bounded reward function. Let $Q^{\star}$ be its optimal $Q$-function, and let $Q^{\pi_{f}}$ be the $Q$-function of $\pi_{f}$ for reward $r$.

Then

$$
\begin{equation*}
\sup _{S \times A}\left|f-Q^{\star}\right| \leq \frac{2}{1-\gamma} \sup _{S \times A}\left|f-Q^{\pi_{f}}\right| \tag{E.1.16}
\end{equation*}
$$

and

$$
\begin{equation*}
\sup _{S \times A}\left|Q^{\pi_{f}}-Q^{\star}\right| \leq \frac{3}{1-\gamma} \sup _{S \times A}\left|f-Q^{\pi_{f}}\right| \tag{E.1.17}
\end{equation*}
$$

Proof of Theorem 5. By construction of the Kantorovich-Rubinstein norm, the second claim of Theorem 5 is a particular case of the third claim, with $\|f\|_{A} \triangleq$ $\max \left(\|f\|_{\infty},\|f\|_{\text {Lip }}\right)$ and $\|\mu\|_{B} \triangleq\|\mu\|_{\text {KR }}$.

Likewise, since $m$ is the density of $M$ with respect to $\rho$, the first claim is an instance of the third, by taking $\|f\|_{A} \triangleq\|f\|_{\infty}$ and $\|\mu\|_{B} \triangleq\left\|\frac{\mathrm{~d} \mu}{\mathrm{~d} \rho}\right\|_{L^{1}(\rho)}$. Therefore, we only prove the third claim.

Let $z \in Z$ and let $r: S \times A \rightarrow \mathbb{R}$ be any reward function. By Proposition 12 with $G=S \times A$ and $\phi=\mathrm{Id}$, the $Q$-function of policy $Q^{\pi_{z}}$ for this reward is

$$
\begin{equation*}
Q^{\pi_{z}}(s, a)=\int r\left(s^{\prime}, a^{\prime}\right) M^{\pi_{z}}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right) \tag{E.1.18}
\end{equation*}
$$

Let $\varepsilon_{z}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right)$ be the difference of measures between the model $F^{\top} B \rho$ and $M^{\pi_{z}}$ :

$$
\begin{align*}
\varepsilon_{z}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right) & \triangleq M^{\pi_{z}}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right)-\hat{M}^{z}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right)  \tag{E.1.19}\\
& =M^{\pi_{z}}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right)-F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right) \tag{E.1.20}
\end{align*}
$$

We want to control the optimality gap in terms of $\sup _{s, a}\left\|\varepsilon_{z}(s, a, \cdot)\right\|_{B}$.
By definition of $\varepsilon_{z}$,

$$
\begin{align*}
Q^{\pi_{z}}(s, a) & =\int r\left(s^{\prime}, a^{\prime}\right) F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)+\int r\left(s^{\prime}, a^{\prime}\right) \varepsilon_{z}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right)  \tag{E.1.21}\\
& =F(s, a, z)^{\top} z_{R}+\int r\left(s^{\prime}, a^{\prime}\right) \varepsilon_{z}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right) \tag{E.1.22}
\end{align*}
$$

since $z_{R}=\int r\left(s^{\prime}, a^{\prime}\right) B\left(s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)$. Therefore,

$$
\begin{aligned}
\left|Q^{\pi_{z}}(s, a)-F(s, a, z)^{\top} z_{R}\right| & =\left|\int r\left(s^{\prime}, a^{\prime}\right) \varepsilon_{z}\left(s, a, \mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime}\right)\right| \\
& \leq\|r\|_{A}\left\|\varepsilon_{z}(s, a, \cdot)\right\|_{B}
\end{aligned}
$$

for any reward $r$ and any $z \in Z$ (not necessarily $z=z_{R}$ ).
Let $Q^{\star}$ be the optimal $Q$-function for reward $r$. Define $f(s, a) \triangleq F\left(s, a, z_{R}\right)^{\top} z_{R}$. By definition, the policy $\pi_{z_{R}}$ is equal to $\operatorname{argmax}_{a} f(s, a)$. Therefore, by Proposition 14,

$$
\begin{equation*}
\sup _{S \times A}\left|Q^{\pi_{z_{R}}}-Q^{\star}\right| \leq \frac{3}{1-\gamma} \sup _{S \times A}\left|f-Q^{\pi_{z_{R}}}\right| \tag{E.1.23}
\end{equation*}
$$

and

$$
\begin{equation*}
\sup _{S \times A}\left|f-Q^{\star}\right| \leq \frac{2}{1-\gamma} \sup _{S \times A}\left|f-Q^{\pi_{z_{R}}}\right| \tag{E.1.24}
\end{equation*}
$$

But by the above,

$$
\begin{aligned}
\sup _{S \times A} \mid f-Q^{\pi_{z_{R}}} & =\sup _{S \times A}\left|F\left(s, a, z_{R}\right)^{\top} z_{R}-Q^{\pi_{z_{R}}(s, a)}\right| \\
& \leq\|r\|_{A} \sup _{S \times A}\left\|\varepsilon_{z_{R}}(s, a, \cdot)\right\|_{B}
\end{aligned}
$$

Therefore, for any reward function $r$,

$$
\begin{equation*}
\sup _{S \times A}\left|Q^{\pi_{z_{R}}}-Q^{\star}\right| \leq \frac{3\|r\|_{A}}{1-\gamma} \sup _{S \times A}\left\|\varepsilon_{z_{R}}(s, a, \cdot)\right\|_{B} \tag{E.1.25}
\end{equation*}
$$

This inequality transfers to the value functions, hence the result. In addition, using again $f(s, a)=F\left(s, a, z_{R}\right)^{\top} z_{R}$, we obtain

$$
\begin{equation*}
\sup _{S \times A}\left|F\left(s, a, z_{R}\right)^{\top} z_{R}-Q^{\star}(s, a)\right| \leq \frac{2\|r\|_{A}}{1-\gamma} \sup _{S \times A}\left\|\varepsilon_{z_{R}}(s, a, \cdot)\right\|_{B} . \tag{E.1.26}
\end{equation*}
$$

Proposition 15 (Pointwise optimality gap). Assume the state space is finite, and view rewards and $Q$-functions as vectors over state-actions.

Let $r_{1}$ and $r_{2}$ be two reward functions, and let $\pi_{1}$ and $\pi_{2}$ be optimal policies for $r_{1}$ and $r_{2}$ respectively. Let $P_{1}$ and $P_{2}$ be the stochastic transition matrices over state-actions induced by $\pi_{1}$ and $\pi_{2}$.

Then the optimality gap of policy $\pi_{2}$ on reward $r_{1}$ is at most

$$
\begin{equation*}
0 \leq Q_{r_{1}}^{\pi_{1}}-Q_{r_{1}}^{\pi_{2}} \leq\left(\left(\operatorname{Id}-\gamma P_{1}\right)^{-1}-\left(\operatorname{Id}-\gamma P_{2}\right)^{-1}\right)\left(r_{1}-r_{2}\right) \tag{E.1.27}
\end{equation*}
$$

where the equality holds componentwise viewing the Q-functions as vectors over stateactions.

Proof. This is a classical result. The inequality $0 \leq Q_{r_{1}}^{\pi_{1}}-Q_{r_{1}}^{\pi_{2}}$ is trivial since $\pi_{1}$ is optimal for $r_{1}$ and the optimal policy is optimal at every state-action. Denote $M_{1} \triangleq\left(\operatorname{Id}-\gamma P_{1}\right)^{-1}$ and likewise for $M_{2}$. Then for any reward function $r$ one has $Q_{r}^{\pi_{1}}=M_{1} r$ and likewise for $M_{2}$. Therefore

$$
\begin{aligned}
Q_{r_{1}}^{\pi_{1}}-Q_{r_{1}}^{\pi_{2}} & =M_{1} r_{1}-M_{2} r_{1} \\
& =M_{1} r_{1}-M_{1} r_{2}+M_{1} r_{2}-M_{2} r_{1} \\
& \leq M_{1} r_{1}-M_{1} r_{2}+M_{2} r_{2}-M_{2} r_{1} \quad \text { since } \pi_{2} \text { is optimal for } r_{2} \\
& =\left(M_{1}-M_{2}\right)\left(r_{1}-r_{2}\right) .
\end{aligned}
$$

Proof of Theorem 6. Let $f: S \times A \rightarrow \mathbb{R}$ be any function. Define the policy $\pi_{f}(s) \triangleq$ $\operatorname{argmax}_{a} f(s, a)$. Define $r^{\prime} \triangleq\left(\operatorname{Id}-\gamma P^{\pi_{f}}\right) f$. The equality $f=r^{\prime}+\gamma P^{\pi_{f}} f$ can be rewritten as

$$
\begin{equation*}
f(s, a)=r^{\prime}(s, a)+\gamma \mathbb{E}_{s^{\prime} \sim P\left(\mathrm{~d} s^{\prime} \mid s, a\right)} f\left(s^{\prime}, \pi_{f}\left(s^{\prime}\right)\right) \tag{E.1.28}
\end{equation*}
$$

But by definition of $\pi_{f}, \pi_{f}\left(s^{\prime}\right)=\operatorname{argmax}_{a^{\prime}} f\left(s^{\prime}, a^{\prime}\right)$. Therefore

$$
\begin{equation*}
f(s, a)=r^{\prime}(s, a)+\gamma \mathbb{E}_{s^{\prime} \sim P\left(\mathrm{~d} s^{\prime} \mid s, a\right)} \max _{a^{\prime}} f\left(s^{\prime}, a^{\prime}\right) \tag{E.1.29}
\end{equation*}
$$

namely, $f$ is the optimal $Q$-function for reward $r^{\prime}$, with $\pi_{f}$ the corresponding optimal policy.

Let $r$ be any reward function and let $Q^{\pi_{f}}$ be the $Q$-function of $\pi_{f}$ for $r$. By definition, it satisfies the Bellman equation $Q^{\pi_{f}}=r+\gamma P^{\pi_{f}} Q^{\pi_{f}}$, namely, $r=$ $\left(\operatorname{Id}-\gamma P^{\pi_{f}}\right) Q^{\pi_{f}}$. Therefore,

$$
\begin{equation*}
r-r^{\prime}=\left(\operatorname{Id}-\gamma P^{\pi_{f}}\right)\left(Q^{\pi_{f}}-f\right) \tag{E.1.30}
\end{equation*}
$$

For any policy $\pi$, denote $M^{\pi} \triangleq\left(\operatorname{Id}-\gamma P^{\pi}\right)^{-1}$. $\left(M^{\pi}\right.$ is the successor measure seen as a matrix.) The $Q$-function of $\pi$ for reward $r$ is $M^{\pi} r$.

Since $\pi^{\star}$ is optimal for $r$, and $\pi_{f}$ is optimal for $r^{\prime}$, Proposition 15 yields

$$
\begin{aligned}
0 \leq Q^{\star}-Q^{\pi_{f}} & \leq\left(M^{\pi^{\star}}-M^{\pi_{f}}\right)\left(r-r^{\prime}\right) \\
& =\left(M^{\pi^{\star}}-M^{\pi_{f}}\right)\left(\operatorname{Id}-\gamma P^{\pi_{f}}\right)\left(Q^{\pi_{f}}-f\right)
\end{aligned}
$$

Since $M^{\pi_{f}}$ is the inverse of $\operatorname{Id}-\gamma P^{\pi_{f}}$ and likewise for $\pi^{\star}$, we have

$$
\begin{aligned}
\left(M^{\pi^{\star}}-M^{\pi_{f}}\right)\left(\operatorname{Id}-\gamma P^{\pi_{f}}\right) & =M^{\pi^{\star}}\left(\operatorname{Id}-\gamma P^{\pi_{f}}\right)-\text { Id } \\
& =M^{\pi^{\star}}\left(\operatorname{Id}-\gamma P^{\pi_{f}}-\left(\operatorname{Id}-\gamma P^{\pi^{\star}}\right)\right) \\
& =\gamma M^{\pi^{\star}}\left(P^{\pi^{\star}}-P^{\pi_{f}}\right)
\end{aligned}
$$

and therefore

$$
\begin{equation*}
0 \leq Q^{\star}-Q^{\pi_{f}} \leq \gamma M^{\pi^{\star}}\left(P^{\pi^{\star}}-P^{\pi_{f}}\right)\left(Q^{\pi_{f}}-f\right) \tag{E.1.31}
\end{equation*}
$$

Now set

$$
\begin{equation*}
f(s, a) \triangleq F\left(s, a, z_{R}\right)^{\top} z_{R} \tag{E.1.32}
\end{equation*}
$$

or in matrix notation, $f=F\left(z_{R}\right)^{\top} z_{R}$. Then $\pi_{f}=\pi_{z_{R}}$ by definition. Thus $Q^{\pi_{f}}=$ $M^{\pi_{z_{R}} r}=m^{\pi_{z_{R}}} \operatorname{diag}(\rho) r$ in matrix notation. Moreover, $z_{R}=\mathbb{E}_{\rho}[B(s, a) r(s, a)]=$ $B \operatorname{diag}(\rho) r$ in matrix notation. So $f=F\left(z_{R}\right)^{\top} B \operatorname{diag}(\rho) r$. Therefore,

$$
\begin{equation*}
Q^{\pi_{f}}-f=\left(m^{\pi_{z_{R}}}-F\left(z_{R}\right)^{\top} B\right) \operatorname{diag}(\rho) r \tag{E.1.33}
\end{equation*}
$$

and

$$
\begin{equation*}
Q^{\star}-Q^{\pi_{f}} \leq \gamma M^{\pi^{\star}}\left(P_{\pi^{\star}}-P_{\pi_{f}}\right)\left(m^{\pi_{z_{R}}}-F\left(z_{R}\right)^{\top} B\right) \operatorname{diag}(\rho) r \tag{E.1.34}
\end{equation*}
$$

and using $M^{\pi^{\star}}=\sum_{t \geq 0} \gamma^{t} P_{\pi_{\star}}^{t}$ provides the required inequality.
As a remark, the same proof works on continuous state spaces, by viewing all matrices as linear operators over functions on $S \times A$.

Proof of Proposition 7. This is just a triangle inequality.

$$
\begin{aligned}
\left\|V_{r}^{\pi_{\hat{z}_{R}}}-V_{r}^{\star}\right\|_{\infty} & \leq\left\|V_{r}^{\pi_{\hat{z}_{R}}}-V_{\hat{r}}^{\pi_{\hat{z}_{R}}}\right\|_{\infty}+\left\|V_{\hat{r}}^{\pi_{\hat{z}_{R}}}-V_{\hat{r}}^{\star}\right\|_{\infty}+\left\|V_{\hat{r}}^{\star}-V_{r}^{\star}\right\|_{\infty} \\
& \leq \frac{\|r-\hat{r}\|_{\infty}}{1-\gamma}+\epsilon_{F B}+\frac{\|r-\hat{r}\|_{\infty}}{1-\gamma} .
\end{aligned}
$$

The last inequality follows from two facts: by assumption, the difference between the value function of $\pi_{\hat{z}_{R}}$ and the optimal value function $V_{\hat{r}}^{\star}$ is at most $\varepsilon_{F B}$, then by Proposition 13, the difference between $V_{\hat{r}}^{\star}$ and $V_{r}^{\star}$ as well as the difference between $V_{r}^{\pi_{\varepsilon_{R}}}$ and $V_{\hat{r}}^{\pi_{2_{R}}}$ are bounded by $\frac{1}{1-\gamma} \sup _{S \times A}|\hat{r}-r|$.

Proof of Theorem 7. We proceed by building a reward function $\hat{r}$ corresponding to $\hat{z}_{R}$. Then we will bound $\hat{r}-r$ and apply Proposition 7 .

First, by Remark 4, up to reducing $d$, we can assume that $B$ is $L^{2}(\rho)$ orthonormal.

For any function $\phi:(s, a) \rightarrow \mathbb{R}$, define $z_{\phi} \triangleq \mathbb{E}_{(s, a) \sim \rho}[\phi(s, a) B(s, a)]$. For each $z \in Z$, define $\phi_{z}$ via $\phi_{z}(s, a) \triangleq B(s, a)^{\top} z$. Then, if $B$ is $L^{2}(\rho)$-orthonormal, we have $z_{\phi_{z}}=z$. (Indeed, $z_{\phi_{z}}=\mathbb{E}\left[\left(B(s, a)^{\top} z\right) B(s, a)\right]=\mathbb{E}\left[B(s, a)\left(B(s, a)^{\top} z\right)\right]=$ $\left(\mathbb{E}\left[B(s, a) B(s, a)^{\top}\right]\right) z$.)

Define the function

$$
\begin{equation*}
\hat{r} \triangleq r+\phi_{\hat{z}_{R}-z_{R}} \tag{E.1.35}
\end{equation*}
$$

using the functions $\phi_{z}$ defined above. By construction, $z_{\hat{\gamma}}=z_{R}+z_{\phi_{2_{R}}-z_{R}}=\hat{z}_{R}$. Therefore, the policy $\pi_{\hat{z}_{R}}$ associated to $\hat{z}_{R}$ is the policy associated to the reward $\hat{r}$.

We will now apply Proposition 7 to $r$ and $\hat{r}$. For this, we need to bound $\left\|\phi_{\hat{z}_{R}-z_{R}}\right\|_{\infty}$.

Let $B_{1}, B_{2}, \ldots, B_{d}$ be the components of $B$ as functions on $S \times A$. For any $z \in Z$, we have

$$
\begin{equation*}
\left\|\phi_{z}\right\|_{L^{2}(\rho)}^{2}=\left\|\sum_{i} z_{i} B_{i}\right\|_{L^{2}(\rho)}^{2}=\sum_{i} z_{i}^{2}=\|z\|^{2} \tag{E.1.36}
\end{equation*}
$$

since the $B_{i}$ are $L^{2}(\rho)$-orthonormal. Moreover, by construction, $\phi_{z}$ lies in the linear span $\langle B\rangle$ of the functions $\left(B_{i}\right)$. Therefore

$$
\begin{equation*}
\left\|\phi_{z}\right\|_{\infty} \leq \zeta(B)\left\|\phi_{z}\right\|_{L^{2}(\rho)}=\zeta(B)\|z\| \tag{E.1.37}
\end{equation*}
$$

by the definition of $\zeta(B)$ (Definition 6).
Therefore,

$$
\begin{equation*}
\left\|\phi_{\hat{z}_{R}-z_{R}}\right\|_{\infty} \leq \zeta(B)\left\|\hat{z}_{R}-z_{R}\right\| . \tag{E.1.38}
\end{equation*}
$$

Let us now bound $\hat{z}_{R}-z_{R}$ :

$$
\begin{aligned}
& \mathbb{E}\left[\left\|\hat{z}_{R}-z_{R}\right\|^{2}\right] \\
& =\mathbb{E}\left[\mathbb{E}\left[\left\|\hat{z}_{R}-z_{R}\right\|^{2} \mid\left(s_{i}, a_{i}\right)\right]\right] \\
& =\mathbb{E}\left[\mathbb{E}\left[\left\|\hat{z}_{R}-\mathbb{E}\left[\hat{z}_{R} \mid\left(s_{i}, a_{i}\right)\right]\right\|^{2}+\left\|\mathbb{E}\left[\hat{z}_{R} \mid\left(s_{i}, a_{i}\right)\right]-z_{R}\right\|^{2} \mid\left(s_{i}, a_{i}\right)\right]\right] \\
& =\mathbb{E}\left[\mathbb{E}\left[\left.\left\|\frac{1}{N} \sum_{i}\left(\hat{r}_{i}-r\left(s_{i}, a_{i}\right)\right) B\left(s_{i}, a_{i}\right)\right\|^{2}+\left\|\frac{1}{N} \sum_{i} r\left(s_{i}, a_{i}\right) B\left(s_{i}, a_{i}\right)-z_{R}\right\|^{2} \right\rvert\,\left(s_{i}, a_{i}\right)\right]\right]
\end{aligned}
$$

The first term satisfies
$\begin{aligned} \mathbb{E}\left[\left.\left\|\frac{1}{N} \sum_{i}\left(\hat{r}_{i}-r\left(s_{i}, a_{i}\right)\right) B\left(s_{i}, a_{i}\right)\right\|^{2} \right\rvert\,\left(s_{i}, a_{i}\right)\right] & =\frac{1}{N^{2}} \sum_{i} \mathbb{E}\left[\left(\hat{r}_{i}-r\left(s_{i}, a_{i}\right)\right)^{2}\left\|B\left(s_{i}, a_{i}\right)\right\|^{2}\right] \\ & \leq \frac{1}{N^{2}} \sum_{i} v\left\|B\left(s_{i}, a_{i}\right)\right\|^{2}\end{aligned}$
because the $\hat{r}_{i}$ are independent conditionally to $\left(s_{i}, a_{i}\right)$, and because $B$ is deterministic. The expectation of this over $\left(s_{i}, a_{i}\right)$ is

$$
\begin{equation*}
\mathbb{E}\left[\frac{1}{N^{2}} \sum_{i} v\left\|B\left(s_{i}, a_{i}\right)\right\|^{2}\right]=\frac{v}{N} \mathbb{E}_{(s, a) \sim \rho}\|B(s, a)\|^{2}=\frac{v}{N}\|B\|_{L^{2}(\rho)}^{2} \tag{E.1.39}
\end{equation*}
$$

which is thus a bound on the first term.
The second term satisfies
$\mathbb{E}\left[\left\|\frac{1}{N} \sum_{i} r\left(s_{i}, a_{i}\right) B\left(s_{i}, a_{i}\right)-z_{R}\right\|^{2}\right]=\frac{1}{N}\left\|r(s, a) B(s, a)-\mathbb{E}_{\rho}[r(s, a) B(s, a)]\right\|_{L^{2}(\rho)}^{2}$
since the $\left(s_{i}, a_{i}\right)$ are independent with distribution $\rho$. By the Cauchy-Schwarz inequality (applied to each component of $B$ ), this is at most

$$
\begin{equation*}
\frac{1}{N}\left\|r(s, a)-\mathbb{E}_{\rho} r\right\|_{L^{2}(\rho)}^{2}\|B\|_{L^{2}(\rho)}^{2} \tag{E.1.41}
\end{equation*}
$$

Therefore,

$$
\begin{equation*}
\mathbb{E}\left[\left\|\hat{z}_{R}-z_{R}\right\|^{2}\right] \leq\left(v+\left\|r(s, a)-\mathbb{E}_{\rho} r\right\|_{L^{2}(\rho)}^{2}\right) \frac{\|B\|_{L^{2}(\rho)}^{2}}{N} \tag{E.1.42}
\end{equation*}
$$

Since $B$ is orthonormal in $L^{2}(\rho)$, we have $\|B\|_{L^{2}(\rho)}^{2}=d$. Putting everything together, we find

$$
\begin{equation*}
\mathbb{E}\left[\|\hat{r}-r\|_{\infty}^{2}\right] \leq \frac{\zeta(B) d}{N}\left(v+\left\|r(s, a)-\mathbb{E}_{\rho} r\right\|_{L^{2}(\rho)}^{2}\right) \tag{E.1.43}
\end{equation*}
$$

Therefore, by the Markov inequality, for any $\delta>0$, with probability $1-\delta$,

$$
\begin{equation*}
\|\hat{r}-r\|_{\infty} \leq \sqrt{\frac{\zeta(B) d}{N \delta}\left(v+\left\|r(s, a)-\mathbb{E}_{\rho} r\right\|_{L^{2}(\rho)}^{2}\right)} \tag{E.1.44}
\end{equation*}
$$

hence the conclusion by Proposition 7.
Proof of Theorem 8. Let

$$
\begin{equation*}
m\left(s, a, s^{\prime}, a^{\prime}\right) \triangleq \sum_{t \geq 0} \gamma^{t} \frac{P_{t}\left(\mathrm{~d} s^{\prime}, \mathrm{d} a^{\prime} \mid s, a, \pi_{z}\right)}{\rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)} \tag{E.1.45}
\end{equation*}
$$

so that

$$
\begin{equation*}
\ell(F, B)=\int\left|F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)-m\left(s, a, s^{\prime}, a^{\prime}\right)\right|^{2} \rho(\mathrm{~d} s, \mathrm{~d} a) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right) \tag{E.1.46}
\end{equation*}
$$

Let us first take the derivative with respect to the parameters of $F$. This is 0 by assumption, so we find

$$
\begin{align*}
0 & =\int \partial_{\theta} F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)\left(F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)-m\left(s, a, s^{\prime}, a^{\prime}\right)\right) \rho(\mathrm{d} s, \mathrm{~d} a) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)  \tag{E.1.47}\\
& =\int \partial_{\theta} F(s, a, z)^{\top} G(s, a) \rho(\mathrm{d} s, \mathrm{~d} a) \tag{E.1.48}
\end{align*}
$$

where

$$
\begin{equation*}
G(s, a) \triangleq \int B\left(s^{\prime}, a^{\prime}\right)\left(F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)-m\left(s, a, s^{\prime}, a^{\prime}\right)\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right) \tag{E.1.49}
\end{equation*}
$$

Since the model is overparameterized, we can realize any $L^{2}$ function $f(s, a)$ as the derivative $\partial_{\theta} F(s, a, z)$ for some direction $\theta$. Therefore, the equation $0=\int \partial_{\theta} F(s, a, z)^{\top} G(s, a) \rho(\mathrm{d} s, \mathrm{~d} a)$ implies that $G(s, a)$ is $L^{2}(\rho)$-orthogonal to any function $f(s, a)$ in $L^{2}(\rho)$. Therefore, $G(s, a)$ vanishes $\rho$-almost everywhere, namely

$$
\begin{equation*}
\int B\left(s^{\prime}, a^{\prime}\right) F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)=\int B\left(s^{\prime}, a^{\prime}\right) m\left(s, a, s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right) \tag{E.1.50}
\end{equation*}
$$

Now, since $F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)$ is a real number, $F(s, a, z)^{\top} B\left(s^{\prime}, a^{\prime}\right)=$ $B\left(s^{\prime}, a^{\prime}\right)^{\top} F(s, a, z)$. Therefore, the right-hand-side above rewrites as

$$
\begin{equation*}
\int B\left(s^{\prime}, a^{\prime}\right) B\left(s^{\prime}, a^{\prime}\right)^{\top} F(s, a, z) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)=(\operatorname{Cov} B) F(s, a, z) \tag{E.1.51}
\end{equation*}
$$

so that

$$
\begin{equation*}
(\operatorname{Cov} B) F(s, a, z)=\int B\left(s^{\prime}, a^{\prime}\right) m\left(s, a, s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right) \tag{E.1.52}
\end{equation*}
$$

Unfolding the definition of $m$ yields the statement for $F$. The proof for $B$ is similar.

Proof of Theorem 9. According to the proof of Theorem 8, if $F$ achieves a local extremum of $\ell(F, B)$ given a fixed $B$, we have

$$
\begin{equation*}
(\operatorname{Cov} B) F(s, a, z)=\int B\left(s^{\prime}, a^{\prime}\right) m^{\pi_{z}}\left(s, a, s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right) \tag{E.1.53}
\end{equation*}
$$

where $m^{\pi_{z}}\left(s, a, s^{\prime}, a^{\prime}\right)=\sum_{t \geq 0} \gamma^{t} \frac{P_{t}\left(\mathrm{ds} s^{\prime}, \mathrm{d} a^{\prime} \mid s, a, \pi_{z}\right)}{\rho\left(\mathrm{ds} s^{\prime}, \mathrm{d} a^{\prime}\right)}$ is the successor state density induced by the policy $\pi_{z}$.

Let $r: S \times A \rightarrow \mathbb{R}$ a bounded reward function that lies in the span of $B$ i.e there exists $\omega \in \mathbb{R}^{d}$ such that $r(s, a)=B(s, a)^{\top} \omega$ for any stateaction pair $(s, a)$. This implies that $(\operatorname{Cov} B) \omega=\mathbb{E}_{(s, a) \sim \rho}\left[B(s, a) B(s, a)^{\top} \omega\right]=$ $\mathbb{E}_{(s, a) \sim \rho}[B(s, a) r(s, a)]=z_{R}$, by definition of $z_{R}$.

Therefore,

$$
\begin{aligned}
F\left(s, a, z_{R}\right)^{\top} z_{R} & =F\left(s, a, z_{R}\right)^{\top}(\operatorname{Cov} B) \omega \\
& =\left((\operatorname{Cov} B) F\left(s, a, z_{R}\right)\right)^{\top} \omega \\
& =\left(\int B\left(s^{\prime}, a^{\prime}\right)^{\top} m^{\pi_{z_{R}}}\left(s, a, s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right)\right) \omega \\
& =\int B\left(s^{\prime}, a^{\prime}\right)^{\top} \omega m^{\pi_{z_{R}}}\left(s, a, s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right) \\
& =\int r\left(s^{\prime}, a^{\prime}\right) m^{\pi_{z_{R}}}\left(s, a, s^{\prime}, a^{\prime}\right) \rho\left(\mathrm{d} s^{\prime}, \mathrm{d} a^{\prime}\right) \\
& =Q^{\pi_{z_{R}}}(s, a)
\end{aligned}
$$

Therefore, $\pi_{z_{R}}$ is the greedy policy with respect to its own Q -value. We conclude that $\pi_{z_{R}}$ is optimal for $r$.

Now, let $r: S \times A \rightarrow \mathbb{R}$ be an arbitrary bounded reward function, and let $r_{B}$ be the $L^{2}(\rho)$-projection of $r$ onto the span of $B$. Both $r$ and $r_{B}$ share the same $z_{R}=\mathbb{E}_{(s, a) \sim \rho}[r(s, a) B(s, a)]=\mathbb{E}_{(s, a) \sim \rho}\left[r_{B}(s, a) B(s, a)\right]$. According to the first part of our proof, $\pi_{z_{R}}$ is optimal for $r_{B}$ since $r_{B}$ lies in the span of $B$. Denote by the subscript $r$ in $V_{r}$ the reward function that a value function corresponds to. By Proposition 15,

$$
\begin{equation*}
0 \leq V_{r}^{\star}-V_{r}^{\pi_{z_{R}}} \leq\left(\operatorname{Id}-\gamma P^{\pi^{\star}}\right)^{-1}\left(r-r_{B}\right)-\left(\operatorname{Id}-\gamma P^{\pi_{z_{R}}}\right)^{-1}\left(r-r_{B}\right) \tag{E.1.54}
\end{equation*}
$$

hence, taking norms,

$$
\begin{equation*}
\left\|V^{\pi_{z_{R}}}-V^{\star}\right\|_{\infty} \leq\left\|\left(\operatorname{Id}-\gamma P^{\pi^{\star}}\right)^{-1}\left(r-r_{B}\right)\right\|_{\infty}+\left\|\left(\operatorname{Id}-\gamma P^{\pi_{z_{R}}}\right)^{-1}\left(r-r_{B}\right)\right\|_{\infty} \tag{E.1.55}
\end{equation*}
$$

as needed.
The bound with $\frac{2}{1-\gamma}$ follows by noting that $\left(\operatorname{Id}-\gamma P^{\pi}\right)^{-1}$ is bounded by $\frac{1}{1-\gamma}$ in $L^{\infty}$ norm for any policy $\pi$.

Proof of Proposition 8. From the definition (7.8.33) of $\hat{z}_{R}$ and the definition (7.8.34) of $\hat{r}$, we find

$$
\begin{aligned}
\hat{z}_{R} & =\mathbb{E}_{\rho}[B(s, a) \hat{r}(s, a)] \\
& =\mathbb{E}_{\rho}\left[B(s, a) B(s, a)^{\top} w\right] \\
& =\left(\operatorname{Cov}_{\rho} B\right) w
\end{aligned}
$$

hence the result given the expression (7.8.34) for $w$.
If $\rho=\rho_{\text {test }}$ then the covariances cancel out: we find $\hat{z}_{R}=$ $\left(\operatorname{Cov}_{\rho} B\right) w=\left(\operatorname{Cov}_{\rho} B\right)\left(\operatorname{Cov}_{\rho_{\text {test }}} B\right)^{-1} \mathbb{E}_{\rho_{\text {test }}}[r(s, a) B(s, a)]=\mathbb{E}_{\rho_{\text {test }}}[r(s, a) B(s, a)]=$ $\mathbb{E}_{\rho}[r(s, a) B(s, a)]=z_{R}$.

If $r$ is linear in $B$, then the linear regression model does not depend on the data distribution $\rho_{\text {test }}$ used: if $r(s, a)=B(s, a)^{\top} w_{0}$ then $w=w_{0}$ for any $\rho_{\text {test, }}$ as long as $\operatorname{Cov}_{\rho_{\text {test }}} B$ is invertible. In that case, both $z_{R}$ and $\hat{z}_{R}$ are equal to $\left(\operatorname{Cov}_{\rho} B\right) w_{0}$.

## E. 2 EXperimental Setup

In this section we provide additional information about our experiments.

## E.2.1 Environments

- Discrete maze: is the $11 \times 11$ classical tabular gridworld with foor rooms. States are represented by one-hot unit vectors, $S=\{0,1\}^{121}$. There are five available actions, $A=\{$ left, right, up, down, do nothing $\}$. The dynamics are deterministic and the walls are impassable.
- Continuous maze: is a two dimensional environment with impassable walls. States are represented by their Cartesian coordinates $(x, y) \in$ $S=[0,1]^{2}$. There are five available actions, $A=\{$ left, right, up, down, do nothing $\}$. The execution of one of the actions moves the agent 0.1 units in the desired direction, and normal random noise with zero mean and standard deviation 0.01 is added to the position of the agent (that is, a move along the x axis would be $x^{\prime}=x \pm 0.1+\mathcal{N}(0,0.01)$, where $\mathcal{N}(0,0.01)$ is a normal variable with mean 0 and standard deviation 0.01 ). If after a move the agent ends up outside of $[0,1]^{2}$, the agent's position is clipped (e.g if $x<0$ then we set $x=0$ ). If a move make the agent cross an interior wall, this move is undone. For all algorithms, we convert a state $s=(x, y)$ into feature vector $\phi(s) \in \mathbb{R}^{441}$ by computing the activations of a regular $21 \times 21$ grid of radial basis functions at the point $(s, y)$. Especially, we use Gaussian functions: $\phi(s)=$
$\left(\exp \left(-\frac{\left(x-x_{i}\right)^{2}+\left(y-y_{i}\right)^{2}}{\sigma}\right), \ldots, \exp \left(-\frac{\left(x-x_{441}\right)^{2}+\left(y-y_{441}\right)^{2}}{2 \sigma^{2}}\right)\right)$ where $\left(x_{i}, y_{i}\right)$ is the center of the $i^{\text {th }}$ Gaussian and $\sigma=0.05$.
- FeatchReach: is a variant of the simulated robotic arm environment from Plappert et al. [2018] using discrete actions instead of continuous actions. States are 10 -dimensional vectors consisting of positions and velocities of robot joints. We discretise the original 3-dimensional action space into 6 possible actions using action stepsize of 1 (The same way as in https://github.com/paulorauber/hpg, the implementation of hindsight policy gradient Rauber et al. [2018]). The goal space is 3-dimensional space representing of the position of the object to reach.
- Ms. Pacman: is a variant of the Atari 2600 game Ms. Pacman Bellemare et al. [2013b], where an episode ends when the agent is captured by a monster Rauber et al. [2018]. States are obtained by processing the raw visual input directly from the screen. Frames are preprocessed by cropping, conversion to grayscale and downsampling to $84 \times 84$ pixels. A state $s_{t}$ is the concatenation of ( $x_{t-12}, x_{t-8}, x_{t-4}, x_{t}$ ) frames, i.e. an $84 \times 84 \times 4$ tensor. An action repeat of 12 is used. As Ms. Pacman is not originally a multi-goal domain, we define the set of goals as the set of the 148 reachable coordinate pairs $(x, y)$ on the screen; these can be reached only by learning to avoid monsters. In contrast with Rauber et al. [2018], who use a heuristic to find the agent's position from the screen's pixels, we use the Atari annotated RAM interface wrapper Anand et al. [2019].


## E.2.2 Architectures

We use the same architecture for discrete maze, continuous maze and FeatchReach. Both forward and backward networks are represented by a feedforward neural network with three hidden layers, each with 256 ReLU units. The forward network receives a concatenation of a state and a $z$ vector as input and has $|A| \times d$ as output dimension. The backward network receives a state as input (or gripper's position for FeatchReach) and has $d$ as output dimension. For goal-oriented DQN , the $Q$-value network is also a feedforward neural network with three hidden layers, each with 256 ReLU units. It receives a concatenation of a state and a goal as input and has $|A|$ as output dimension.

For Ms. Pacman, the forward network is represented by a convolutional neural network given by a convolutional layer with 32 filters ( $8 \times 8$, stride 4 ); convolutional layer with 64 filters ( $4 \times 4$, stride 2 ); convolutional layer with 64 filters ( $3 \times 3$, stride 1); and three fully-connected layers, each with 256 units. We use ReLU as activation function. The $z$ vector is concatenated with the output of the third convolutional layer. The output dimension of the final linear layer is
$|A| \times d$. The backward network acts only on agent's position, a 2-dimensional input. It is represented by a feedforward neural network with three hidden layers, each with 256 ReLU units. The output dimension is $d$. For goal-oriented DQN, the $Q$-value network is represented by a convolutional neural network with the same architecture as the one of the forward network. The goal's position is concatenated with the output of the third convolutional layer. The output dimension of the final linear layer is $|A|$.

## E.2.3 Implementation Details

For all environments, we run the algorithms for 800 epochs. Each epoch consists of 25 cycles where we interleave between gathering some amount of transitions, to add to the replay buffer $\mathcal{D}$ (old transitions are thrown when we reach the maximum of its size), and performing 40 steps of stochastic gradient descent on the model parameters. To collect transitions, we generate episodes using some behavior policy. For both mazes, we use a uniform policy while for FetchReach and Ms. Pacman, we use an $\epsilon$-greedy policy ( $\varepsilon=0.2$ ) with respect to the current approximation $F(s, a, z)^{\top} z$ for a sampled $z$. At evaluation time, $\varepsilon$ greedy policies are also used, with a smaller $\varepsilon=0.02$ for all environments except from discrete maze where we use Boltzmann policy with temperature $\tau=1$. We train each models for three different random seeds.

For generality, we will keep using the notation $B(s, a)$ while in our experiments $B$ acts only on $\phi(s, a)$, a part of the state-action space. For discrete and continuous mazes, $\phi(s, a)=s$, for FetchReach, $\phi(s, a)$ the position of arm's gripper and for Ms. Pacman, $\phi(s, a)$ is the 2-dimensional position $(x, y)$ of the agent on the screen.

We denote by $\theta$ and $\omega$ the parameters of forward and backward networks respectively and $\theta^{-}$and $\omega^{-}$the parameters of their corresponding target networks. Both $\theta^{-}$and $\omega^{-}$are updated after each cycle using Polyak averaging; i.e $\theta^{-} \leftarrow \alpha \theta^{-}+(1-\alpha) \theta$ and $\omega^{-} \leftarrow \alpha \omega^{-}+(1-\alpha) \omega$ where $\alpha=0.95$ is the Polyak coefficient.

During training, we sample $z$ from a rescaled Gaussian that we denote $v$. Especially, we sample a $d$-dimensional standard Gaussian variable $x \sim \mathcal{N}(0$, Id $) \in$ $\mathbb{R}^{d}$ and a scalar centered Cauchy variable $u \in \mathbb{R}$ of scale 0.5 , then we set $z=\sqrt{d} u \frac{x}{\|x\|}$. We use a Cauchy distribution to ensure that the norm of $z$ spans the non-negative real numbers space while having a heavy tail: with a pure Gaussian, the norm of $z$ would be very concentrated around a single value. We also scale by $\sqrt{d}$ to ensure that each component of $z$ has an order of magnitude of 1 .

Before being fed to $F, z$ is preprocessed by $z \leftarrow \frac{z}{\sqrt{1+\|z\|_{2}^{2} / d}}$; this way, $z$ ranges over a bounded set in $\mathbb{R}^{d}$, and this takes advantage of optimal policies being equal for a reward $R$ and for $\lambda R$ with $\lambda>0$.

To update network parameters, we compute an empirical loss by sampling 3 mini-batches, each of size $b=128$, of transitions $\left\{\left(s_{i}, a_{i}, s_{i+1}\right)\right\}_{i \in I} \subset \mathcal{D}$, of target state-action pairs $\left\{\left(s_{i}^{\prime}, a_{i}^{\prime}\right)\right\}_{i \in I} \subset \mathcal{D}$ and of $\left\{z_{i}\right\}_{i \in I} \sim v$ :
$\mathscr{L}(\theta, \omega)=-\frac{1}{b} \sum_{i \in I} F_{\theta}\left(s_{i}, a_{i}, z_{i}\right)^{\top} B_{\omega}\left(s_{i}, a_{i}\right)$
$\frac{1}{2 b^{2}} \sum_{i, j \in I^{2}}\left(F_{\theta}\left(s_{i}, a_{i}, z_{i}\right)^{\top} B_{\omega}\left(s_{j}^{\prime}, a_{j}^{\prime}\right)-\gamma \sum_{a \in A} \pi_{z_{i}}\left(a \mid s_{i+1}\right) \cdot F_{\theta^{-}}\left(s_{i+1}, a, z_{i}\right)^{\top} B_{\omega^{-}}\left(s_{j}^{\prime}, a_{j}^{\prime}\right)\right)^{2}$
where we use the Boltzmann policy $\pi_{z_{i}}\left(\cdot \mid s_{i+1}\right)=\operatorname{softmax}\left(F_{\theta^{-}}\left(s_{i+1}, \cdot, z_{i}\right)^{\top} z_{i} / \tau\right)$ with fixed temperature $\tau=200$ to avoid the instability and discontinuity caused by the argmax operator.

Since there is unidentifiability between $F$ and $B$ (Appendix, Remark 4), we include a gradient to make $B$ closer to orthonormal, $\mathbb{E}_{(s, a) \sim \rho} B(s, a) B(s, a)^{\top} \approx \mathrm{Id}$ :

$$
\begin{align*}
& \frac{1}{4} \partial_{\omega}\left\|\mathbb{E}_{(s, a) \sim \rho} B_{\omega}(s, a) B_{\omega}(s, a)^{\top}-\operatorname{Id}\right\|^{2}= \\
& \quad \mathbb{E}_{(s, a) \sim \rho,\left(s^{\prime}, a^{\prime}\right) \sim \rho} \partial_{\omega} B_{\omega}(s, a)^{\top}\left(B_{\omega}(s, a)^{\top} B_{\omega}\left(s^{\prime}, a^{\prime}\right) \cdot B_{\omega}\left(s^{\prime}, a^{\prime}\right)-B(s, a)\right) \tag{E.2.2}
\end{align*}
$$

To compute an unbiased estimate of the latter gradient, we use the following auxiliary empirical loss:
$\mathscr{L}_{\text {reg }}(\omega)=-\frac{1}{b} \sum_{i \in I} B_{\omega}\left(s_{i}, a_{i}\right)^{\top}$ stop-gradient $\left(B_{\omega}\left(s_{i}, a_{i}\right)\right)$
$\frac{1}{b^{2}} \sum_{i, j \in I^{2}} B_{\omega}\left(s_{i}, a_{i}\right)^{\top}$ stop-gradient $\left(B_{\omega}\left(s_{j}^{\prime}, a_{j}^{\prime}\right)\right) \cdot$ stop-gradient $\left(B_{\omega}\left(s_{i}, a_{i}\right)^{\top} B_{\omega}\left(s_{j}^{\prime}, a_{j}^{\prime}\right)\right)$
Finally, we use the Adam optimizer and we update $\theta$ and $\omega$ by taking a gradient step on $\mathscr{L}(\theta, \omega)$ and $\mathscr{L}(\theta, \omega)+\lambda \cdot \mathscr{L}_{\text {reg }}(\omega)$ respectively, where $\lambda$ is a regularization coefficient that we set to 1 for all experiments.

We summarize the hyperparameters used for FB algorithm and goal-oriented DQN in table E. 1 and E. 2 respectively.

## E.2.4 Experimental results

In this section, we provide additional experimental results.

## Goal-Oriented Setup: Quantitative Comparisons

More Complex Rewards: Qualitative Results
Embedding Visualization

| Hyperparameters | Discrete Maze | Continuous Maze | FetchReach | Ms. Pacman |
| :--- | :--- | :--- | :--- | :--- |
| number of cycles per epoch | 25 | 25 | 25 | 25 |
| number of episodes per cycles | 4 | 4 | 2 | 2 |
| number of timesteps per episode | 50 | 30 | 50 | 50 |
| number of updates per cycle | 40 | 40 | 40 | 40 |
| exploration $\boldsymbol{\varepsilon}$ | 1 | 1 | 0.2 | 0.2 |
| evaluation $\boldsymbol{E}$ | Boltzman with $\tau=1$ | 0.02 | 0.02 | 0.02 |
| temperature $\tau$ | 200 | 200 | 200 | 200 |
| learning rate | 0.001 | 0.0005 | 0.0005 | 0.0001 if $d=100$ |
| mini-batch size | 128 | 128 | 128 | else 0.0005 |
| regularization coefficient $\lambda$ | 1 | 1 | 1 | 1 |
| Polyak coefficient $\alpha$ | 0.95 | 0.95 | 0.95 | 0.95 |
| discount factor $\gamma$ | 0.99 | 0.99 | 0.9 | 0.9 |
| replay buffer size | $10^{6}$ | $10^{6}$ | $10^{6}$ | $10^{6}$ |

Table E.1: Hyperparameters of the FB algorithm

| Hyperparameters | Discrete Maze | Continuous Maze | FetchReach | Ms. Pacman |
| :--- | :--- | :--- | :--- | :--- |
| number of cycles per epoch | 25 | 25 | 25 | 25 |
| number of episodes per cycles | 4 | 4 | 2 | 2 |
| number of timesteps per episode | 50 | 30 | 50 | 50 |
| number of updates per cycle | 40 | 40 | 40 | 40 |
| exploration $\varepsilon$ | 0.2 | 0.2 | 0.2 | 0.2 |
| evaluation $\varepsilon$ | Boltzman with $\tau=1$ | 0.02 | 0.02 | 0.02 |
| learning rate | 0.001 | 0.0005 | 0.0005 | 0.0005 |
| mini-batch size | 128 | 128 | 128 | 128 |
| Polyak coefficient $\alpha$ | 0.95 | 0.95 | 0.95 | 0.95 |
| discount factor $\gamma$ | 0.99 | 0.99 | 0.9 | 0.9 |
| replay buffer size | $10^{6}$ | $10^{6}$ | $10^{6}$ | $10^{6}$ |
| ratio of hindsight replay | - | - | - | 0.8 |

Table E.2: Hyperparameters of the goal-oriented DQN algorithm


Figure E.1: Discrete maze: Comparative performance of FB for different dimensions and DQN. Left: the policy quality averaged over 20 randomly selected goals as function of the training epochs. Right: the policy quality averaged over the goal space after 800 training epochs.


Figure E.2: Continuous maze: Comparative performance of FB for different dimensions and DQN. Left: the success rate averaged over 20 randomly selected goals as function of the training epochs. Right: the success rate averaged over 1000 randomly sampled goals after 800 training epochs.


Figure E.3: FetchReach: Comparative performance of FB for different dimensions and DQN. Left: the success rate averaged over 20 randomly selected goals as function of the training epochs. Right: the success rate averaged over 1000 randomly sampled goals after 800 training epochs.


Figure E.4: Ms. Pacman: Comparative performance of FB for different dimensions and DQN. Left: the success rate averaged over 20 randomly selected goals as function of the training epochs. Right: the success rate averaged over the 184 handcrafted goals after training epochs. Note that FB-50 and F-100 have been trained only for 200 epochs.


Figure E.5: Distance to goal of FB for different dimensions and DQN as function of training epochs. Left: Continuous maze. Right: FetchReach.



Figure E.6: Discrete Maze: Heatmap plots of $\max _{a \in A} F\left(s, a, z_{R}\right)^{\top} z_{R}$ (left) and trajectories of the Boltzmann policy with respect to $F\left(s, a, z_{R}\right)^{\top} z_{R}$ with temperature $\tau=1$ (right). Top row: for the task of reaching a target while avoiding a forbidden region, Middle row: for the task of reaching the closest goal among two equally rewarding positions, Bottom row: choosing between a small, close reward and a large, distant one.


Figure E.7: Continuous Maze: Contour plots plot of $\max _{a \in A} F\left(s, a, z_{R}\right)^{\top} z_{R}$ (left) and trajectories of the $\varepsilon$ greedy policy with respect to $F\left(s, a, z_{R}\right)^{\top} z_{R}$ with $\varepsilon=0.1$ (right). Left: for the task of reaching a target while avoiding a forbidden region, Middle: for the task of reaching the closest goal among two equally rewarding positions, Right: choosing between a small, close reward and a large, distant one..


Figure E.8: Ms. Pacman: Trajectories of the $\varepsilon$ greedy policy with respect to $F\left(s, a, z_{R}\right)^{\top} z_{R}$ with $\varepsilon=0.1$ (right). Top row: for the task of reaching a target while avoiding a forbidden region, Middle row: for the task of reaching the closest goal among two equally rewarding positions, Bottom row: choosing between a small, close reward and a large, distant one..


Figure E.9: Full series of frames in Ms. Pacman along the trajectory generated by the $F^{\top} B$ policy for the task of reaching a target position (star shape $\hat{\sim}$ ) while avoiding forbidden positions (red circle).


Figure E.10: Discrete maze: Visualization of FB embedding vectors after projecting them in twodimensional space with t-SNE. Left: the $F$ embedding for $z=0$. Right: the $B$ embedding. Note how both embeddings recover the foor-room and door structure of the original environment. The spread of B embedding is due to the regularization that makes B closer to orthonormal.


Figure E.11: Continuous maze: Visualization of FB embedding vectors after projecting them in two-dimensional space with t-SNE. Left: the states to be mapped. Middle: the $F$ embedding. Right: the $B$ embedding.


Figure E.12: Ms. Pacman: Visualization of FB embedding vectors after projecting them in two-dimensional space with t-SNE. Left: the agent's position corresponding to the state to be mapped. Middle: the $F$ embedding for $z=0$. Right: the $B$ embedding. Note how both embeddings recover the cycle structure of the environment. F acts on visual inputs and B acts on the agent's position.


Figure E.13: Continuous maze: visualization of $F$ embedding vectors for different $z$ vectors, after projecting them in two-dimensional space with t-SNE.

## Bibliography

Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. In Advances in Neural Information Processing Systems, pages 2312-2320, 2011.
A. Abdolmaleki, J. T. Springenberg, Y. Tassa, R. Munos, N. Heess, and M. A. Riedmiller. Maximum a posteriori policy optimisation. In 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings. OpenReview.net, 2018. URL https://openreview.net/forum?id=S1ANxQW0b.
J. Achiam, D. Held, A. Tamar, and P. Abbeel. Constrained policy optimization. In Proceedings of the 34th International Conference on Machine Learning-Volume 70, pages 22-31. JMLR. org, 2017.
A. Anand, E. Racah, S. Ozair, Y. Bengio, M.-A. Côté, and R. D. Hjelm. Unsupervised state representation learning in atari. arXiv preprint arXiv:1906.08226, 2019.
M. Andrychowicz, D. Crow, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba. Hindsight experience replay. In NIPS, 2017.
K. Asadi, D. Misra, and M. L. Littman. Lipschitz continuity in model-based reinforcement learning. arXiv preprint arXiv:1804.07193, 2018.
P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multiarmed bandit problem. SIAM journal on computing, 32(1):48-77, 2002.
A. Ayoub, Z. Jia, C. Szepesvari, M. Wang, and L. Yang. Model-based reinforcement learning with value-targeted regression. In International Conference on Machine Learning, pages 463-474. PMLR, 2020.
M. G. Azar, V. Gómez, and H. J. Kappen. Dynamic policy programming. Journal of Machine Learning Research, 13(Nov):3207-3245, 2012.
M. G. Azar, A. Lazaric, and E. Brunskill. Online stochastic optimization under correlated bandit feedback. In ICML, pages 1557-1565, 2014.
M. G. Azar, I. Osband, and R. Munos. Minimax regret bounds for reinforcement learning. In Proceedings of the 34th International Conference on Machine LearningVolume 70, pages 263-272. JMLR. org, 2017.
L. Baird et al. Residual algorithms: Reinforcement learning with function approximation. In Proceedings of the twelfth international conference on machine learning, 1995.
A. Barreto, W. Dabney, R. Munos, J. J. Hunt, T. Schaul, D. Silver, and H. P. van Hasselt. Successor features for transfer in reinforcement learning. In NIPS, 2017.
M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: An evaluation platform for general agents. J. Artif. Int. Res., 47 (1):253-279, May 2013a. ISSN 1076-9757.
M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: An evaluation platform for general agents. Journal of Artificial Intelligence Research, 47:253-279, 2013b.
R. Bellman. Dynamic Programming. Princeton University Press, 1957.
M. Benzi and V. Simoncini. On the eigenvalues of a class of saddle point matrices. Numerische Mathematik, 2006.
D. P. Bertsekas. Temporal difference methods for general projected equations. IEEE Transactions on Automatic Control, 2011.
D. P. Bertsekas and J. N. Tsitsiklis. Neuro-dynamic programming: an overview. In Decision and Control, 1995., Proceedings of the 34th IEEE Conference on. IEEE, 1995.
L. Blier, C. Tallec, and Y. Ollivier. Learning successor states and goal-dependent values: A mathematical viewpoint. arXiv preprint arXiv:2101.07123, 2021.
V. I. Bogachev. Measure theory. Springer, 2007.
V. S. Borkar and S. P. Meyn. The o.d.e. method for convergence of stochastic approximation and reinforcement learning. SIAM Journal on Control and Optimization, jan 2000.
D. Borsa, A. Barreto, J. Quan, D. Mankowitz, R. Munos, H. van Hasselt, D. Silver, and T. Schaul. Universal successor features approximators. arXiv preprint arXiv:1812.07626, 2018.
P. Brémaud. Markov chains: Gibbs fields, Monte Carlo simulation, and queues, volume 31. 1999.
G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym, 2016.
S. Bubeck, G. Stoltz, C. Szepesvári, and R. Munos. Online optimization in xarmed bandits. In Advances in Neural Information Processing Systems, pages 201-208, 2009.
Q. Cai, Z. Yang, C. Jin, and Z. Wang. Provably efficient exploration in policy optimization. arXiv preprint arXiv:1912.05830, 2019.
T. Cao and A. Krishnamurthy. Provably adaptive reinforcement learning in metric spaces. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, Advances in Neural Information Processing Systems, 2020.
J. Chen and N. Jiang. Information-theoretic considerations in batch reinforcement learning. In International Conference on Machine Learning, pages 10421051. PMLR, 2019.
Y. Chen, G. Lan, and Y. Ouyang. Optimal primal-dual methods for a class of saddle point problems. SIAM Journal on Optimization, 2014.
W. C. Cheung, D. Simchi-Levi, and R. Zhu. Learning to optimize under nonstationarity. In The 22nd International Conference on Artificial Intelligence and Statistics, pages 1079-1087, 2019.
W. C. Cheung, D. Simchi-Levi, and R. Zhu. Reinforcement learning for nonstationary markov decision processes: The blessing of (more) optimism. arXiv preprint arXiv:2006.14389, 2020.
Y. Chow and H. Teicher. Probability theory: Independence, interchangeability, martingales. Journal of the American Statistical Association, 93, 12 1998. doi: 10.2307/2670078.
B. Dai, A. Shaw, L. Li, L. Xiao, N. He, Z. Liu, J. Chen, and L. Song. Sbeed: Convergent reinforcement learning with nonlinear function approximation. In International Conference on Machine Learning, pages 1125-1134, 2018.
G. Dalal, B. Szorenyi, G. Thoppe, and S. Mannor. Finite sample analysis of twotimescale stochastic approximation with applications to reinforcement learning. arXiv preprint arXiv:1703.05376, 2017.
C. Dann, T. Lattimore, and E. Brunskill. Unifying pac and regret: Uniform pac bounds for episodic reinforcement learning. In NIPS, 2017.
P. Dayan. Improving generalization for temporal difference learning: The successor representation. Neural Computation, 5(4):613-624, 1993.
A. Defazio, F. Bach, and S. Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In Advances in neural information processing systems, 2014.
O. D. Domingues, P. Ménard, M. Pirotta, E. Kaufmann, and M. Valko. A kernelbased approach to non-stationary reinforcement learning in metric spaces. arXiv preprint arXiv:2007.05078, 2020a.
O. D. Domingues, P. Ménard, M. Pirotta, E. Kaufmann, and M. Valko. Regret bounds for kernel-based reinforcement learning. arXiv preprint arXiv:2004.05599, 2020b.
K. Dong, J. Peng, Y. Wang, and Y. Zhou. Root-n-regret for learning in markov decision processes with function approximation and low bellman rank. In Conference on Learning Theory, pages 1554-1557. PMLR, 2020.
M. D. Donsker and S. S. Varadhan. Asymptotic evaluation of certain markov process expectations for large time. iv. Communications on Pure and Applied Mathematics, 36(2):183-212, 1983.
M. P. Drazin. Pseudo-inverses in associative rings and semigroups. The American Mathematical Monthly, aug 1958.
S. S. Du, J. Chen, L. Li, L. Xiao, and D. Zhou. Stochastic variance reduction methods for policy evaluation. In International Conference on Machine Learning, 2017.
S. S. Du, J. D. Lee, G. Mahajan, and R. Wang. Agnostic q-learning with function approximation in deterministic systems: Near-optimal bounds on approximation error and sample complexity. Advances in Neural Information Processing Systems, 33, 2020.
L. Faury, Y. Russac, M. Abeille, and C. Calauzènes. A technical note on nonstationary parametric bandits: Existing mistakes and preliminary solutions. In $A L T, 2021$.
N. Ferns, P. Panangaden, and D. Precup. Metrics for finite markov decision processes. In Proceedings of the 20th conference on Uncertainty in artificial intelligence, pages 162-169. AUAI Press, 2004.
D. Foster and P. Dayan. Structure in the space of value functions. Machine Learning, 49(2):325-346, 2002.
P. Gajane, R. Ortner, and P. Auer. A sliding-window algorithm for markov decision processes with arbitrarily changing rewards and transitions. arXiv preprint arXiv:1805.10066, 2018.
M. Geist, B. Scherrer, and O. Pietquin. A theory of regularized markov decision processes. In International Conference on Machine Learning, pages 2160-2169, 2019.
I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In Advances in neural information processing systems, pages 2672-2680, 2014.
C. Grimm, I. Higgins, A. Barreto, D. Teplyashin, M. Wulfmeier, T. Hertweck, R. Hadsell, and S. Singh. Disentangled cumulants help successor representations transfer to new tasks. arXiv preprint arXiv:1911.10866, 2019.
C. M. Grinstead and J. L. Snell. Introduction to probability. American Mathematical Soc., 1997.
A. Gruslys, W. Dabney, M. G. Azar, B. Piot, M. Bellemare, and R. Munos. The reactor: A fast and sample-efficient actor-critic agent for reinforcement learning. arXiv preprint arXiv:1704.04651, 2017.
T. Haarnoja, H. Tang, P. Abbeel, and S. Levine. Reinforcement learning with deep energy-based policies. In Proceedings of the 34th International Conference on Machine Learning-Volume 70, pages 1352-1361. JMLR. org, 2017.
T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In International Conference on Machine Learning, pages 1861-1870, 2018.
S. Hansen, W. Dabney, A. Barreto, T. Van de Wiele, D. Warde-Farley, and V. Mnih. Fast task inference with variational intrinsic successor features. arXiv preprint arXiv:1906.05030, 2019.
A. Harutyunyan, M. G. Bellemare, T. Stepleton, and R. Munos. $Q(\lambda)$ with offpolicy corrections. In International Conference on Algorithmic Learning Theory. Springer, 2016.
C.-W. Huang, A. Touati, L. Dinh, M. Drozdzal, M. Havaei, L. Charlin, and A. C. Courville. Learnable explicit density for continuous latent space and variational inference. ArXiv, abs/1710.02248, 2017a.
C.-W. Huang, A. Touati, P. Vincent, G. K. Dziugaite, A. Lacoste, and A. Courville. Stochastic neural network with kronecker flow. In International Conference on Artificial Intelligence and Statistics, pages 4184-4194. PMLR, 2020.
G. Huang, H. Berard, A. Touati, G. Gidel, P. Vincent, and S. Lacoste-Julien. Parametric adversarial divergences are good task losses for generative modeling. arXiv preprint arXiv:1708.02511, 2017b.
T. Jaakkola, M. I. Jordan, and S. P. Singh. On the convergence of stochastic iterative dynamic programming algorithms. Neural computation, 6(6):1185-1201, 1994.
T. Jaksch, R. Ortner, and P. Auer. Near-optimal regret bounds for reinforcement learning. Journal of Machine Learning Research, 2010.
Q. Jian, R. Fruit, M. Pirotta, and A. Lazaric. Exploration bonus for regret minimization in discrete and continuous average reward mdps. In Advances in Neural Information Processing Systems, pages 4891-4900, 2019.
N. Jiang, A. Krishnamurthy, A. Agarwal, J. Langford, and R. E. Schapire. Contextual decision processes with low bellman rank are pac-learnable. In International Conference on Machine Learning, pages 1704-1713. PMLR, 2017.
C. Jin, Z. Allen-Zhu, S. Bubeck, and M. I. Jordan. Is q-learning provably efficient? In Advances in Neural Information Processing Systems, pages 4863-4873, 2018.
C. Jin, A. Krishnamurthy, M. Simchowitz, and T. Yu. Reward-free exploration for reinforcement learning. In International Conference on Machine Learning, pages 4870-4879. PMLR, 2020a.
C. Jin, Z. Yang, Z. Wang, and M. I. Jordan. Provably efficient reinforcement learning with linear function approximation. In Conference on Learning Theory, pages 2137-2143. PMLR, 2020b.
R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In Advances in neural information processing systems, 2013.
S. Kakade and J. Langford. Approximately optimal approximate reinforcement learning. In ICML, volume 2, pages 267-274, 2002.
S. Kakade, M. J. Kearns, and J. Langford. Exploration in metric state spaces. In Proceedings of the 20th International Conference on Machine Learning (ICML-03), pages 306-312, 2003.
S. M. Kakade. On the sample complexity of reinforcement learning. PhD thesis, UCL (University College London), 2003.
N. R. Ke, A. Singh, A. Touati, A. Goyal, Y. Bengio, D. Parikh, and D. Batra. Modeling the long term future in model-based reinforcement learning. In International Conference on Learning Representations, 2018.
N. R. Ke, A. Singh, A. Touati, A. Goyal, Y. Bengio, D. Parikh, and D. Batra. Modeling the long term future in model-based reinforcement learning. In ICLR, 2019.
M. Kearns and S. Singh. Near-optimal reinforcement learning in polynomial time. Machine learning, 2002.
J. G. Kemeny and J. L. Snell. Finite Markov Chains. Van Nostrand, New York, 1960.
B. Kim and A. Tewari. Randomized exploration for non-stationary stochastic linear bandits. In Conference on Uncertainty in Artificial Intelligence, pages 7180. PMLR, 2020.
R. Kleinberg, A. Slivkins, and E. Upfal. Multi-armed bandits in metric spaces. In Proceedings of the fortieth annual ACM symposium on Theory of computing, pages 681-690. ACM, 2008.
N. Kodali, J. Abernethy, J. Hays, and Z. Kira. On convergence and stability of gans. arXiv preprint arXiv:1705.07215, 2017.
V. R. Konda and J. N. Tsitsiklis. Actor-critic algorithms. In Advances in neural information processing systems, pages 1008-1014. Citeseer, 2000.
I. Kostrikov. Pytorch implementations of reinforcement learning algorithms. https://github.com/ikostrikov/pytorch-a2c-ppo-acktr-gail, 2018.
I. Kostrikov, O. Nachum, and J. Tompson. Imitation learning via off-policy distribution matching. arXiv preprint arXiv:1912.05032, 2019.
K. Lakshmanan, R. Ortner, and D. Ryabko. Improved regret bounds for undiscounted continuous reinforcement learning. In International Conference on Ma chine Learning, pages 524-532, 2015.
C. Lakshminarayanan and C. Szepesvári. Linear stochastic approximation: Constant step-size and iterate averaging. arXiv preprint arXiv:1709.04073, 2017.
T. Lattimore, M. Hutter, P. Sunehag, et al. The sample-complexity of general reinforcement learning. In Proceedings of the 30th International Conference on Machine Learning. Journal of Machine Learning Research, 2013.
S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. The Journal of Machine Learning Research, 17(1):1334-1373, 2016.
L.-J. Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching. Machine Learning, may 1992.
B. Liu, J. Liu, M. Ghavamzadeh, S. Mahadevan, and M. Petrik. Finite-sample analysis of proximal gradient td algorithms. In UAI. Citeseer, 2015.
S. Liu and H. Su. Regret bounds for discounted mdps. arXiv preprint arXiv:2002.05138, 2020.
Y. Liu, A. Swaminathan, A. Agarwal, and E. Brunskill. Provably good batch reinforcement learning without great exploration. arXiv preprint arXiv:2007.08202, 2020.
S. V. Macua, J. Chen, S. Zazo, and A. H. Sayed. Distributed policy evaluation under multiple behavior strategies. IEEE Transactions on Automatic Control, 2015.
H. R. Maei. Gradient temporal-difference learning algorithms. 2011.
H. R. Maei and R. S. Sutton. Gq ( $\lambda$ ): A general gradient algorithm for temporaldifference prediction learning with eligibility traces. In Proceedings of the Third Conference on Artificial General Intelligence, 2010.
S. Mahadevan and M. Maggioni. Proto-value functions: A laplacian framework for learning representation and control in markov decision processes. Journal of Machine Learning Research, 8(10), 2007.
A. R. Mahmood, H. Yu, and R. S. Sutton. Multi-step off-policy learning without importance sampling ratios. arXiv preprint arXiv:1702.03006, 2017.
V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. nature, 518(7540):529-533, 2015.
V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In International Conference on Machine Learning, 2016.
R. Munos. Error bounds for approximate value iteration. In Proceedings of the National Conference on Artificial Intelligence, volume 20, page 1006. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2005.
R. Munos and C. Szepesvári. Finite-time bounds for fitted value iteration. Journal of Machine Learning Research, 9(5), 2008.
R. Munos, T. Stepleton, A. Harutyunyan, and M. Bellemare. Safe and efficient off-policy reinforcement learning. In Advances in Neural Information Processing Systems, 2016.
R. Munos et al. From bandits to monte-carlo tree search: The optimistic principle applied to optimization and planning. Foundations and Trends® in Machine Learning, 7(1):1-129, 2014.
O. Nachum, Y. Chow, B. Dai, and L. Li. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. In Advances in Neural Information Processing Systems, pages 2315-2325, 2019a.
O. Nachum, B. Dai, I. Kostrikov, Y. Chow, L. Li, and D. Schuurmans. Algaedice: Policy gradient from arbitrary experience. arXiv preprint arXiv:1912.02074, 2019b.
A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. SIAM Journal on optimization, 2009.
G. Neu and J. Olkhovskaya. Online learning in mdps with linear function approximation and bandit feedback. arXiv preprint arXiv:2007.01612, 2020.
G. Neu, A. Jonsson, and V. Gómez. A unified view of entropy-regularized markov decision processes. arXiv preprint arXiv:1705.07798, 2017.
X. Nguyen, M. J. Wainwright, M. I. Jordan, et al. On surrogate loss functions and f-divergences. The Annals of Statistics, 37(2):876-904, 2009.
R. Ortner. Pseudometrics for state aggregation in average reward markov decision processes. In International Conference on Algorithmic Learning Theory, pages 373-387. Springer, 2007.
R. Ortner and D. Ryabko. Online regret bounds for undiscounted continuous reinforcement learning. In Advances in Neural Information Processing Systems, pages 1763-1771, 2012.
R. Ortner, P. Gajane, and P. Auer. Variational regret bounds for reinforcement learning. In UAI, page 16, 2019.
I. Osband, D. Russo, and B. Van Roy. (more) efficient reinforcement learning via posterior sampling. In Advances in Neural Information Processing Systems, 2013.
I. Osband, B. Van Roy, D. J. Russo, and Z. Wen. Deep exploration via randomized value functions. Journal of Machine Learning Research, 20(124):1-62, 2019.
B. Palaniappan and F. Bach. Stochastic variance reduction methods for saddlepoint problems. In Advances in Neural Information Processing Systems, 2016.
J. Pazis and R. Parr. Pac optimal exploration in continuous space markov decision processes. In Twenty-Seventh AAAI Conference on Artificial Intelligence, 2013.
Z. Peng, A. Touati, P. Vincent, and D. Precup. Svrg for policy evaluation with fewer gradient evaluations. arXiv preprint arXiv:1906.03704, 2019.
M. Pirotta, M. Restelli, A. Pecorino, and D. Calandriello. Safe policy iteration. In International Conference on Machine Learning, pages 307-315, 2013.
M. Plappert, M. Andrychowicz, A. Ray, B. McGrew, B. Baker, G. Powell, J. Schneider, J. Tobin, M. Chociej, P. Welinder, et al. Multi-goal reinforcement learning: Challenging robotics environments and request for research. arXiv preprint arXiv:1802.09464, 2018.
D. Pollard. Empirical processes: theory and applications. In NSF-CBMS regional conference series in probability and statistics, pages i-86. JSTOR, 1990.
D. Precup. Eligibility traces for off-policy policy evaluation. Computer Science Department Faculty Publication Series, 2000.
D. Precup, R. S. Sutton, and S. Dasgupta. Off-policy temporal difference learning with function approximation. In Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01, 2001.
M. L. Puterman. Markov decision processes. Handbooks in operations research and management science, 2:331-434, 1990.
M. L. Puterman. Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley \& Sons, Inc., New York, NY, USA, 1st edition, 1994. ISBN 0471619779.
J. Qian, R. Fruit, M. Pirotta, and A. Lazaric. Exploration bonus for regret minimization in undiscounted discrete and continuous markov decision processes. arXiv preprint arXiv:1812.04363, 2018.
P. Rauber, A. Ummadisingu, F. Mutz, and J. Schmidhuber. Hindsight policy gradients. In International Conference on Learning Representations, 2018.
N. Roese. Counterfactual thinking and decision making. Psychonomic bulletin $\mathcal{E}$ review, 6(4):570-578, 1999.
J. Romoff, P. Henderson, A. Touati, E. Brunskill, J. Pineau, and Y. Ollivier. Separating value functions across time-scales. In International Conference on Machine Learning, pages 5468-5477. PMLR, 2019.
J. Romoff, P. Henderson, D. Kanaa, E. Bengio, A. Touati, P.-L. Bacon, and J. Pineau. Tdprop: Does adaptive optimization with jacobi preconditioning help temporal difference learning? In Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems, pages 1082-1090, 2021.
L. Rosasco, S. Villa, and B. C. Vũ. Stochastic forward-backward splitting for monotone inclusions. Journal of Optimization Theory and Applications, 2016.
Y. Russac, C. Vernade, and O. Cappé. Weighted linear bandits for non-stationary environments. In Advances in Neural Information Processing Systems, pages 12017-12026, 2019.
T. Schaul, D. Horgan, K. Gregor, and D. Silver. Universal value function approximators. In F. Bach and D. Blei, editors, Proceedings of the 32nd International Conference on Machine Learning, volume 37 of Proceedings of Machine Learning Research, pages 1312-1320, Lille, France, 07-09 Jul 2015. PMLR. URL http://proceedings.mlr.press/v37/schaul15.html.
J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz. Trust region policy optimization. In International conference on machine learning, pages 1889-1897, 2015.
J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017.
S. E. Shreve and D. P. Bertsekas. Alternative theoretical frameworks for finite horizon discrete-time stochastic optimal control. SIAM Journal on control and optimization, 16(6):953-978, 1978.
S. Sinclair, T. Wang, G. Jain, S. Banerjee, and C. Yu. Adaptive discretization for model-based reinforcement learning. Advances in Neural Information Processing Systems, 33, 2020.
S. R. Sinclair, S. Banerjee, and C. L. Yu. Adaptive discretization for episodic reinforcement learning in metric spaces. arXiv preprint arXiv:1910.08151, 2019.
A. Slivkins. Contextual bandits with similarity information. The Journal of Machine Learning Research, 15(1):2533-2568, 2014.
Z. Song and W. Sun. Efficient model-free reinforcement learning in metric spaces. arXiv preprint arXiv:1905.00475, 2019.
B. K. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf, and G. R. Lanckriet. On integral probability metrics, $\phi$-divergences and binary classification. arXiv preprint arXiv:0901.2698, 2009.
K. L. Stachenfeld, M. M. Botvinick, and S. J. Gershman. The hippocampus as a predictive map. Nature neuroscience, 20(11):1643, 2017.
A. L. Strehl and M. L. Littman. A theoretical analysis of model-based interval estimation. In Proceedings of the 22nd international conference on Machine learning. ACM, 2005.
M. Strens. A bayesian framework for reinforcement learning. In ICML, 2000.
R. S. Sutton. Learning to predict by the methods of temporal differences. Machine learning, 1988.
R. S. Sutton. Introduction to reinforcement learning with function approximation. Tutorial Session of the Neural Information Processing Systems Conference, 2015.
R. S. Sutton and A. G. Barto. Introduction to Reinforcement Learning. MIT Press, Cambridge, MA, USA, 1st edition, 1998. ISBN 0262193981.
R. S. Sutton and A. G. Barto. Reinforcement Learning: An Introduction. MIT Press, Cambridge, MA, USA, 2nd edition, Near-final draft - May 27, 2018.
R. S. Sutton and B. Tanner. Temporal-difference networks. In Advances in Neural Information Processing Systems 17, 2004.
R. S. Sutton, D. A. McAllester, S. P. Singh, Y. Mansour, et al. Policy gradient methods for reinforcement learning with function approximation. In NIPs, volume 99, pages 1057-1063. Citeseer, 1999a.
R. S. Sutton, D. Precup, and S. Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. Artificial intelligence, 112(1-2):181-211, 1999b.
R. S. Sutton, D. Precup, and S. Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. Artificial Intelligence, aug 1999c.
R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. In Advances in neural information processing systems, pages 1057-1063, 2000.
R. S. Sutton, H. R. Maei, D. Precup, S. Bhatnagar, D. Silver, C. Szepesvári, and E. Wiewiora. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In Proceedings of the 26th Annual International Conference on Machine Learning - ICML. ACM Press, 2009a.
R. S. Sutton, H. R. Maei, D. Precup, S. Bhatnagar, D. Silver, C. Szepesvári, and E. Wiewiora. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In Proceedings of the 26th Annual International Conference on Machine Learning. ACM, 2009b.
R. S. Sutton, H. R. Maei, and C. Szepesvári. A convergent $o(n)$ temporaldifference algorithm for off-policy learning with linear function approximation. In Advances in neural information processing systems, 2009c.
R. S. Sutton, J. Modayil, M. Delp, T. Degris, P. M. Pilarski, A. White, and D. Precup. Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS '11, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.
R. S. Sutton, A. R. Mahmood, and M. White. An emphatic approach to the problem of off-policy temporal-difference learning. The Journal of Machine Learning Research, 2015.
E. Talvitie. Self-correcting models for model-based reinforcement learning. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 31, 2017.
W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika, 1933.
A. Touati and Y. Ollivier. Learning one representation to optimize all rewards. arXiv preprint arXiv:2103.07945, 2021.
A. Touati and P. Vincent. Efficient learning in non-stationary linear markov decision processes. arXiv preprint arXiv:2010.12870, 2020a.
A. Touati and P. Vincent. Sharp analysis of smoothed bellman error embedding. arXiv preprint arXiv:2007.03749, 2020b.
A. Touati, P.-L. Bacon, D. Precup, and P. Vincent. Convergent tree backup and retrace with function approximation. In International Conference on Machine Learning, pages 4955-4964. PMLR, 2018.
A. Touati, H. Satija, J. Romoff, J. Pineau, and P. Vincent. Randomized value functions via multiplicative normalizing flows. In Uncertainty in Artificial Intelligence, pages 422-432. PMLR, 2020a.
A. Touati, A. A. Taiga, and M. G. Bellemare. Zooming for efficient model-free reinforcement learning in metric spaces. arXiv preprint arXiv:2003.04069, 2020b.
A. Touati, A. Zhang, J. Pineau, and P. Vincent. Stable policy optimization via off-policy divergence regularization. In Conference on Uncertainty in Artificial Intelligence, pages 1328-1337. PMLR, 2020c.
J. N. Tsitsiklis, B. Van Roy, et al. An analysis of temporal-difference learning with function approximation. IEEE transactions on automatic control, 1997.
L. Van der Maaten and G. Hinton. Visualizing data using t-sne. Journal of machine learning research, 9(11), 2008.
H. van Hasselt, A. R. Mahmood, and R. S. Sutton. Off-policy td ( $\lambda$ ) with a true online equivalence. In Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence, Quebec City, Canada, 2014.
N. Vieillard, O. Pietquin, and M. Geist. Deep conservative policy iteration. arXiv preprint arXiv:1906.09784, 2019.
H.-T. Wai, M. Hong, Z. Yang, Z. Wang, and K. Tang. Variance reduced policy evaluation with smooth function approximation. Advances in Neural Information Processing Systems, 32:5784-5795, 2019.
M. Wang and D. P. Bertsekas. Stabilization of stochastic iterative methods for singular and nearly singular linear systems. Mathematics of Operations Research, 2013.
Q. Wang, Y. Li, J. Xiong, and T. Zhang. Divergence-augmented policy optimization. In Advances in Neural Information Processing Systems, pages 6097-6108, 2019a.
R. Wang, R. R. Salakhutdinov, and L. Yang. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. Advances in Neural Information Processing Systems, 33, 2020.
Y. Wang, W. Chen, Y. Liu, Z.-M. Ma, and T.-Y. Liu. Finite sample analysis of the gtd policy evaluation algorithms in markov setting. In Advances in Neural Information Processing Systems, pages 5510-5519, 2017.
Y. Wang, H. He, and X. Tan. Truly proximal policy optimization. arXiv preprint arXiv:1903.07940, 2019b.
Z. Wang, V. Bapst, N. Heess, V. Mnih, R. Munos, K. Kavukcuoglu, and N. de Freitas. Sample efficient actor-critic with experience replay. arXiv preprint arXiv:1611.01224, 2016.
C. J. Watkins and P. Dayan. Q-learning. Machine learning, 8(3-4):279-292, 1992.
C.-Y. Wei and H. Luo. Non-stationary reinforcement learning without prior knowledge: An optimal black-box approach. In COLT, 2021.
R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. Machine learning, 8(3-4):229-256, 1992.
L. F. Yang and M. Wang. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. arXiv preprint arXiv:1905.10389, 2019.
L. F. Yang, C. Ni, and M. Wang. Learning to control in metric space with optimal regret. arXiv preprint arXiv:1905.01576, 2019.
A. Zanette, D. Brandfonbrener, E. Brunskill, M. Pirotta, and A. Lazaric. Frequentist regret bounds for randomized least-squares value iteration. In International Conference on Artificial Intelligence and Statistics, pages 1954-1964, 2020a.
A. Zanette, A. Lazaric, M. Kochenderfer, and E. Brunskill. Learning near optimal policies with low inherent bellman error. arXiv preprint arXiv:2003.00153, 2020b.
J. Zhang, J. T. Springenberg, J. Boedecker, and W. Burgard. Deep reinforcement learning with successor features for navigation across similar environments. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 2371-2378. IEEE, 2017.
P. Zhao and L. Zhang. Non-stationary linear bandits revisited. arXiv preprint arXiv:2103.05324, 2021.
P. Zhao, L. Zhang, Y. Jiang, and Z.-H. Zhou. A simple approach for nonstationary linear bandits. In International Conference on Artificial Intelligence and Statistics, pages 746-755. PMLR, 2020.
H. Zhou, J. Chen, L. R. Varshney, and A. Jagmohan. Nonstationary reinforcement learning with linear function approximation. arXiv preprint arXiv:2010.04244, 2020.
X. Zhu and D. Dunson. Stochastic lipschitz q-learning. 2019.


[^0]:    ${ }^{1}$ Intuitively, for a metric space to be complete means that there are no "points missing" from it. Bounded real-valued functions can be shown to be complete, and therefore the space of $Q$ functions with $Q$ equipped with the $l_{\infty}$-norm is a complete metric space.

[^1]:    ${ }^{1}$ We overload our notation over linear operators and their corresponding matrix representation.

[^2]:    ${ }^{1}$ The result is not novel, it can be found as intermediate step in proof of theorem 1 in Achiam et al. [2017], for example.

[^3]:    ${ }^{2}$ Both regularized $L_{\pi_{i}}^{\text {clip }}$ and $L_{\pi_{i}}$ are lower bounds on policy performance in Lemma 2 . We

[^4]:    ${ }^{3}$ Code: https:/ / github.com/facebookresearch/ppo-dice

[^5]:    ${ }^{1}$ A concrete case that would satisfy these assumptions, is if $\forall i \in[d], \phi_{i}(s, a) \geq 0$ and $\sum_{i=1}^{d} \phi_{i}(s, a)=1$, and $\forall i \in[d], \mu_{t, h}^{(i)}$ is a probability measure. In this case $\phi(\mathrm{s}, \mathrm{a})$ can be understood as providing the mixture coefficients with which to mix the $d$ measures in $\boldsymbol{\mu}_{t, h}$.

[^6]:    ${ }^{1}$ Here DQN is short for goal-oriented DQN, $Q(s, a, g)$.

[^7]:    ${ }^{2}$ This also holds for continuous rewards, but the Lipschitz assumption yields an explicit bound in Theorem 5.
    ${ }^{3}$ The Kantorovich-Rubinstein norm is closely related to the $L^{1}$ Wasserstein distance on probability distributions, but slightly more general as it does not require the distance functions to be integrable: the Wasserstein distance metrizes weak convergence among those probability measures such that $\mathbb{E}\left[d\left(x, x_{0}\right)\right]<\infty$.

[^8]:    ${ }^{4}$ Intuitively, a parametric function $f$ is overparameterized if every possible small change of $f$ can be realized by a small change of the parameter. Formally, we say that a parametric family of functions $\theta \in \Theta \mapsto f_{\theta} \in L^{2}\left(X, \mathbb{R}^{d}\right)$ smoothly parameterized by $\theta$, on some space $X$, is overparameterized if, for any $\theta$, the differential $\partial_{\theta} f_{\theta}$ is surjective from $\Theta$ to $L^{2}\left(X, \mathbb{R}^{d}\right)$. For finite $X$, this implies that the dimension of $\theta$ is larger than $\# X$. For infinite $X$, this implies that $\operatorname{dim}(\theta)$ is infinite, such as parameterizing functions on $[0 ; 1]$ by their Fourier expansion.

