Université de Montréal

AI-based modeling of brain and behavior: Combining neuroimaging, imitation learning and video games

par

Anirudha Kemtur

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

Mémoire présenté en vue de l'obtention du grade de
Maître ès sciences (M.Sc.)
en Informatique

Juillet 2023

Université de Montréal

Faculté des arts et des sciences

Ce mémoire intitulé

AI-based modeling of brain and behavior: Combining neuroimaging, imitation learning and video games

présenté par

Anirudha Kemtur

a été évalué par un jury composé des personnes suivantes :

Dr. Aaron Courville
(président-rapporteur)

Dr. Pierre Bellec
(directeur de recherche)

Dr. Karim Jerbi
(codirecteur)

Dr. Guillaume Lajoie
(membre du jury)

Résumé

Les récentes avancées dans le domaine de l'intelligence artificielle ont ouvert la voie au développement de nouveaux modèles d'activité cérébrale. Les réseaux neuronaux artificiels (RNA) formés à des tâches complexes, telles que la reconnaissance d'images, peuvent être utilisés pour prédire la dynamique cérébrale en réponse à une série de stimuli avec une précision sans précédent, un processus appelé encodage cérébral. Les jeux vidéo ont fait l'objet d'études approfondies dans le domaine de l'intelligence artificielle, mais n'ont pratiquement pas été utilisés pour l'encodage cérébral. Les jeux vidéo offrent un cadre prometteur pour comprendre l'activité cérébrale dans un environnement riche, engageant et actif, contrairement aux tâches essentiellement passives qui dominent actuellement le domaine, telles que la visualisation d'images. Un défi majeur soulevé par les jeux vidéo complexes est que le comportement individuel est très variable d'un sujet à l'autre, et nous avons émis l'hypothèse que les RNAs doivent prendre en compte le comportement spécifique du sujet afin de capturer correctement les dynamiques cérébrales. Dans cette étude, nous avons cherché à utiliser des RNAs pour modéliser l'imagerie par résonance magnétique fonctionnelle (IRMf) et les données comportementales des participants, que nous avons collectées pendant que les sujets jouaient au jeu vidéo Shinobi III. En utilisant l'apprentissage par imitation, nous avons entraîné un RNA à jouer au jeu vidéo en reproduisant fidèlement le style de jeu unique de chaque participant. Nous avons constaté que les couches cachées de notre modèle d'apprentissage par imitation parvenaient à encoder des représentations neuronales pertinentes pour la tâche et à prédire la dynamique cérébrale individuelle avec une plus grande précision que divers modèles de contrôle, y compris des modèles entraînés sur les actions d'autres sujets. Les corrélations les plus fortes entre les activations des couches cachées et les signaux cérébraux ont été observées dans des zones cérébrales biologiquement plausibles, à savoir les réseaux somatosensoriels, attentionnels et visuels. Nos résultats soulignent le potentiel de la combinaison de l'apprentissage par imitation, de l'imagerie cérébrale et des jeux vidéo pour découvrir des relations spécifiques entre le cerveau et le comportement.

Mots clés: Apprentissage par imitation, Réseaux de neurones artificiels, Codage cerveau, Jeux vidéos, IRMf

Abstract

Recent advances in the field of Artificial Intelligence have paved the way for the development of novel models of brain activity. Artificial Neural networks (ANN) trained on complex tasks, such as image recognition and language processing, can be used to predict brain dynamics in response to wide range of stimuli with unprecedented accuracy, a process called brain encoding. Videogames have been extensively studied in the AI field, but have hardly been used yet for brain encoding. Videogames provide a promising framework to understand brain activity in rich, engaging and active environments, in contrast to mostly passive tasks currently dominating the field, such as image viewing. A major challenge raised by complex videogames is that individual behavior is highly variable across subjects, and we hypothesized that ANNs need to account for subject-specific behavior in order to properly capture brain dynamics. In this study, we aimed to use ANNs to model functional magnetic resonance imaging (fMRI) and behavioral gameplay data, which we collected while subjects played the Shinobi III videogame. Using imitation learning, we trained an ANN to play the game closely replicating the unique gameplay style of individual participants. We found that hidden layers of our imitation learning model successfully encode task-relevant neural representations and predict individual brain dynamics with higher accuracy than various control models, including models trained on other subjects' actions. The highest correlations between layer activations and brain signals were observed in biologically plausible brain areas, i.e. somatosensory, attentional and visual networks. Our results highlight the potential of combining imitation learning, brain imaging, and videogames to uncover subject-specific relationships between brain and behavior.

Keywords: Imitation Learning, Artificial Neural Networks, Brain encoding, Videogames, fMRI

Contents

List of tables

List of figures

List of acronyms and abbreviations

IL	Imitation Learning
RL	Reinforcement Learning
ANN	Artificial Neural Networks
fMRI	Functional magnetic resonance imaging
CNN	Convolutional neural network
RNN	Recurrent Neural Network
LSTM	Long-short term Memory Network
BOLD	Blood Oxygen Level Dependent
PCA	Principal component analysis

Gratitude

I would like to sincerely thank my supervisors Dr. Pierre Bellec and Dr. Karim Jerbi for giving me the opportunity to work on this exciting project. They have constantly supported me, patiently discussed ideas, provided insightful insights, and most importantly taught me how to do good research.

I would also like to thank all my lab members for their invaluable feedback and critical comments on the project. Not only they all have helped me to understand and do research in a better way, but also as friends have constantly supported me and made my stay here a memorable experience.

I am grateful to my family members for their constant love, support, and encouragement throughout my career.

Thank you.

Chapter 1

Introduction

1.1. Overview

The brain is a remarkable and intricate system consisting of countless interconnected neurons. These neurons work together in various organized components to produce specific processes, which in turn lead to our thoughts and actions. Human intelligence is mainly attributed to the brain, and throughout history, people have been fascinated by understanding how the brain supports intelligent behavior. In the past, this endeavor was mostly a topic of philosophical speculation, but today, with advanced neuroimaging technology, we can record brain activity to quantitatively model its structure and function. Neuroimaging data have enabled substantial advancements in our comprehension of brain organization and in the diagnosis of certain medical conditions. Such progress has been made possible due to the development of a wide array of computational modeling tools for neuroimaging.

Over the past ten years, deep learning has become the dominant force in the field of Artificial Intelligence (AI) and is playing an increasingly central role in the analysis of neuroimaging data. Artificial neural networks (ANNs) have gained prominence due to their unprecedented performance on challenging tasks. This success is based on the ability of ANNs to learn from data-driven training methods, allowing them to construct effective representations of input data through various computational hierarchies. By learning meaningful transformations of inputs that correspond to the tasks on which they were trained, these networks can go beyond simply making predictions for the given task, offering broader applications for downstream tasks.

Since its inception, AI has looked to neuroscience for inspiration, particularly in the subfield of machine learning (ML). As the brain is the only example of a physical system that displays human-level intelligence, it's natural to try to replicate some of its design principles in order to achieve human-level AI. This influence is most noticeable in artificial neural networks, such as the convolution filters used in convolutional neural networks (CNNs). These filters imitate the functional properties of neurons in the visual cortex and were instrumental in the deep learning breakthrough of 2012, particularly with the creation of AlexNet

[**38**]. Additionally, long-short term memory networks (LSTMs) [**29**], which are based on the concept of working memory, have demonstrated remarkable performance in many tasks involving time series, including natural language processing.

Even though the brain has inspired a lot of deep learning techniques, the parallel between brain mechanisms and ANNs is often abstract, with many key differences in implementation and architectures. Comparing how brains and deep learning models are similar/ dissimilar in their computations, internal representations, and responses might help us close this gap, understand our brains better and at the same time lead to better AI models. Such a comparison often takes the form of an encoding analysis: a computational model is first trained to do a task of interest, and then that model is used to predict the brain responses while performing the same task. Many brain encoding studies have successfully applied deep neural networks to model the functioning of different brain functional systems, such as vision [**81, 30**], audition [**35**], and language [**10**], leading to a better understanding of the neural basis of these cognitive functions.

These previous works have focused mainly on sensory processing and much less on decision-making and complex, integrative behavior. Hence the tasks studied and stimuli used have been mostly passive, consisting in looking at pictures, videos or listening to stories. These tasks are far from the interactive environment the brain normally operates in. Therefore, to properly study and compare brains and ANNs, it might be very important to consider experiments involving integration between different sensory modalities, decision-making as well as agency in the environment. Videogames have been shown to strongly engage participants' attention, emotions, motor and decision-making abilities [**3, 5, 56**] and hence serve as a rich framework to understand brain dynamics. Recent work showed that ANNs can be used as an encoding model for complex tasks like videogames [**15**]. Cross et.al. used a reinforcement learning model to perform encoding on functional neuroimaging data acquired while participants played Atari games. They used the encoding model behaviorally was not similar to the subject's gameplay, which might have been sufficient for simple atari videogames.

In this study, we set out to use ANNs to model brain and behavioral data simultaneously, while subjects do naturalistic tasks such as playing a complex videogame. We first investigate if we can use imitation learning to model behavioral data, by closely replicating the unique gameplay style of individual participants. Further, we see if such a model, trained to imitate subject-specific behavior, is a better encoding model of the subject's brain dynamics than models trained to replicate the game style of other subjects.

The initial sections of this thesis aim to equip the reader with a basic understanding of the methodological building blocks of this research work. In this introductory chapter, we first introduce functional magnetic resonance imaging which is the form of brain data used in this study. Then we introduce the deep learning model architectures and learning techniques (Reinforcement learning and Imitation learning) which were used in training our behavioral model. Further, we discuss how representations of AI models and brains can be studied using encoding models, and present an overview of related works in this area. In the end, we briefly describe the Courtois-Neuromod databank which provided the data for this study.
In chapter two, we present our modeling and analysis as a scientific article.

1.2. Background

1.2.1. Human brain imaging

This part of the paper provides an overview of fundamental concepts and vocabulary used in functional magnetic resonance imaging (fMRI) to help understand the neuroimaging data that we use to model in our work.

1.2.1.1. Functional Neuroimaging

Functional neuroimaging refers to a group of techniques that employ neuroimaging to observe brain activity within specific regions, with the goal of comprehending the mechanisms underlying various cognitive functions. These approaches that capture data pertaining to neural activity in the brain are at the center of research methods used to advance our understanding of the brain.

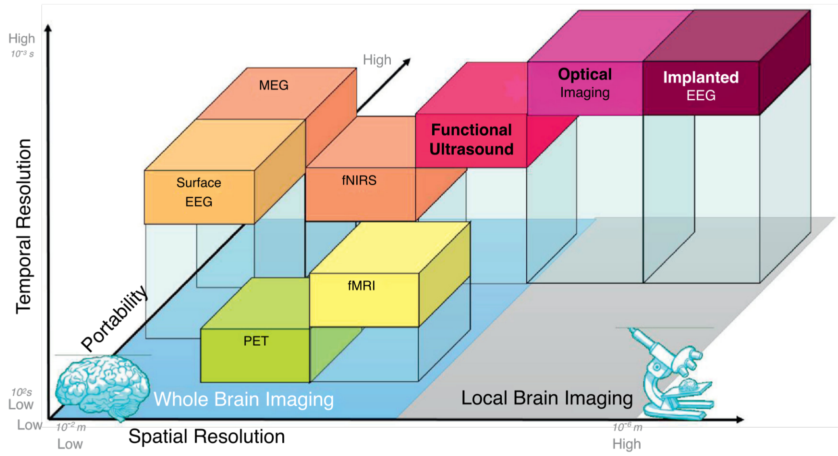

Fig. 1.1. An overview of various functional neuroimaging techniques depicted to indicate their spatial and temporal resolutions along with the portability of the recording technique (from *Deffieux et al.* [**16**] - CC BY 4.0)

Functional neuroimaging encompasses a range of techniques, including functional magnetic resonance imaging (fMRI), electroencephalography (EEG), and magnetoencephalography (MEG), that are used to record and analyze brain activity data in order to gain insights into cognitive functions. These quantitative methods rely on statistical and computational analyses, which presents opportunities for contributions from computer scientists to advance the field.

As depicted in Figure 1.1, various functional neuroimaging methods are situated along a spectrum of portability and spatial and temporal resolution. For instance, MEG signals offer high temporal resolution , whereas fMRI provides higher spatial resolution. Notably, each technique has its own strengths and limitations and may be useful in different contexts.

fMRI has been extensively utilized and has played a dominant role in brain research in recent years, owing to its non-invasive nature that poses no harm or adverse effects to participants, and its ability to model the entire brain at a reasonable temporal resolution.

1.2.1.2. Magnetic Resonance Imaging (MRI)

Some nuclei that contain an odd number of protons (such as water) align along an external constant magnetic field. Under this condition, they possess a spin property that causes them to emit a characteristic electromagnetic signal when they are perturbed by a weak magnetic field oscillating at a certain specific frequency characteristic of the nuclei. The protons in these kinds of nuclei are always spinning about their axis producing a net magnetic moment along the direction of the axis of the spins. This happens when this characteristic frequency matches that of the perturbing magnetic pulse, causing a resonance effect that is referred to as Nuclear Magnetic Resonance (NMR).

Fig. 1.2. A view of a MRI scanner showing the table extended out of the bore (attribution: Liz West from Boxborough, MA - CC BY 2.0)

Magnetic Resonance Imaging (MRI) uses this principle of NMR to measure the net magnetization of all such nuclei within a given space. A powerful magnet in the bore of an MRI scanner, as illustrated in Figure 1.2, generates a magnetic field around the participant lying on a table. Radio-frequency (RF) coils are utilized to excite the nuclei at specific locations, causing them to emit a radio frequency signal, also known as an echo pulse, which is detected by another coil. The frequency of the received signal at each location is translated into an intensity value that is mapped to pixels, producing an image.

During the experiment, the Time to Echo (TE) is a setting that designates the duration between the transmission of the RF pulse and the receipt of the echo signal. Different types of tissues produce diverse contrasts in the resulting image, as tissues that require imaging have varying characteristic frequencies at which NMR can occur.

1.2.1.3. Blood-oxygen-level-dependent (BOLD) Signal

The firing activity of neurons in the brain is dependent on oxygen in the blood. When neurons are stimulated, they fire at a higher rate and require more energy. The body responds by supplying more oxygen to these active neurons via the hemodynamic response, which causes a change in the proportions of oxyhemoglobin and deoxyhemoglobin in the blood vessels. This change in proportions can be used as a proxy to measure neuronal activity.

Seiji Ogawa and his colleagues discovered that there is a difference in the magnetic properties of oxygenated and deoxygenated hemoglobin [**53**] and were able to detect this signal variation in an MRI scanner. This signal is known as the Blood Oxygen Level Dependent (BOLD) signal. The BOLD fMRI signal represents the proportion of oxygenated and deoxygenated hemoglobin in the blood and is observed in areas of the brain where neurons are active. Therefore, the BOLD signal can be considered as stimuli-driven activations.

It has been empirically established that the BOLD response is indicative of an increase in neuronal activity [**51**]. The hemodynamics phenomena is characterized by a *hemodynamic*

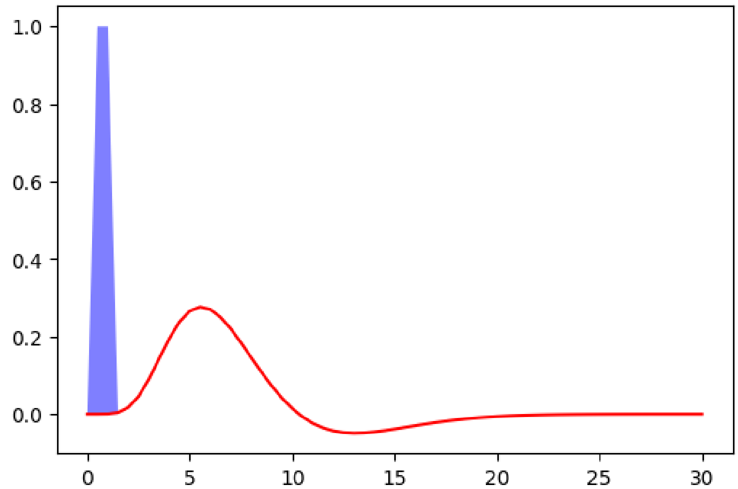

Fig. 1.3. Hemodynamic response to a unit pulse lasting one second

response function (HRF) as seen in Figure 1.3 where it has a slow activation peak and eventual fall. It can be seen that a stimulus that lasts for a brief period of few seconds acts as an impulse to produce the response lasting till 25 s. It can be viewed as a general model of the BOLD response to an impulse neural input (triggered by the stimulus).

1.2.1.4. fMRI Experiment

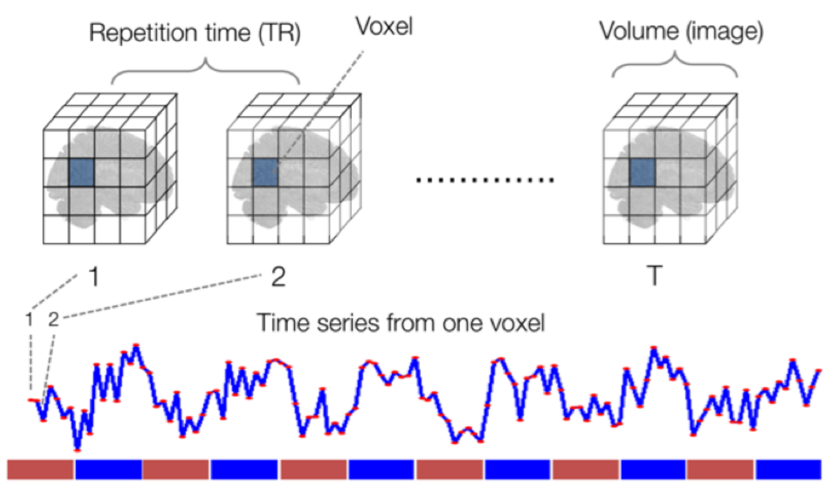

Fig. 1.4. An illustration of the BOLD time series corresponding to a single voxel from the recorded spatial volumes across multiple time steps (TRs). The colorbar indicates different task conditions denoted by red and blue. (*© 2015 Tor D. Wager and Martin A. Lindquist* [**42**])

In an fMRI experiment, the participant lies on a table inside the MRI scanner with their head placed in a head coil. The participant is provided with a screen and other necessary tools to perform the task. Each recording session is a single stretch of recording on a given day and is divided into a fixed number of runs, which are composed of many blocks of different task and rest periods. The participant is presented with stimuli and/or asked to perform certain actions at certain time intervals within each block. At fixed time intervals, called the repetition time (TR), a three-dimensional image of the BOLD activity across the entire volume of the brain is captured by the scanner. The images are acquired as axial slices of certain thickness known as the slice thickness. The vertical extent of the brain present inside the image is called the field of view (FOV), and the number of grids in the axial slice images is called the matrix size. The individual units of recorded data are called voxels, with their intensities representing the strength of the BOLD signal at that location. The slice thickness, FOV, and matrix size determine the voxel dimensions (usually in mm). At the end of each fMRI data recording session, the output produced is a time-series of image volumes representing the intensities for each voxel that is part of the three-dimensional brain volume. This is shown in Figure 1.4.

The main purpose of recording the fMRI data during experiments is to analyze the data and explain brain function and behavior. One key goal of fMRI data analysis is the localization of brain regions that are active during a specific task, to understand their cognitive function.

1.2.1.5. fMRI data Preprocessing

The fMRI data collected from brain scans is typically not immediately suitable for analysis. If the data does not meet certain criteria, it can make many statistical methods unusable. Therefore, a series of processing steps must be taken to eliminate false artifacts, ensure certain statistical assumptions, and standardize the spatial locations of brain regions. This is an important step in guaranteeing the accuracy of the recorded data, and some of the necessary steps can be performed using software packages like fMRIPrep [**22, 21**]. However, these steps must be carefully chosen and executed properly to ensure that the data is useful. Not all steps may be necessary in every case, and some may require modifications. While there is an extensive list of these steps, we will briefly summarize some important ones below.

Slice-time Correction

It is important to make sure that all voxels in a particular brain volume were recorded at the same time. However, in some cases, the activity is recorded as 2D axial slices of the brain volume with some delay between them, which can cause issues. This can be resolved by interpolating the values for a common time point using the entire time course. However, if multiple slices can be acquired simultaneously in a fast manner, this step may not be necessary.

Motion Correction

The assumption that signals recorded at a particular voxel represent only its actual location is not always true, particularly when there is head motion during or between scan sessions. To maintain the consistency of voxel values, a rigid body transformation is applied to align each volume and remove any motion effects. This is done using a reference volume in the time course, typically the first. [**32**]

Susceptibility Distortion Correction

The images produced by the scanner can sometimes be distorted due to the magnetic field not being uniform. This can cause errors in the way the values are mapped to the voxel locations. To counter this, an inhomogeneity map of the field is created, which helps to adjust the mapping of locations and calculate the displacements of the voxels to reduce the effect of these distortions.

Co-registration

The process of registration involves aligning an image to fit into another space. In the case of fMRI BOLD images, they may not be properly aligned with the structural features of the brain. To solve this problem, anatomical images are used as they have clear boundaries. As a result, functional volumes are mapped to the anatomical image and the fMRI signal is projected onto the surfaces that are generated in the anatomical images using a regression.

Spatial Normalization

The brain shapes and sizes vary between individuals, and to make data analysis easier, it is helpful to use a common space with fixed dimensions that maps to the same anatomical or functional structures. To achieve this, spatial normalization is performed, which involves warping the data of each subject into a standard template space, such as the one developed by the Montreal Neurological Institute[**24**]. This step involves a warping algorithm that uses a complicated nonlinear normalization with a large set of parameters. Nonlinear algorithms based on diffeomorphic registration, which is an invertible transform that maps from subject space to template space and back, have been successful. Advanced Normalization Tools (ANTs)[**4**] algorithm is an example of such algorithms.

Smoothing

It is a common practice to apply a spatial smoothing kernel to the fMRI BOLD data in order spread out the intensities in the neighborhood of voxels. This improves the process of registration by eliminating intra-subject differences. A full width of the kernel at half its maximum height (FWHM) is used for this process.

Confounds Removal

The BOLD signal that is measured can potentially be influenced by various sources such as physiological (breathing, cardiac activity, etc.), hardware (coil heating up, frame displacement etc.) or other sources that contribute to the global noise signal. Due to this, the signal has a small amplitude and can be confounded. In order to overcome this issue, a set of confound time series are identified as noise components through various methods such as ICA or CompCor [**6**]. These nuisance parameters are then removed by performing regression to eliminate their effects.

1.2.1.6. Spatial Orientations

The analysis of fMRI data often generates maps of the brain that display statistical values for each location in space. Since the brain is a complex three-dimensional structure, it is helpful to have a shared understanding of spatial orientations and names for different regions. This section will introduce the fundamental terminology used for this purpose.

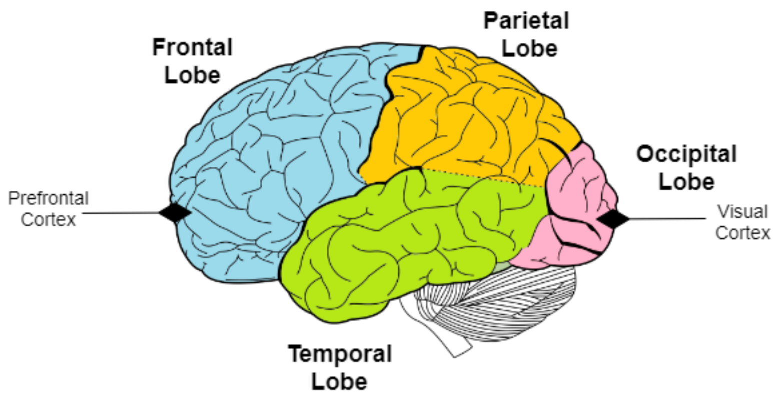

Fig. 1.5. A diagram of the human brain showing the different lobes and some important cortices. The 4 lobes of the brain are occipital (pink), parietal (orange), temporal (green) and frontal (blue). The two important cortices (among many other) that are highlighted are the prefrontal coretex located in the frontal lobe and the visual cortex in the occipital lobe (Colored and labelled on a sketch by Henry Vandyke Carter)

The brain is composed of several distinct lobes that contain specialized regions for different cognitive functions. These lobes include the occipital, parietal, temporal, and frontal lobes as shown in Figure 1.5 . In addition to these structures, the brain surface also includes curved shapes that can be used to locate areas on the surface, with each outward folding referred to as a gyrus and each inward fold referred to as a sulcus.

When displaying results in two dimensions, cross-sectional planes are used to represent the values for each spatial location in the volume along the three directions. The planes used have specific names and orientations with respect to the overall head and brain surface.

To describe different directions (axes) and planes in the three dimensions of the space representing the brain, specific names are given. The X axis is the left-to-right dimension of the participant, while the Y axis spans from the back of the brain (posterior) to the front (anterior). The Z axis is the bottom-to-top dimension, extending from the bottom (inferior) to the top (superior) locations. These locations can also be referred to as ventral (top-to-bottom) and dorsal (bottom-to-top).

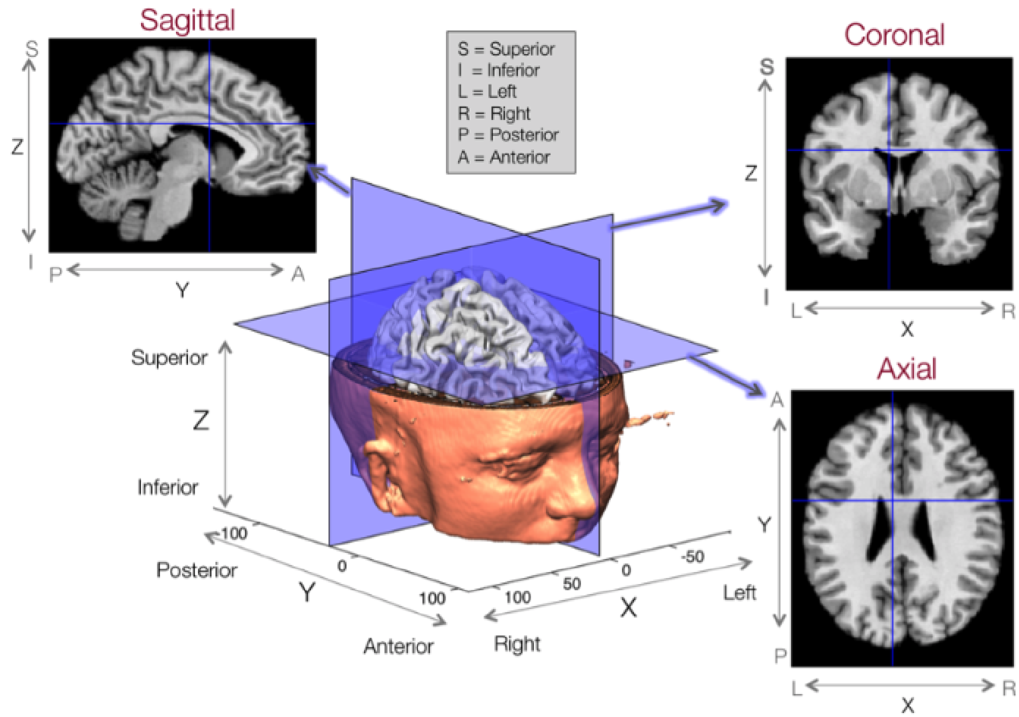

Fig. 1.6. A depiction of the three types of cross-sectional planes in a 3D brain volume, along with the associated axes indicating the various terms associated with the directions in brain anatomy (legend in gray box in the top). The images corresponding to the three cross-sectional planes - sagittal, coronal and axial are also shown. (*© 2015 Tor D. Wager and Martin A. Lindquist* [**42**])

Referring to Figure 1.6 and Figure 1.5 might be helpful to connect specific vocabulary related to identifying locations or orientations within the brain.

1.2.2. AI Techniques

In recent years, deep learning (DL) has proven to be one of the most successful machine learning methods for various prediction problems, with specific applications in computer vision, natural language understanding as well as in the area of neuroscience [**61**]. DL methods allow the identification of optimal discriminative patterns in a given minimally pre-processed dataset, thus reducing the reliance on a priori selection of features.

1.2.2.1. Convolutional Neural networks

CNNs are a class of DL architectures used to model image data. An example CNN model is shown in Figure 1.7. CNNs, owing to their ability to learn representations that are invariant to affine transformations, have revolutionized the field of computer vision [**40**].

There are mainly three types of functions used in a CNN. The main component is the Conv layer which performs convolution operations using kernels whose parameters can be

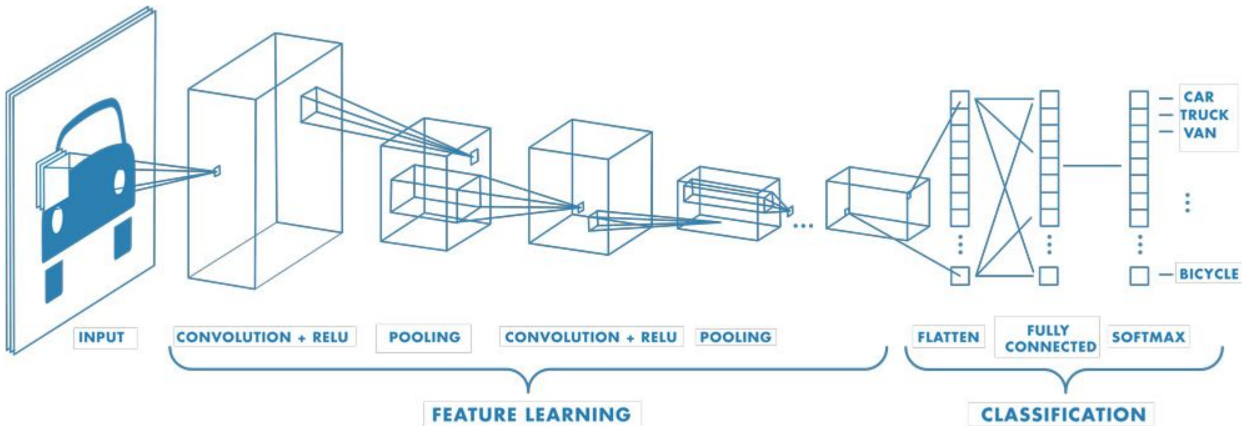

Fig. 1.7. Example of a network with many convolutional layers. Filters are applied to each training image at different resolutions, and the output of each convolved image is used as the input to the next layer. Figure from Syulistyo et. al.[**60**] - CC BY 3.0.

learned from the data. These kernels can be thought of as feature identifiers, and they are able to capture more and more complex features as we stack more layers. The spatial pooling layer helps the network to reduce the size of the feature representation, thereby reducing its dimensionality. Commonly used is MaxPool pooling operation that calculates the maximum value in each patch of each feature map. ReLU Layer adds in the non-linearity that might be required to model the data [**52**]. Normally these three are stacked multiple times, before flattening the feature map to a 1-D vector. Finally, it is common to have Fully connected layers in which every neuron unit is connected to all the neurons in the subsequent layer.

1.2.2.2. Recurrent Neural Network

Recurrent neural networks (RNNs) are distinct from CNN/feed-forward neural networks in that they include extra self or temporal connections in the hidden layers. These connections link the same hidden layers from a previous time step, resulting in recurrence in these networks. This characteristic allows RNNs to model temporal dynamics of the inputs that feed-forward networks cannot.

The Long Short-Term Memory (LSTM) is a type of Recurrent Neural Network (RNN) that is more advanced and can consolidate information over longer periods. It stores "long term" memory in a memory cell vector, which is denoted as c_t, in addition to the output hidden state, denoted as h_t, which together form the complete state of each LSTM cell (Figure 1.8).
While there are various LSTM architectures with different connectivity structures and activation functions, they all utilize memory cells to store information for extended periods. The LSTM can overwrite, retrieve, or maintain the memory cell for the next time step. The flow of information within the network is managed by a mechanism known as gating, which

employs additional parameters to operate gates that compute additional variables. A typical LSTM cell with three gates - forget gate, input gate, and output gate. The forget gate determines which portion of the information from the previously remembered past (previous cell state c_{t-1}) should be further remembered, the input gate controls which portion of the current input (x_t) should be retained, and the output gate controls which portion of the presently remembered information (cell state c_t) should be output as the hidden state h_t.

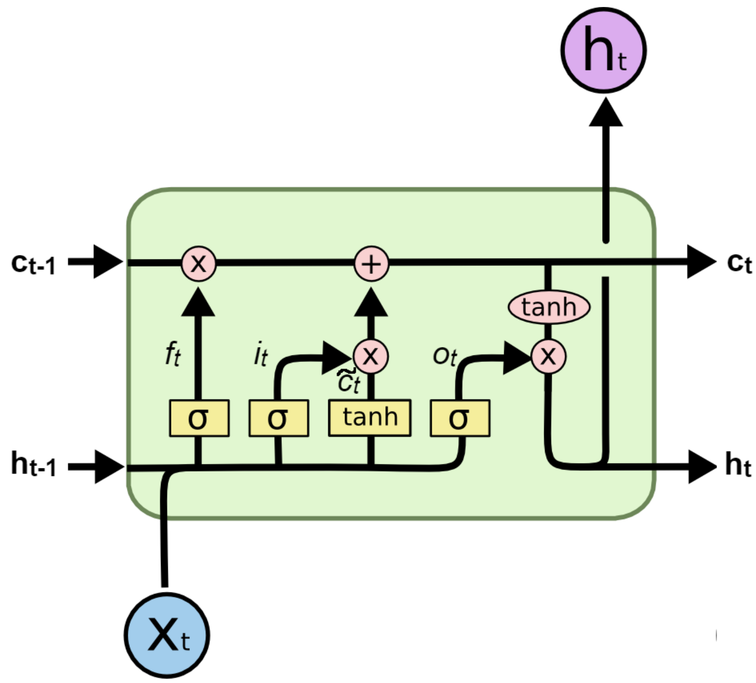

Fig. 1.8. Schematic diagram of a simple LSTM cell [**54**]

1.2.2.3. Reinforcement Learning

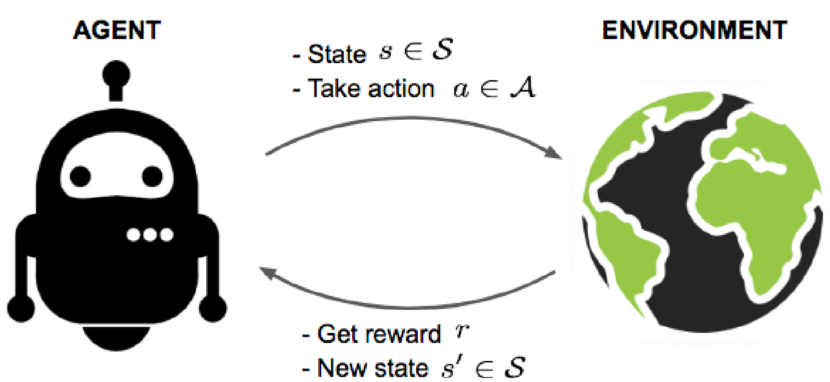

Fig. 1.9. Typical RL scenario: An agent interacts with the environment, trying to take smart actions to maximize cumulative rewards [**79**]

Reinforcement Learning (RL) is a machine learning technique commonly applied to design ML models when raw data is not labeled but we have access to an environment where we can take actions and get rewards . They are generally applied to the class of problems that come under Markov Decision Processes(MDP). MDP has set of states(S), a set of actions(A), a transition function(T: S * A -> prob(S)), and a reward function R(for each S * A). In an RL setting as shown in Figure 1.9 , an agent learns by interacting with the environment. An agent takes actions that are applied to the environment, and the environment in turn returns the reward and the next environment state. Given such a setting, algorithms such as Q-Learning [**12**] can find effective solutions to it.

Earlier most successful RL applications required carefully handcrafted features and relied on linear policy(function from S -> A) representations. With the advances in Deep Learning(DL), supervised learning applications were being solved with very high success rates. But that required a lot of labeled data, which was not available for RL problems.

The first work to bring the power of DL into RL was from Deepmind and this set off major advances in the field. Playing Atari with Deep Reinforcement Learning [**46**] demonstrated that Convolutional neural networks(CNN's) could successfully overcome the challenges of non-labeled data, delayed reward, etc. The paper's main contribution was to combine DL with traditional Q-learning (thus coming up with Deep Q-Learning Networks -DQN). The authors were able to train the same algorithm on seven different games and the network performed exceedingly well (SOTA on most games).

After this paper, there have been many other improvements to the algorithms by using recurrent models in the function (DRQN [**27**]), Having two Q networks to increase the stability of training [**74**], etc. These RL algorithms estimate the value of the state. Another popular class RL algorithms are those that directly learn the Policy(what action to take given a state) rather than learning the value functions of each S * A, come under Policy gradient approaches. Actor critic models like A3C [**45**] or algorithms like PPO [**66**] are a few examples of methods that directly learn the policy.

RL can directly interact with the environment and figure out good policies. But this has a few limitations, the main one being the design of reward functions. This designing of the reward function to get the desired behavior can be very tricky and sometimes not possible at all.

1.2.2.4. Imitation Learning

In the last section, we saw that designing a reward function for RL can be challenging. Hence one approach to follow that enables humans to learn: is to learn by seeing others. So a solution could be to get an expert to do the desired behavior and make our policy imitate the expert. This would fall under the category of machine learning algorithms known as Imitation learning (IL) [**87**].

The simplest and one of the earliest approaches to IL is Behaviour Cloning (BC). BC aims to learn the expert's policy by treating it as a supervised learning problem. So given the expert state-action pairs, we treat them as i.i.d examples and apply normal supervised learning by minimizing the loss function.

The classic paper that introduced this method is ALVINN: An Autonomous Land Vehicle In A Neural Network [**58**] published in 1988. In this paper, they used a simple 2-layer neural network as the policy network and taught the network to map road images to steering angles. They not only used the expert driver's data but also created a road simulator to convert the input images to different simulated conditions. Though due to a lack of computing resources then, the speed was limited and the test road was very constrained (eg. no traffic, humans, etc) this work was definitely inspirational to recent successes in IL.

The BC approach though very simple suffers from a common problem. If we encounter a state that is different from the expert data, then the network might not know what is the right action and it action might result in a state that is further away from those seen in expert training data. This might lead to undefined and harmful behaviors. But in applications where an expert's data can cover the state space, and where an error might not be that harmful, BC would be a very effective model.

An extension of BC is have a demonstrator that encodes the expert's policy rather than having an expert's state-action pairs explicitly. Hence the training algorithm can query the demonstrator and get the optimal action for that state. One of the popular methods of this type is DAGGER [**62**]. Though this helps us solve the problem of BC if the expert is good, the main hindrance would be that an expert that can evaluate agents at all times might not be available.

Another major approach to IL is Inverse Reinforcement Learning (IRL) [**1**]. Here we learn the reward function that underlies expert behavior. This makes some assumptions about the reward function and its parameters. Now given an initial reward function, we train an agent with RL to maximize that reward function and get the policy. Further, we see how the behavior of the policy differs from the expert and update the reward function. This continues till the policies become similar. Here as RL optimization is present in the inner loop, which leads to very high training times.

1.2.3. Neuro-AI

In the next sections, we discuss the potential benefits of utilizing computational models, specifically deep neural networks, in generating useful representations. Further, we explain how encoding analysis can compare these representations to those found in the brain.

1.2.3.1. Representations of AI Models and Brains

Deep neural networks are computational models that use optimized parameters to transform data into useful outputs. The network's hidden layers learn representations from the input data, which are referred to as features because they highlight important aspects of the input data. This eliminates the need for explicit feature engineering, making trained deep neural networks very useful for extracting features from inputs that can be used for further downstream tasks [**83, 84**].

The brain can be considered as a system that generates responses when stimulated with specific input conditions, just like computational models. The brain's activity can be measured through various means like electrical, magnetic, and vascular activity, while its behavior can be observed through actions and decisions in the real world. By comparing the brain's activity patterns with representations from other deep neural networks or computational models, we can identify similarities between them. To do so, we need standard methods to compare two representations of the same or similar inputs.

While there are a number of ways to compare representations across the Deep neural network and brain activity [**18**], the method we use here is based on an encoding approach which we summarise in the next section.

1.2.3.2. Encoding analysis

The field of Neuroscience has extensively researched Encoding models[**49, 77**], which provide valuable insights into the brain's behavior and function. The process involves constructing an encoding model to estimate the similarity of representations. These models predict how the brain responds to specific input stimuli or data. The approach relies on studying the mapping of the DNN feature space to the brain activity space to reveal the similarity between the two representations.

Learning from data to map one vector space to another can be seen as a regression problem. A range of statistical learning methods, as discussed in [**70**], can be utilized for this purpose. However, using a linear model for this task is particularly advantageous as it is a well-established method, as supported by various studies such as [**49, 44, 80, 34**]. This is because a linear relationship is the simplest kind of relationship that can exist between two vector spaces, making it easier to identify similarities. This would enable us to gauge how similar the feature space is to the brain space. Additionally, linear models have a lower capacity, which means they require less data to fit the model.

Regularized linear models have been found to be better for linear encoding, as shown by [**18**]. These models have been used in several works [**31, 76, 35**].

Once the regression has been performed, the next step is to use metrics to assess the linear predictive model's generalization using validation strategies. Pearson's correlation coefficient, coefficient of determination, and mean absolute error are commonly used metrics to evaluate the model's performance. The score obtained on the validation set is a measure of the similarity between the representations.

1.2.3.3. Related works

Brain encoding and images

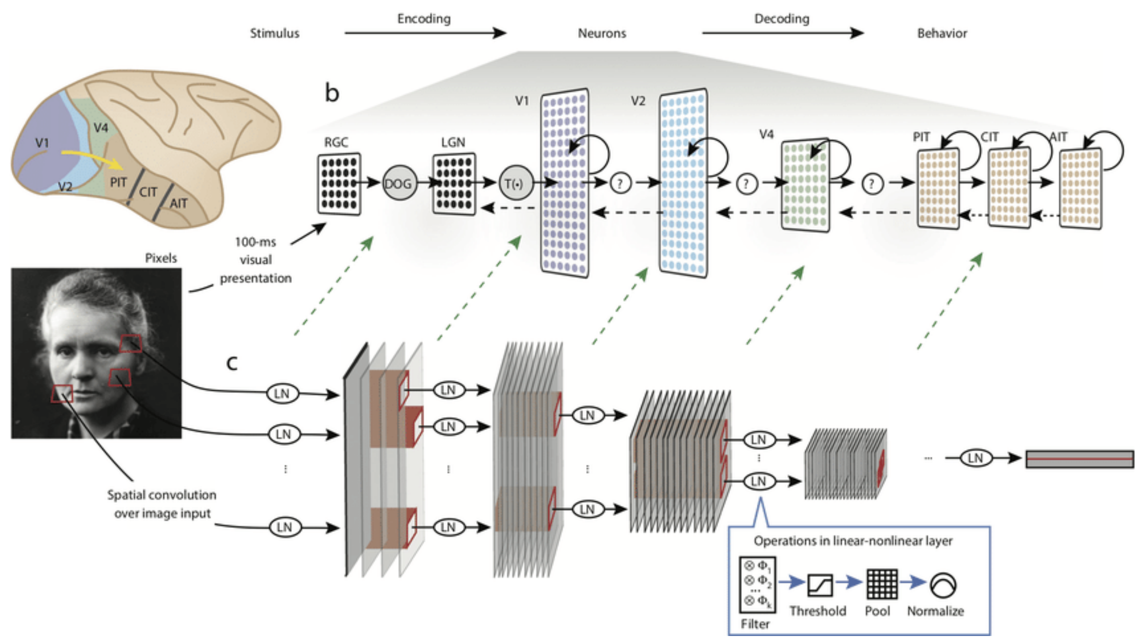

Fig. 1.10. Comparing the visual system and deep convolutional neural networks. The same image can be passed through the visual cortex (top) and a deep convolutional neural network (bottom), allowing for side-by-side comparisons between biological and artificial neural networks. Figure used from Yamins and DiCarlo [**81**].

Previous studies have found functional similarities between artificial neural networks (ANNs) and the brain. In particular, when both the visual cortex and a CNN trained on object recognition are presented with the same image, there is a linear correlation between the activation patterns of the neurons in the CNN and those measured in the visual cortex using fMRI or MEG. This correlation enables the use of CNN activations to predict the brain's response to stimuli in specific regions, which is known as brain encoding [**82, 13, 20**]. The general principle of a brain encoding task from CNN activation is depicted in Figure 1.10.

CNNs can encode brain activity, which suggests that even though the brain and ANNs learn differently, they come up with similar solutions to solve tasks like object recognition. For instance, both artificial and biological networks extract similar features from images to arrive at comparable solutions in vision-related tasks. The extracted features resemble Gabor filter-like patterns, which are low-level features that cortical columns in V1 are believed to be sensitive to, according to current theoretical models of visual receptive fields. The feature maps learned in the early layers of CNNs visually resemble these low-level features as shown in Figure 1.11.

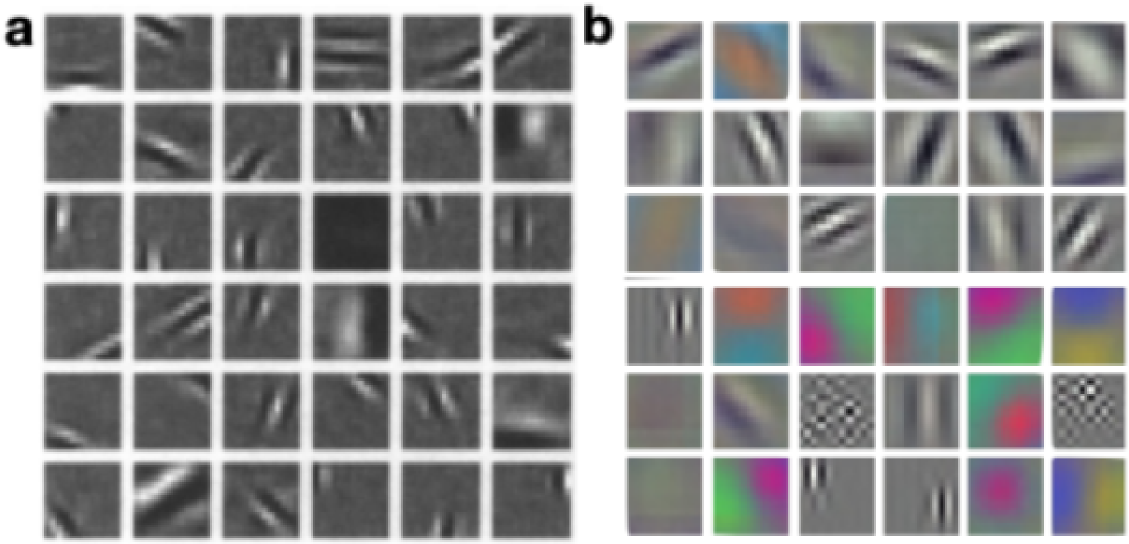

Fig. 1.11. Similarity between low level brain visual features and features captured in initial layers of CNN. (a) Gabor filters used for an analytical model of V1 by Olshausen and Field [**55**]. (b) Kernels learned by the first layer of AlexNet [**28**].

CNN activations also encode brain activity in areas of the brain linked to higher-level features, and they exhibit a hierarchical organization, where deeper layers of the CNN correspond more closely with voxels that are associated with higher-level features [**20**]. There is even more similarity in the way features are embedded in artificial and biological neural networks, as evidenced by the similar specialization that occurs when networks are trained on multiple tasks. Dobs et al. [**19**] demonstrated this with CNNs trained to recognize both objects and faces, which displayed functional segregation similar to the spatial specialization in the visual cortex. In addition to the organizational similarity, complex behavioral similarity has been observed in PredNets, which are CNNs trained on predicting the next frame and have been observed to react to visual illusions in ways that are similar to how the brain processes them [**43**].

Brain encoding and language

Studies using fMRI and EEG for brain encoding tasks have found significant similarities in brain activation between Artificial Neural Networks (ANNs) and the human brain, not only for visual tasks but also for natural language processing (NLP) tasks [**64, 67, 9**]. These findings suggest that deep NLP models have internal representations of words that are similar to those of the brain. Although the learning process of ANNs is different from that of the brain, both converge towards similar word embedding. The models' architecture is partly responsible for this similarity, as even models with random weights can predict a significant

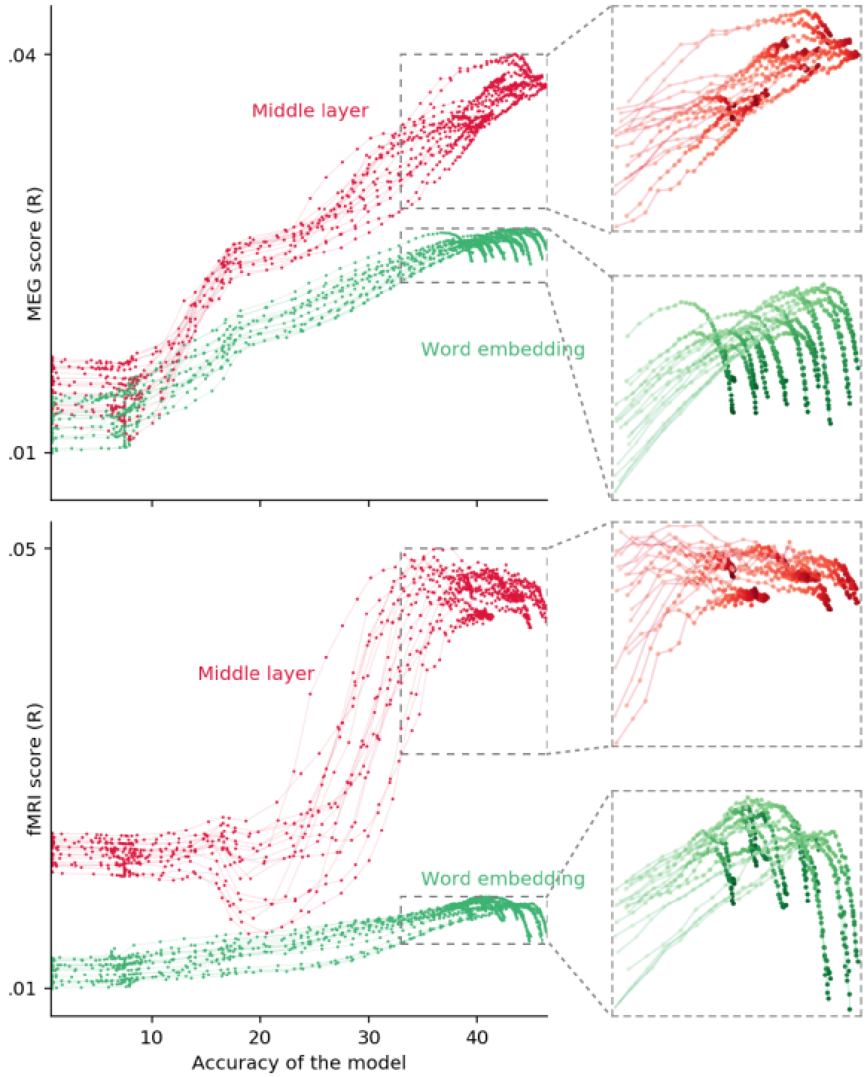

Fig. 1.12. Correlation between task performance and brain encoding accuracy for NLP transformers. Plots adapted from Caucheteux and King (2021)[**9**] show the MEG(Top) and fMRI(Bottom) encoding score of various NLP transformers as a function of the accuracy of the models on the next-word prediction training task. Green dots represent the word embeddings (first layer) of the models, and red dots represent the activation of the middle layer of the models, which was found to be the best layer for brain encoding.

portion of the brain's response to words. However, trained models perform better than untrained models, indicating that the training process increases the brain similarity of the models. Moreover, the accuracy of brain encoding by the models was positively correlated with their performance in their training task of next-word prediction. This correlation is shown in Figure 1.12.

Brain encoding and video games

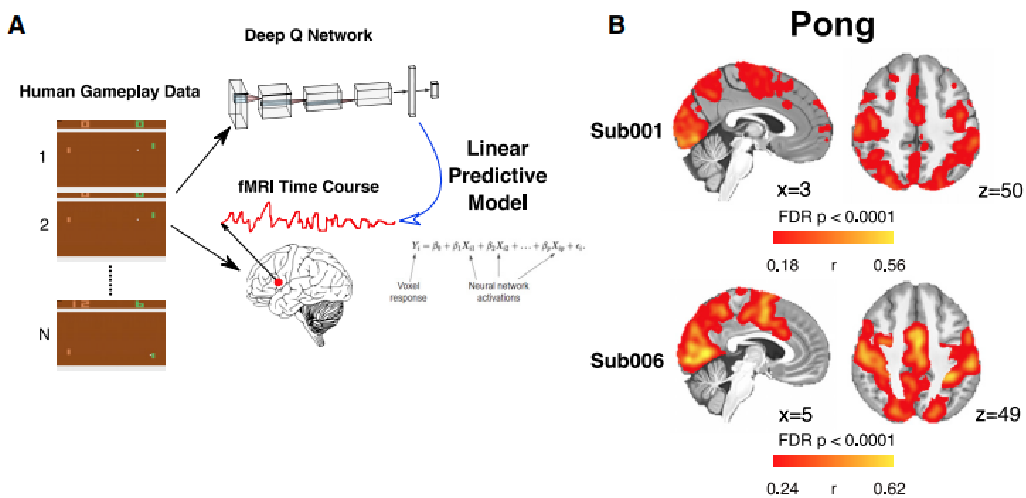

Fig. 1.13. Figure from Cross et. al. [**14**]. (A) Visualization of encoding model analysis. (B) Pearson correlation between the predicted and actual voxel responses for Pong videogame.

The paper "Using deep reinforcement learning to reveal how the brain encodes abstract state-space representations in high-dimensional environments" from Cross et. al.(2020) [**14**], is quite closely related to our study. This study trains DQN model on the atari game and compares it to the brain activations of subjects playing the same game inside an fMRI scanner.
Three games: Pong, Enduro and Space invaders and record 4.5 hrs of data per subject. They also trained 3 DQN networks in each of the games. However, one significant difference from the approach we are trying to take is that the authors here as don't train algorithms that imitate the subject. Hence for comparing the activations with the brain data they only consider the actions where the subject and the DQN network take the same action. Further, they successfully predict the brain voxel data from the trained DQN. It was found that the DQN model outperformed all the control models (that used basic visual input / game actions) in predicting brain responses.

1.2.4. Brain data: the Courtois-Neuromod datasets

This thesis uses the shinobi dataset, which was obtained from the CNeuroMod databank [**8**]. The primary aim of the CNeuroMod initiative is to create a vast and diverse neuroimaging data repository for individual subjects. This resource can facilitate the study of brain activity through the use of integrative artificial intelligence models, while also enabling investigation into the scaling properties of these models. The CNeuroMod project collects hundreds of hours of neuroimaging scans for each subject, including both structural and functional MRI scans, as well as various biophysiological measures, during naturalistic and controlled stimuli. It is one of the most extensive individual fMRI datasets currently available, encompassing diverse cognitive domains such as movie watching, video game playing, and functional tasks developed for the Human Connectome Project. The CNeuroMod project's objective is to provide a distinctive neuroimaging resource that contains vast amounts of data at an individual level and covers a wide range of cognitive domains.

Acquiring fMRI data can be expensive, but luckily for this Masters thesis, the Courtois-Neuromod project has already gathered the largest individual fMRI datasets and will continue to collect more data for several years. The Courtois neuromod project has specifically collected the largest individual video game datasets. This dataset includes fMRI data from four subjects who played the game Shinobi III: Return of the Ninja Master for around 10 hours each, and the project is currently in the process of collecting fMRI data from five subjects who will play Super Mario Bros for 15 hours each. Unlike typical datasets in cognitive neuroscience which involve many subjects but only a short amount of recording per subject, this dataset allows us to build subject-specific models by using each subject's fMRI data independently. Further details on the shinobi dataset used in this study is described in Section 2.1

Chapter 2

Article.

AI-based modeling of brain and behavior: Combining neuroimaging, imitation learning and video games

by

Anirudha Kemtur[1], Francois Paugam[2], Basile Pinsard[3], Pravish Sainath[4], Maximilien Le Clei [5], Julie Boyle[6], Karim Jerbi[7], and Pierre Bellec[8]

([1]) Computer Science Dept., Université de Montréal
Mila – Quebec AI Institute
Centre de Recherche de l'institut Universitaire de Gériatrie de Montréal (CRIUGM)

([2]) Computer Science Dept., Université de Montréal
Mila – Quebec AI Institute
Centre de Recherche de l'institut Universitaire de Gériatrie de Montréal (CRIUGM)

([3]) Centre de Recherche de l'institut Universitaire de Gériatrie de Montréal (CRIUGM)

([4]) Computer Science Dept., Université de Montréal
Mila – Quebec AI Institute
Centre de Recherche de l'institut Universitaire de Gériatrie de Montréal (CRIUGM)

([5]) Centre de Recherche de l'institut Universitaire de Gériatrie de Montréal (CRIUGM)

([6]) Centre de Recherche de l'institut Universitaire de Gériatrie de Montréal (CRIUGM)

([7]) Computer Science Dept., Université de Montréal
Psychology Dept., Université de Montréal
Mila – Quebec AI Institute
Centre de Recherche de l'institut Universitaire de Gériatrie de Montréal (CRIUGM)

([8]) Computer Science Dept., Université de Montréal
Psychology Dept., Université de Montréal
Centre de Recherche de l'institut Universitaire de Gériatrie de Montréal (CRIUGM)

Short version of this article was accepted and presented at Conference on Cognitive Computational Neuroscience 2022. The full version will be submitted soon to eLife journal.
.

The majority of the work described in this article was done by Anirudha Kemtur.
Dr. Basile Pinsard helped with the preprocessing of fMRI data.
Francois Paugam, Pravish Sainath, Maximilien Le Clei and Julie Boyle provided useful feedback throughout the study.
Dr. Karim Jerbi and Dr. Pierre Bellec supervised the work, provided the initial idea for the project, and the data in addition to mentoring and continued feedback.

Abstract.

Recent advances in the field of Artificial Intelligence have paved the way for the development of novel models of brain activity. Artificial Neural networks (ANN) trained on complex tasks, such as image recognition and language processing, can be used to predict brain dynamics in response to a range of stimuli with unprecedented accuracy, a process called brain encoding. Videogames have been extensively studied in the AI field, but have hardly been used yet for brain encoding. Videogames provide a promising framework to understand brain activity in a rich, engaging, and active environment, in contrast to mostly passive tasks currently dominating the field, such as image viewing. A major challenge raised by complex videogames is that individual behavior is highly variable across subjects, and we hypothesized that ANNs need to account for subject-specific behavior in order to properly capture brain dynamics. In this study, we aimed to use ANNs to model functional magnetic resonance imaging (fMRI) and behavioral gameplay data, which we collected while subjects played the Shinobi III videogame. Using imitation learning, we trained an ANN to play the game closely replicating the unique gameplay style of individual participants. We found that hidden layers of our imitation learning model successfully encode task relevant neural representations and predict individual brain dynamics with higher accuracy than various control models, including models trained on other subjects' actions. The highest correlations between layer activations and brain signals were observed in biologically plausible brain areas, i.e. somatosensory, attentional, and visual networks. Our results highlight the potential of combining imitation learning, brain imaging, and videogames to uncover subject-specific relationships between brain and behavior.

Keywords: Imitation Learning, Artificial Neural Networks ,Brain encoding , Videogames, fMRI

1. Introduction

Artificial Neural Networks (ANNs) trained with deep learning have emerged as good encoding models of brain activity in varied cognitive domains such as visual object recognition [**17**], audition [**35**], and natural language processing [**10**]. Previous applications of ANNs have focused mainly on passive sensory processing and much less on active tasks involving decision-making. Agency in the environment is a key aspect of human behavior, which needs to be incorporated in this nascent field using ANNs for brain encoding [**85**]. Most videogames strongly engage attention, emotions, along with motor and decision-making abilities and hence show a lot of promise to model brain dynamics [**7**]. A recent work notably showed that ANNs are suitable for brain encoding in classic Atari video games [**15**]. The study by Cross and colleagues however relied on a generic reinforcement learning model whose behavior was not specifically tuned to individual subject gameplay, which may have been adequate for relatively simple games. In this work, we aimed to study whether ANNs trained to closely imitate individual gameplay are better brain encoding models for complex video game environments.

Even though the brain has inspired a lot of deep learning techniques, the parallel between brain mechanisms and ANNs is often abstract, which is partly due to key differences in implementation and architecture. Comparing how brains and deep learning models are similar/ dissimilar in their computations, internal representations, and responses might help us close this gap, understand our brains better and at the same time lead to better AI models [**26**]. Such a comparison often takes the form of an encoding analysis: a computational model is first trained to do a task of interest using a rich sensory data stream as input (e.g. annotation of objects from pixel-level image data), and then the same model is used to predict the brain responses while performing the same task [**65**]. Many brain encoding studies have successfully applied deep neural networks to model the functioning of different brain functional systems, such as vision [**81, 30**], audition [**35**], and language [**10**], leading to a better understanding of the neural basis of these cognitive functions.

In most of the previous works the tasks studied and stimuli used have been mostly passive, consisting in looking at pictures and videos, or listening to stories. These tasks are far from the interactive environment the brain normally operates in. Videogames have been shown to strongly engage participants' attention, emotions, motor and decision-making abilities [**3, 5, 56**] and hence serve as a rich framework to understand brain dynamics. Embodied cognition is a key aspect of human behavior, and likely a critical step to building informative models of the human brain, as was recently proposed by a consortium of researchers calling for an embodied Turing test [**26**].

Recent work showed that ANNs can be used as an encoding model for complex tasks like videogames [**15**] where a general reinforcement learning model was used to model brain and behavior. Human behavior in video gameplay can be very complex and their style very subject-specific. Hence it might be interesting to study if models that closely replicate the subject's behavior might serve as better individual encoding models, especially in complex and interactive environments.

In this work, we start by first examining if we can train individual imitation models to behaviorally replicate subject gameplay style. A successfully trained imitation learning model should not only be able to play the game but play it in the style whose gameplay it was trained upon. Next, we see if this model that learns the behavior can also encode the brain data of the subject recorded while playing the video game. We hypothesize that an ANN model that was trained to plan and play the game like that subject should serve as a good encoding model of the brain dynamics of that subject.

Further, we wanted to see if this behavior and brain models are subject-specific. To investigate this we first see how a model trained on the gameplay of an individual behaviorally compares to a model with different gameplay styles. If the gameplay style of the subjects is indeed subject-specific and our imitation learning model is able to capture the behavior well, then we expect that the model trained on a particular subjects gameplay style is able to play the game more similarly to that subject when compared to a model trained on different subjects gameplay. Next, we test if this subject-specific behavioral model is better at predicting the brain data of that subject than a model that plays the game with a different gameplay style. We hypothesize that if imitation learning models the behavior of that subject better, it should serve as a better encoding model of that subject brain than an ANN trained to model different subjects' behavior.

Our last aim was to understand whether individual brain encoding models are driven mainly by simple, low-level features of the game. We thus compared the imitation model's brain encoding with a series of control models designed to make sure that our model is not just capturing variance that could be explained with low-level pixel/motor output features.

2. Materials and Methods

The summary of our experiment pipeline is shown in Figure 0.1.

In the sections below, we describe in detail, the dataset and its collection (STEP-1 in Figure 0.1), followed by a description of the Imitation learning neural network and its training (STEP-2) and then how the brain encoding was performed (STEP-3).

2.1. Dataset

The shinobi dataset used in this project was obtained from the Courtois-Neuromod databank [**8**]. the shinobi dataset consists of both fMRI and behavioral data acquired while participants played the videogame Shinobi III: Return of the Ninja Master (Sega, 1993). Data

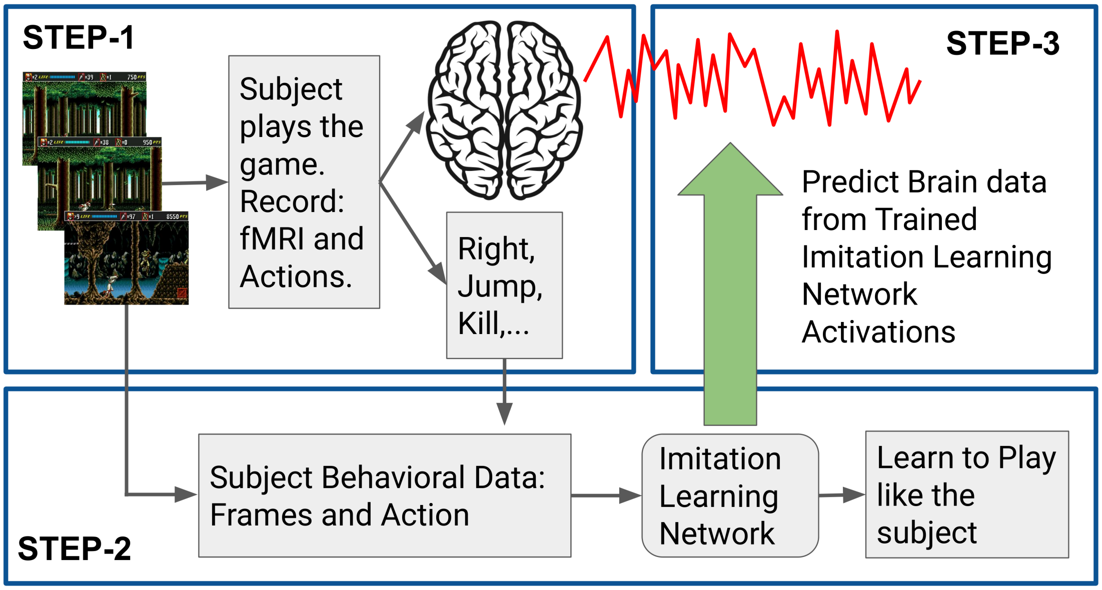

Fig. 0.1. Experiment Pipeline. **Step-1**: Subject's videogame frames, fMRI data, and key presses are recorded. **Step-2**: Train an imitation learning neural network using frames and action data, such that it learns to play with a gameplay style similar to the subject. **Step-3**: Use activations from the neural network model to linearly predict the fMRI data.

was collected on four right-handed subjects (2 women) aged from 33 to 49 years old during acquisition.

2.1.1. Shinobi videogame

In this task, the subjects had to play Shinobi III: Return of the Ninja Master, an action-scrolling platformer in which the player incarnates a ninja trying to free their village from a fictitious group of mercenaries. We selected three levels of the game for their gameplay similarity; levels 1, 4, and 5. In these levels, a player starts on the far left of the map and has to reach and complete a boss-fight located at the far right of the map. To do so, a player can move, crouch, jump, and hit enemies while collecting bonuses and avoiding various obstacles. The full completion of a level usually takes about 2 to 3 minutes, and a repetition of the level was considered complete when a) a player successfully beat the boss of that level, or b) a player dies.

The game was emulated using OpenAI's gym-retro [**50**], a toolbox that allowed the extraction of game frames and memory state variables, and presented using PsychoPy [**57**].

2.1.2. Experiment setup

In phase one of the experiment participants were first instructed to play outside of the MRI scanner our version of the Shinobi video game using a computer monitor or TV, and an

Intel NUC computer equipped with a commercially available replica of a SNES controller. Subjects played until they reached a proficient level of skill, determined using performance metrics such as maximum score reached and speed of level completion.

Once subjects reached the proficient level of play they were invited to our MRI facility to play more sessions of the shinobi task for phase two. We collected behavioral data (videogame frames and controller key presses) and fMRI BOLD activation data as they played in the scanner. Around 10 sessions per subject were recorded with each session lasting for about 40 minutes and composed of multiple runs. One run of the Shinobi task comprises sequential repetitions of levels 1, 4, and 5 (in this order) looping back to level 1 when the first three levels are completed. The runs lasted for at least 10 minutes and ended when the current repetition ended after the 10 minutes were elapsed.

For this project, only Phase two data (behavioral and fMRI) was used. Few runs and sessions were not usable based on quality checks and hence were excluded. The data used per subject is shown in Table 0.1.

Table 0.1. Available data

Subject	Total sessions	Total runs	Data (in hours)
01	12	53	7.20
02	8	37	5.41
04	8	40	6.02
06	10	43	6.35

2.1.3. fMRI data

MRI data was collected using a 3T Siemens Prima Fit MR scanner, a 64-channel head coil. fMRI data was collected using an accelerated simultaneous multi-slice imaging sequence [**68**] with a spatial resolution of 2mm isotropic, and a Repetition Time (TR) of 1.49s. All subjects wore custom CaseForge headcases to minimize motion. Visual stimuli were projected onto a screen via a waveguide and presented to participants on a mirror mounted on the head coil, and the sound was delivered using S15 Sensimetrics headphone inserts. Structural scans were collected using a T1-weighted MPRAGE 3D sagittal sequence, (TR = 2.4 s, TE = 2.2 ms, flip angle = 8 deg, voxel size = 0.8 mm). For a more detailed description of the MRI or fMRI sequences and set up are described at Courtois Neuromod's Project documentation page: `https://docs.cneuromod.ca`.

2.1.4. Data Preprocessing

Behavioral data was recorded at 60Hz. However, it was observed that humans play the game at about 10 to 15Hz due to which each controller key press was repeated 4 to 6 times in our recordings. Hence, we downsampled the Behavioral data to 12Hz. This resulted in having 18 videogame frames + controller key presses for every 1.49 seconds (Duration of each fMRI TR). Videogame frames were converted from RGB to grayscale images. Further, we also cropped out the top portion of the videogame frame where the game score is displayed.

fMRI data was preprocessed using the fMRIprep pipeline[**23**], using the "long term support" version 20.2.3. A high pass filter of 0.01 Hertz and Spatial smoothing of fwhm = 8mm was applied to the data. Further was further cleaned by using [motion + wm_csf + global] confounds using load_confounds package. Finally, voxel space fMRI data was converted into parcel space using MIST- 444 atlas [**73**].

2.2. Imitation learning

Reinforcement Learning (RL) is a machine learning technique commonly applied to design ML models when raw data is not labeled but we have access to an environment where we can take actions and get rewards. They are generally applied to the class of problems that come under Markov Decision Processes (MDP). MDP has a set of states (S), a set of actions (A), a transition function (T: S * A -> prob(S)), and a reward function R (for each S * A). An agent takes actions that are applied to the environment, and the environment in turn returns the reward and the next environment state. Given such a setting, algorithms such as Q-Learning [**12**] can find effective solutions such that the actions the agent takes maximize the cumulative reward.

RL can directly interact with the environment and figure out good policies. But this has a few limitations, first, when the environment rewards are very sparse, we might need to manually define the reward function. This designing of the reward function to get the desired behavior can be very tricky and sometimes not possible at all.

One approach is to follow what enables humans to learn, to learn by seeing others. So a solution could be to get an expert to do the desired behavior and make our policy imitate the expert. This would fall under the category of RL algorithms known as Imitation learning (IL).

In this study, we want to see if training a network to play with the gameplay style of the subject better encodes the subject's brain data, and the fact that we have subjects' gameplay trajectories available, we use imitation learning to train our behavior encoding model.

2.2.1. Model architecture

We use a Deep neural network-based imitation learning model. The network architecture of our model is shown in Figure 0.2. As the task here is to take in a videogame frame and learn which action to take such that it imitated the subject's actions, the model needs to process images as well as keep track of its actions taken at the previous timestep. Hence we use Convolutional neural network (CNN) layers as well as Long short-term memory (LSTM) layers in our architecture.

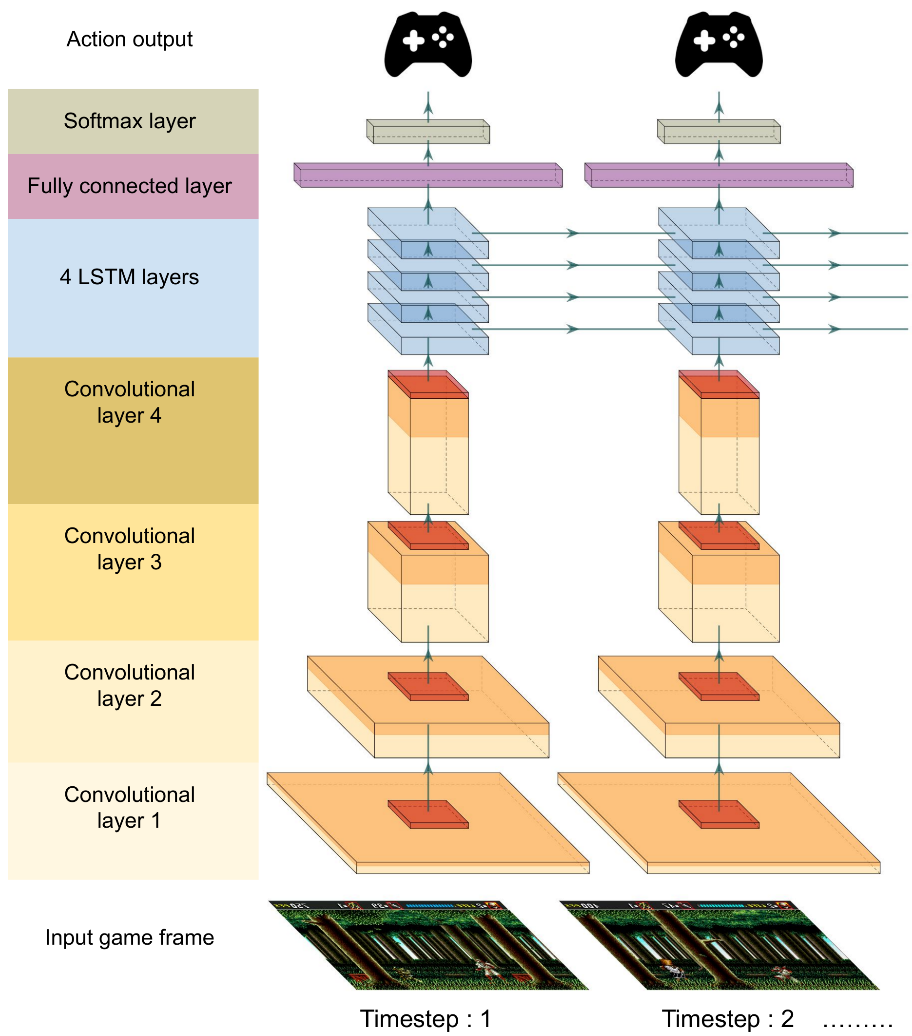

Fig. 0.2. Model Architecture composed of four Convolutional layers, four LSTM layers, fully connected layer to output softmax probabilities over the action space.

CNNs are one of the best architectures to classify image/video data mainly due to their ability to extract representations that are robust to partial translation and deformation of input patterns [**39**]. LSTM layers serve as a memory unit that enables the network to learn the temporal nature of gameplay as complex strategies require a careful sequence of particular actions.

Our imitation learning model is an ANN consisting of four stacked Convolutional layers, followed by four long short-term memory (LSTM) layers. Finally, the output from the LSTM is passed into a fully connected (FC) layer that outputs softmax probabilities over the action space. The architecture details are shown in Table 0.2.

Layer	Type	Filter Size	Number of filters	Output Size
0	Input	-	-	[1,50,100]
1	Conv+Relu	[5,5]	32	[32,16,32]
2	Conv+Relu	[5,5]	64	[64,12,14]
3	Conv+Relu	[5,5]	256	[256,4,5]
4	Conv+Relu	[4,5]	200	[200,1,1]
8	Flatten	-	-	[200]
9	4 LSTM Layers	-	-	[200]
10	Fully connected layer	-	-	[64]
11	Softmax layer	-	-	[1]

Table 0.2. Architecture details

2.2.2. Training

Behaviour Cloning The simplest and one of the earliest approaches to IL is Behaviour Cloning (BC). BC aims to learn the expert's policy by treating it as a supervised learning problem. So given the expert state-action pairs, we treat them as i.i.d examples and apply supervised learning by minimizing the loss between the model's predicted action and the subject's recorded action for that game frame.

Loss function There is a huge class imbalance in the action sequences, as going right is the most frequent step (due to the nature of a right-scroller game). This leads to a class imbalance that cannot be removed by removing some percent of "right" actions from the data, as a gameplay trajectory requires every action in a particular sequence, and removing some actions would result in an entirely different trajectory. Hence we used a weighted negative log-likelihood loss function (based on the frequency of each action in the data) so that if it makes errors on frequent actions the penalty is lesser than if it makes mistakes on infrequent actions. This allows the model to focus more on the infrequent actions as well while maintaining the original game trajectory.

Implementation The model was implemented using PyTorch 1.6 in Python 3.6. The loss was then backpropagated through the model and weights were updated using the Adam optimizer [**37**]. Early stopping [**86**] was employed to prevent overfitting and the best model was stored for validation. Hyperparameter selection is one of the important steps in obtaining a good deep-learning model. We carried out an exhaustive search for the optimal hyperparameters following well-accepted principles and approaches in the field [**69**]. The hyperparameter values found are provided in Table 0.3.

Hyperparameter	Value
LSTM sequence length	18
Batch size	50
Training epochs	3000
Learning rate	0.00005

Table 0.3. Training Hyperparameters

2.3. Brain encoding

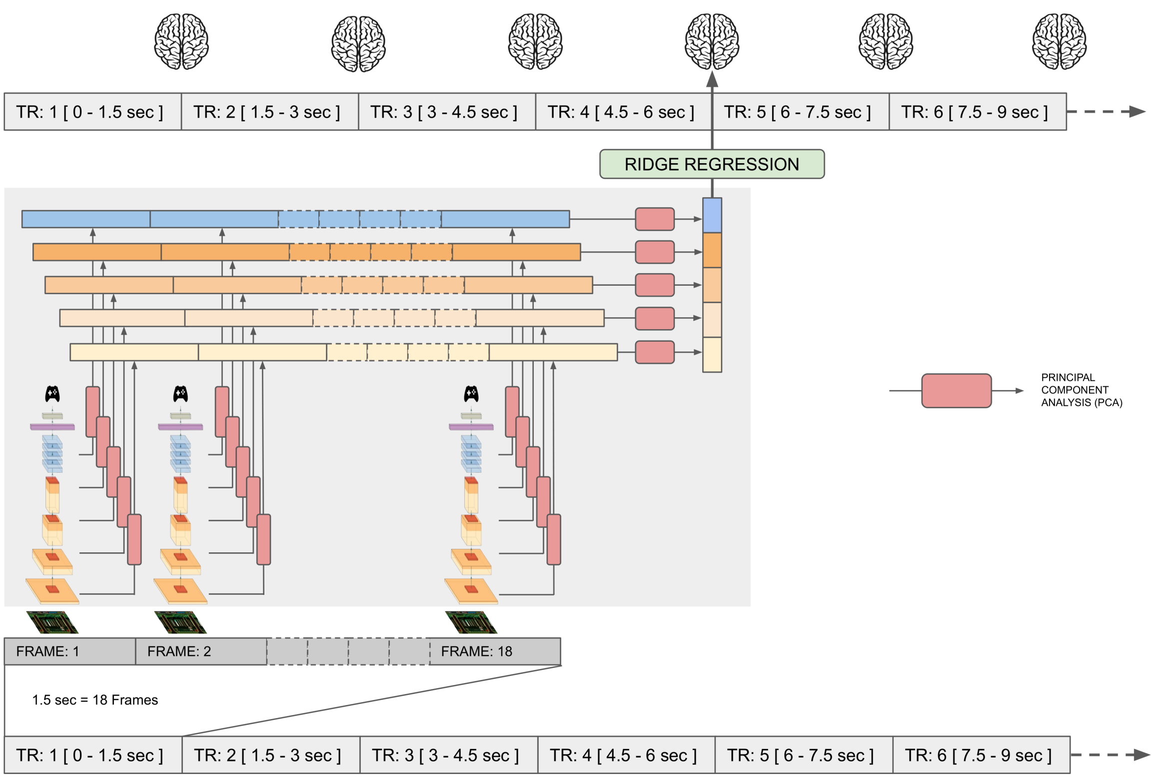

Fig. 0.3. Brain encoding framework. Subject's gameplay frames are passed into a trained imitation learning model and its internal network activations are reduced in dimension using PCA. These resulting features are used to predict fMRI BOLD responses using ridge regression.

Here we analyze the relationship between the feature space learned by the ANN and the subject's brain data [**78, 47**]. To compare the representation we use the brain encoding approach. Now as we want to see how close the features of a model trained to imitate behavior are to the brain data, we use only linear mapping to predict the brain data.

The mapping of activations to predict the BOLD responses is summarised in Figure 0.3. We reduce the dimensions of the activations as we have limited data points and hence cannot use very large feature vector dimensions to train the linear regression model. Activations from the four CNN layers, LSTM layer, and FC layer are reduced in dimension using a principal component analysis (PCA). To match the timescale between the video game frequency and fMRI frequency, we then make use of another PCA to compress information from 1.49-second windows (length of 1 TR), within each layer. Finally, the PCA vectors from each layer are concatenated to get the final feature space vector that is used for ridge regression. To account for the hemodynamic delay, the final feature space vector of the imitation model at time t is mapped to fMRI data at time t+X seconds, where the best encoding was found to be at X=6 (section 3.1.4). A separate ridge regression model is fit for each parcel.

2.4. Control models:

To test whether training an ANN to mimic the behavior of a given individual endows it with higher brain encoding performance in that individual, we compared brain encoding performance across a range of control models. The models we used for comparison are described below.

Untrained model: When the pixel input is passed through an untrained model, due to the nature of Convolutional and LSTM layers, useful features can be captured. To understand how much of variance is captured just due to the architecture of the model, we use an untrained model with randomly initialized weights as one of our control models.

Pixel space PCA model: To capture the amount of brain data variance explained just by the videogame frames, we do a PCA on the pixel RGB input and use it to train the ridge regression model.

Action space model As our final control model, we use the controller key presses of the subject to predict the brain data. This allows us to capture the amount of brain data variance explained just by the final action space.

3. Results

3.1. Individual imitation and brain encoding

3.1.1. Behavioral results

We first aimed to train ANNs that could replicate the gameplay of individual subjects. We trained a deep ANN composed of vision (CNN) and recurrent (LSTM) layers using behavioral cloning, to predict button presses directly from pixel-level video frames from the game. Leave-one-session-out cross-validation was performed and the obtained accuracy of predicted action vs actual subject action is plotted for each individual model. There are 128 possible actions to choose from, out of all possible combinations of 7 button presses. We saw that the model imitate the subject with a accuracy of 62-68% (Subject wise accuracies - Sub 01: 62%, Sub 02: 64%, Sub 04: 68%, Sub 06: 66%). As the classes are imbalanced due to the nature of the gameplay, we also evaluated accuracy scores re-weighted to give equal importance to all actions, which was found to be similar to the results shown above. To complement these quantitative performance metrics, we also used trained models to play the game for entire levels, that is input the actions predicted by the model back into the game emulator. Qualitatively we found that trained models performed adequately, being able to progress through levels and often complete them, and taking sequences of action analogous to individual subject gameplay. A demo video of imitation learning model playing the game with a subject-specific style is available at this Link.

3.1.2. Brain encoding distributions

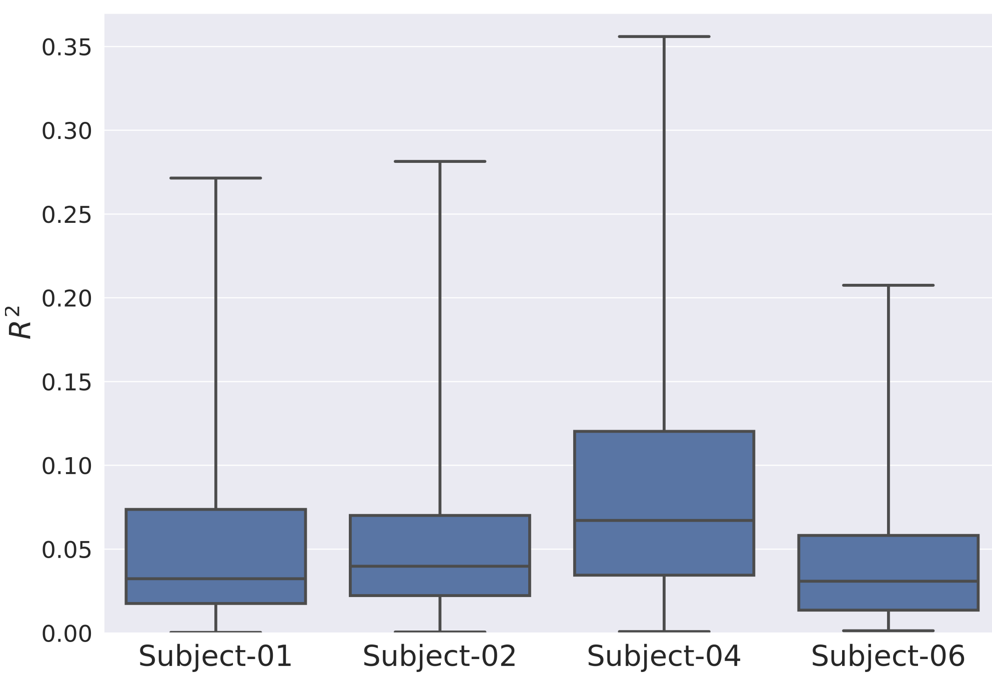

Fig. 0.4. Distribution of brain encoding R^2 (correlation between actual fMRI BOLD data and predicted BOLD responses) shown across four subjects. For each subject distribution of R^2 across 444 parcels is plotted using a box plot showing the minimum, first quartile, median, third quartile, and maximum.

Next, we tested whether the representations learned by individual imitation models could be used to predict task-related fluctuations in brain activity, for the same subjects. For this purpose, we implemented a brain encoding analysis, where the video frames of the game of a player were fed into the imitation model, the activity of layers of the model was used as input to a Ridge regression that predicted brain activity in a series of functional parcels. The coefficient of determination R^2 was calculated between the predicted fMRI data from the ridge regression model and the actual fMRI data for each parcel. Distribution of R^2 across the parcels is plotted for each subject and shown in Figure 0.4. The range of R^2 was consistent across subjects, with most values distributed from negligible (0 or lower) to moderate (0.2), and values up to 0.35 observed depending on the brain region.

3.1.3. Brain encoding maps

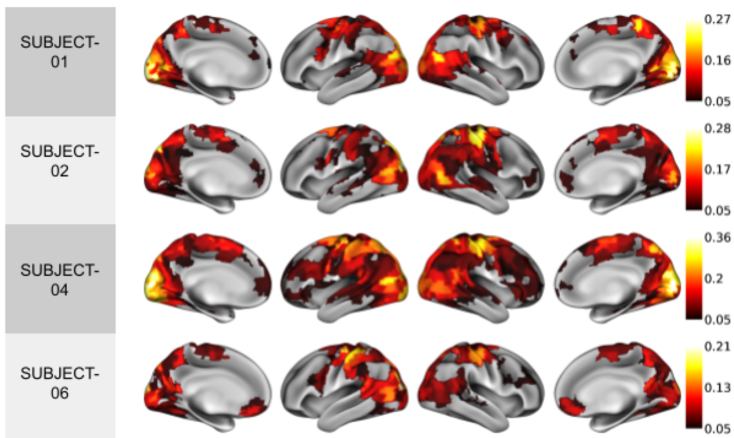

Fig. 0.5. Distribution of brain encoding R^2(correlation between actual and predict fMRI BOLD data), across four subjects plotted on a brain map. Network internal representations from separate Imitation learning models, each trained on that subject's gameplay was used to predict the fMRI data.

To understand in which brain regions the imitation learning model was able to predict the brain data, we projected the R^2 values on a brain map for each subject, shown in Figure 0.5. We see that the imitation learning model predicts brain dynamics with high accuracy in brain regions largely consistent across subjects. These included the visual cortex (in particular the dorsal stream), sensorimotor cortex, as well as some frontoparietal areas. Other parts of the brain, such as the prefrontal cortices, achieved low prediction accuracy (below 0.05).

3.1.4. Analysis of hemodynamic delay

The brain encoding model relies on a critical parameter, called the hemodynamic delay: the effects of the stimuli at a time T is expected to be reflected in the brain data at $T + X$ seconds. This delay X is known to be around 5 seconds [**41**], but variable across subjects as well as brain regions. Results above were shown for a delay of 4.5-6 seconds. We did test the impact of this parameter systematically, with delay X ranging from 0 to 12 seconds, which

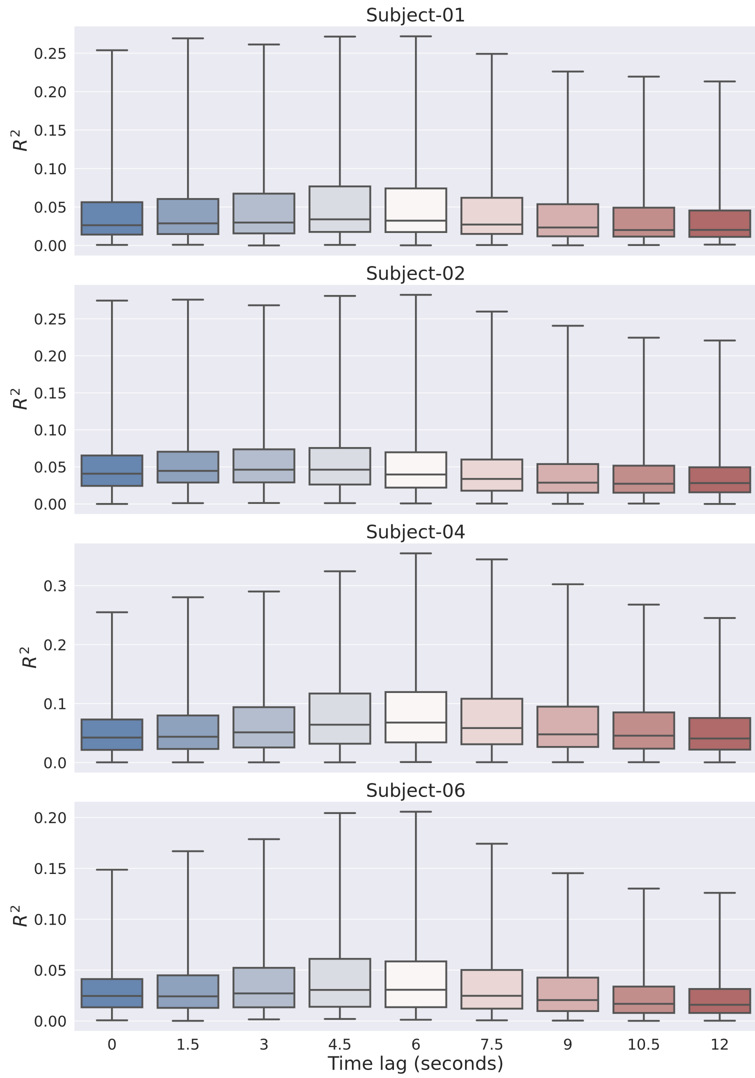

Fig. 0.6. Distribution of R^2 across the parcels plotted for each encoding model (differing in time-lag hyperparameter) and shown for all four subjects.

not only allows us to select the best X for this dataset but also serves as a sanity check to see if the effect of hemodynamic delay is visible.

The coefficient of determination R^2 was calculated between the predicted fMRI data from the ridge regression model and the actual fMRI data for each parcel. Distribution of R^2 across the parcels is plotted for each encoding model(differing in time-lag hyperparameter X) shown in Figure 0.6.

From Figure 0.6 we clearly see the effect of the hemodynamic delay where encoding gets better when the time lag used is increased from 0 seconds, peaks at 6 seconds, and then gradually worsen if we increase the lag further.

3.2. Individual brain encoding: Layer wise analysis

3.2.1. Brain encoding distributions

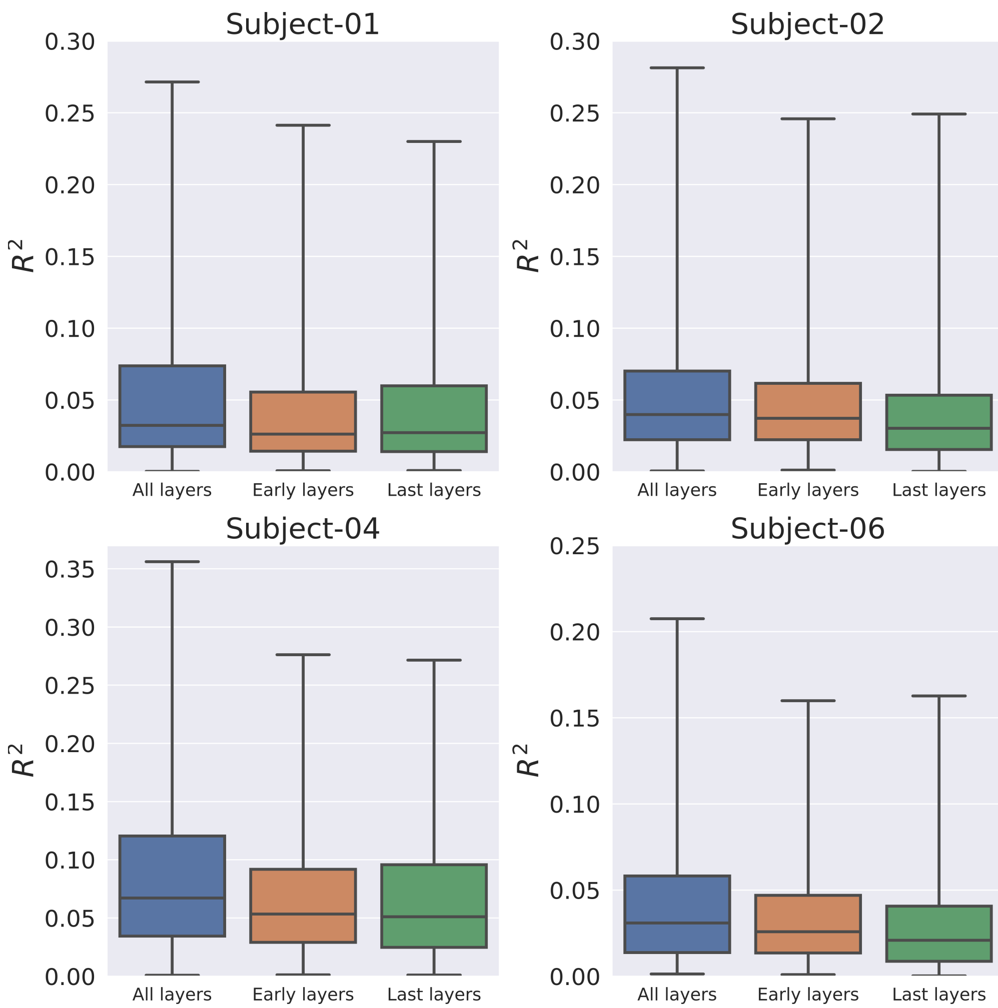

Fig. 0.7. Distribution of brain R^2 comparing encoding performed using trained model where (A) All the layers features were used, (B) Only early two convolutional layer features were used and (C) Last convolutional layer and LSTM layer features were used

In the previous section, we presented the encoding results from the trained imitation learning model where features from all the layers (Convolutional layer-1, Convolutional layer-2, Convolutional layer-3, Convolutional layer-4, and LSTM layer) was used to predict the brain data. In this section, we analyze what effect the early layers vs later layers of the model have on brain encoding and which brain regions each of them help to encode better. Apart from the encoding model that uses all the layer features we consider two additional models. Early layer model, where only the initial two layers (Convolutional layer-1, Convolutional layer-2) features were used, and last layers model where only the last two layers (Convolutional layer-4 and LSTM layer) features were used to predict the brain data using ridge regression. Distribution of R^2 across the parcels for the three models are plotted in Figure 0.7. We see that across all four subjects, encoding performed using all the layers features was better at predicting the brain data rather than using just the early layer features or just the latter layer features.

3.2.2. Brain encoding maps

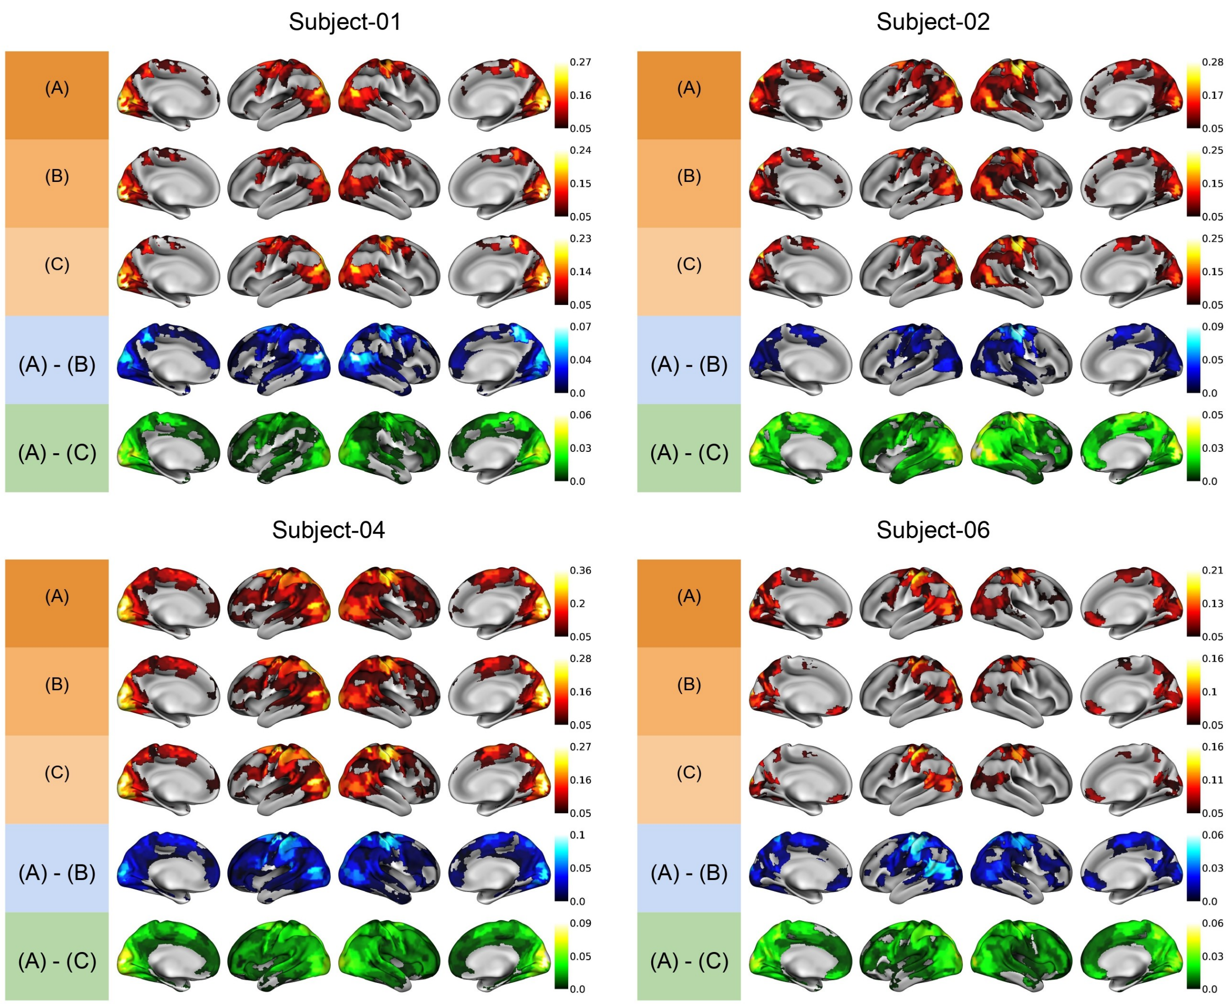

Fig. 0.8. R^2 values on a brain map shown for each of the four subjects. Comparison of encoding based on which layers from the trained model were used. **A**: All the layers features were used. **B**: Only early two convolutional layer features were used. **C**: Last convolutional layer and LSTM layer features were used. **A-B**: R^2(All layers) - R^2(Early layers). This shows the effect of the latter layers in the model. **A-C**: R^2(All layers) - R^2(Last layers). This shows the effect of early layers in the model.

In the above section, we saw that using all the layers features helped encode the brain data better. In order to see which regions of the brain were predicted better by which layers of the model, we projected the R^2 values in the brain map for each subject as shown in Figure 0.8. We also contrasted R^2 values derived from all layer model and partial layer models, and only show values for parcels with a significant difference (q<0.05). We found that the early layer predicted the early visual cortex activation better (shown in Figure 0.8 by rows (B) and (A-C)), and the latter layers of the model better predicted the higher visual cortex, motor, and somatosensory regions (shown in the Figure 0.8 by rows (C) and (A-B)).

3.3. Subject specificity

3.3.1. Behavioral results

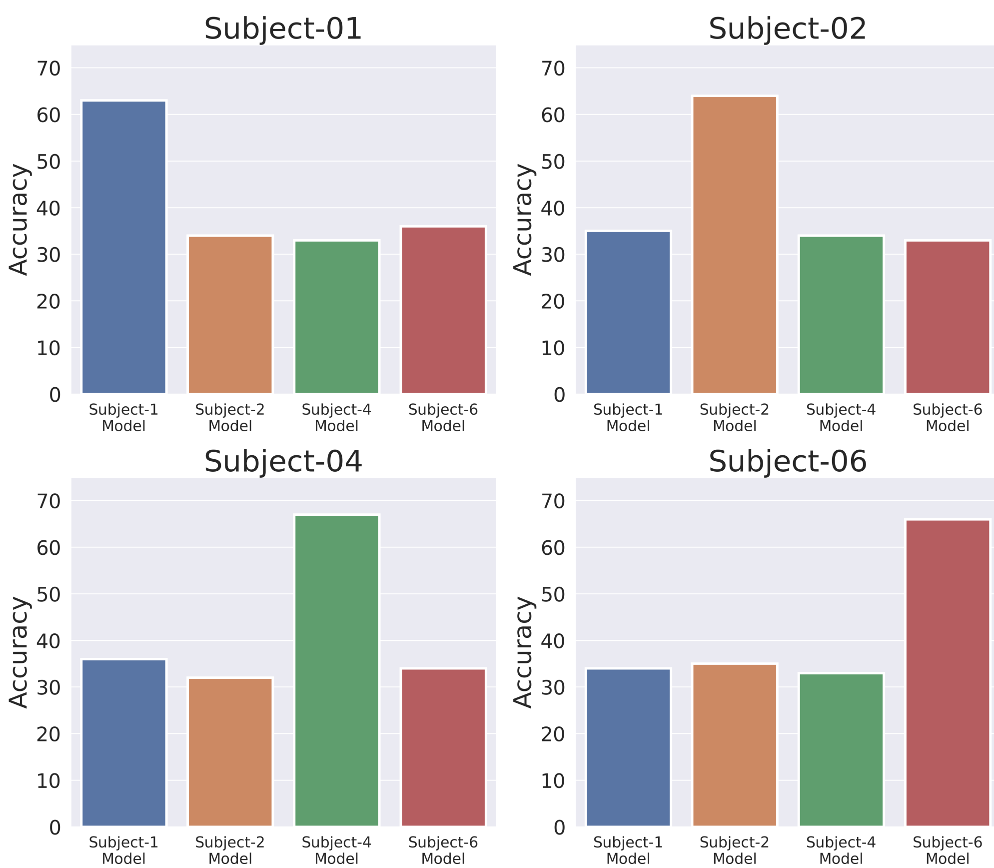

Fig. 0.9. Behavioral imitation learning accuracies across four subjects. X-axis showing which subject the model was trained on.

In section 3.1.1, we saw that the imitation learning model is able to predict actions up to 65% accuracy. Next, we tested whether these imitation models are subject-specific. Specifically, we evaluated if the imitation learning model trained on a subject's gameplay can model the behavior of that same subject better than a model that knows to play the game but in a different gameplay style, drawn from the other three subjects. The results for each of the four subjects are shown in Figure 0.9. While the imitation learning model trained on the same-subject gameplay was able to replicate the gameplay up to 65% accuracy, the models trained on a different gameplay style only achieved 35% accuracy. This result was observed consistently across all four subjects. Overall, we found behavioral imitation models to be highly subject-specific.

3.3.2. Brain encoding distributions

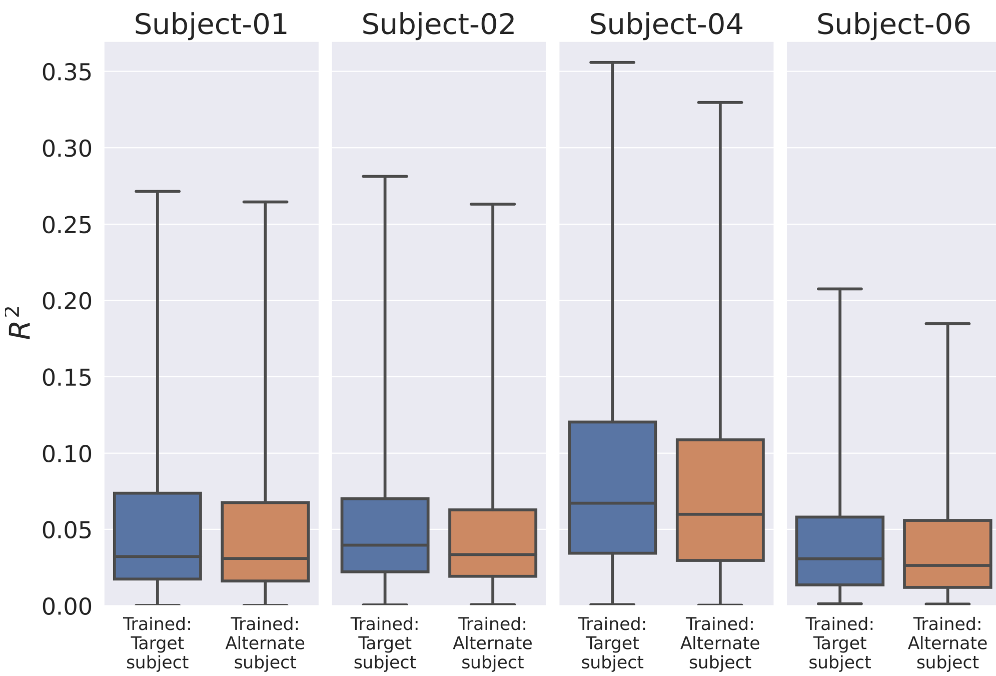

Fig. 0.10. Distribution of brain R^2 comparing encoding performed using model trained on target(same) subject vs alternate(best out of rest three) subject.

Subject	Percentage of parcels significant
01	5.6%
02	37.3%
04	51.5%
06	23.1%

Table 0.4. Significance test results. Percentage of parcels significant(q < 0.05): Table shows in how many parcels the encoding using the Imitation model trained on the same subject gameplay is better than a model trained on different subject gameplay.

Next, we wanted to see if the subject-specificity of imitation models translated into better brain encoding. Specifically, we tested if a subject-specific imitation model is able to encode the brain data of that target individual better than imitation learning models trained on different, alternative subject gameplay. This alternative was selected as the subject most behaviorally similar to the target. We passed prerecorded gameplay frames of the target subject into two models: (1) the imitation model trained on the target subject and (2) the imitation model trained on the alternative subject. The internal activations from these models were used to predict brain data using ridge regression.

Distribution of R^2 across the parcels for the target and alternative models are plotted in Figure 0.10 for each subject. We see that across all four subjects, the target (same-subject) imitation model is better at encoding the brain data than the alternative (different subject) model. These differences were significant at q<0.05 (Wilcoxon signed-rank test with false discovery rate correction for multiple comparisons across parcels). Same-subject imitation learning models better encode brain data in 12%, 35%, 37%, 26% of parcels in Subjects 01, 02, 04 and 06 respectively, see Table 0.4. Overall, we found that imitation models which are specific of individual behavior are better to predict brain activity of this individual.

3.3.3. Brain encoding maps

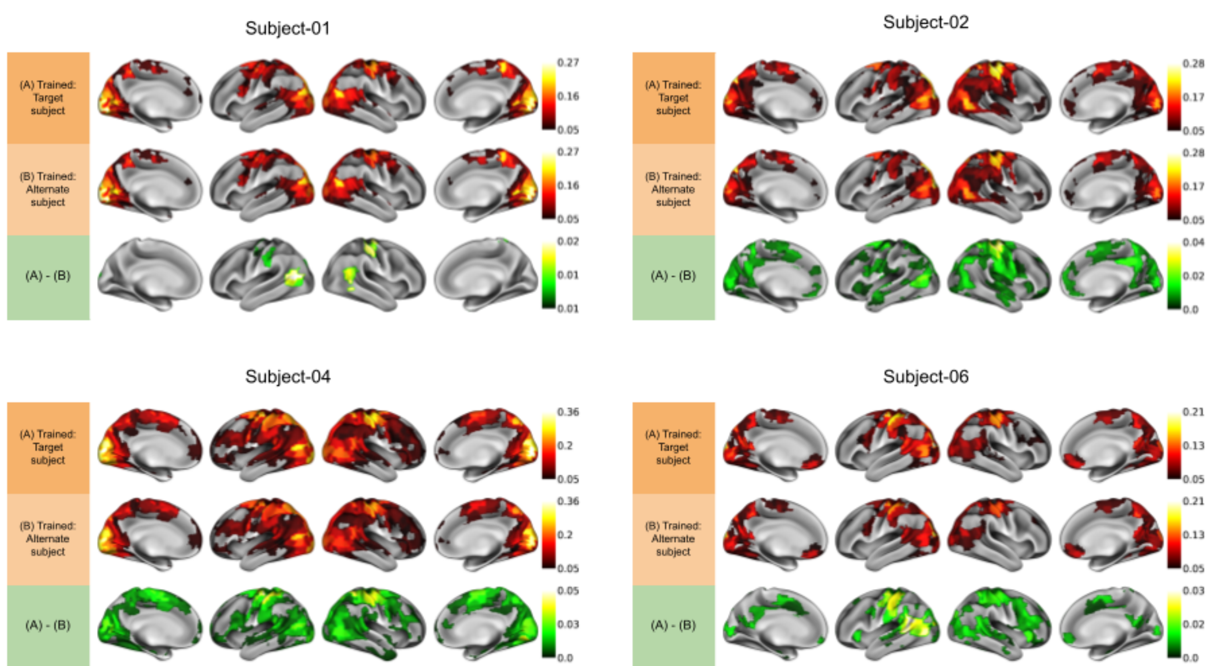

Fig. 0.11. R^2 values on a brain map shown for each of the four subjects. **First row**: Brain map of R^2 from the Trained model (Same subject gameplay). **Second row**: Brain map of R^2 from the Trained model (Different gameplay). **Third row**: Brain map of R^2 (Trained model: Same subject gameplay) - R^2 (Trained model: different gameplay) , where the R^2 values are subtracted parcel wise and plotted.

Next, we wanted to understand which brain regions were better encoded using subject-specific imitation models. We projected the R^2 values on a brain map for each subject and each type of imitation model, see Figure 0.11. We also contrasted the R^2 values derived from two imitation models: target subject minus alternate subject, and only show values for parcels with a significant difference (q<0.05). We found that the imitation model trained on the target subject was able to better predict the brain data, particularly in visual, somatosensory and motor cortex regions. These regions were consistent across all four subjects.

3.4. Control experiments

3.4.1. Brain encoding distributions

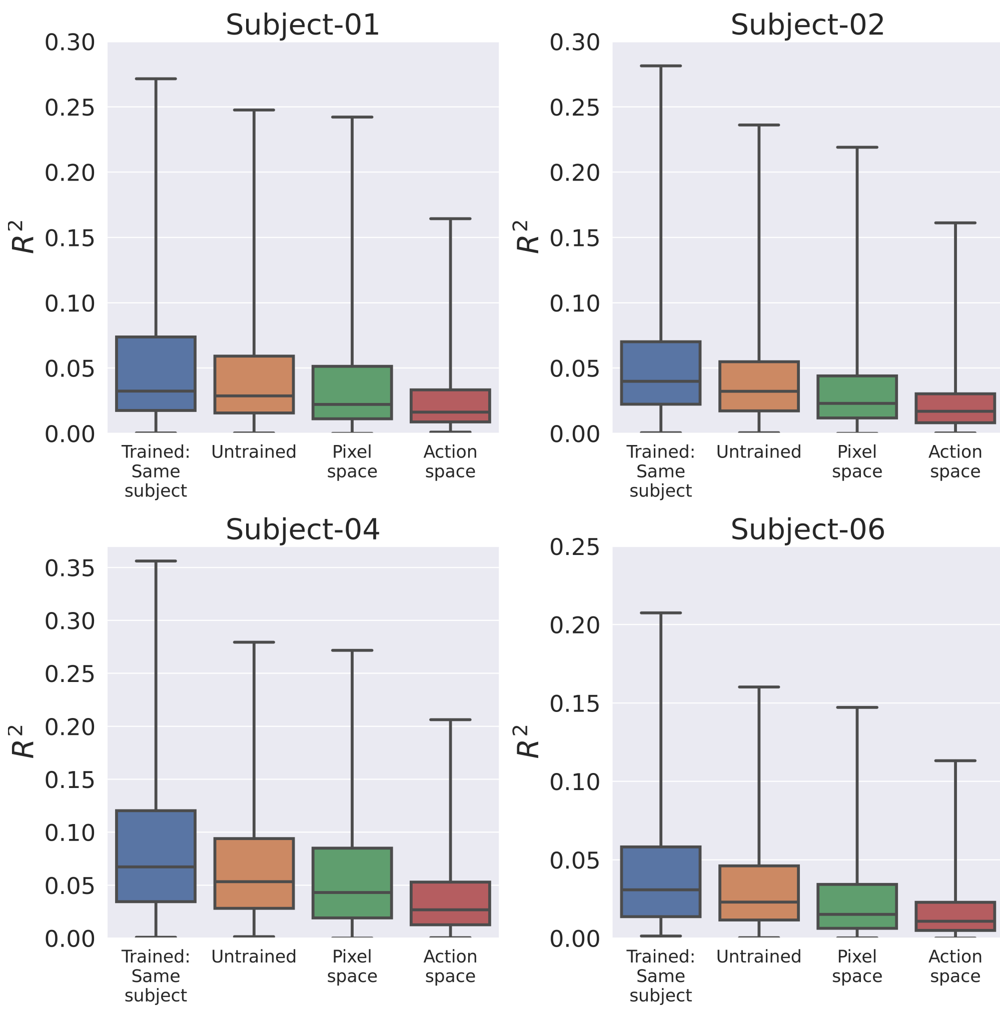

Fig. 0.12. Distribution of brain R^2 comparing encoding performed using model trained on target(same) subject vs control models.

We finally implemented some control experiments. We first wanted to test if the imitation learning models were not just encoding brain data based on basic features like visual input or final gameplay actions. Second, we wanted to see how much of the brain encoding results

Subject	Untrained model	Pixel-space PCA model	Action-space model
01	41.6%	88.1%	85.5%
02	60.8%	92.5%	90.7%
04	75.6%	97.1%	93.0%
06	56.9%	89.8%	87.8%

Table 0.5. Significance test results. Percentage of parcels where the imitation model trained on the same subject gameplay was significantly better (q<0.05) than the control model.

was just due to the properties of CNN-LSTM architecture itself, rather than training with behavioral cloning.

For these aims, we compared brain encoding with subject-specific imitation models to a series of control models: (1) PCA features derived from gameplay input pixel (Pixel-space model), (2) one-hot encoding of the controller button presses (Action-space model) and, (3) passing the subject's gameplay frames to an CNN-LSTM model with random weights (Untrained model).

Distribution of R^2 across the parcels for imitation and control models are shown in Figure 0.12. Across all four subjects, the subject-specific imitation model was better at encoding brain data than all three controls. Significant improvement in brain encoding using imitation models was found in about 40% of parcels (average across subjects) when compared to Untrained model, and in about 85% of parcels (average across subjects) when compared with the Pixel-space and Action-space control models, see Table 0.5. No single parcel was significantly better encoding by the control models than by the imitation model.

3.4.2. Brain encoding maps

We examined in which brain regions imitation models encoded brain data better than controls, using R^2 values projected as brain maps, shown in Figure 0.13. We also derived contrast maps by subtracting the R^2 values imitation model minus controls. When compared to Untrained, Pixel-space and Action-space models, the imitation learning model was able to predict brain data significantly better (q<0.05) in most regions of the brain, with highest differences found in the same areas as the contrast between target-subject minus alternate-subject models: visual, sensorimotor and dorsal attentinal.

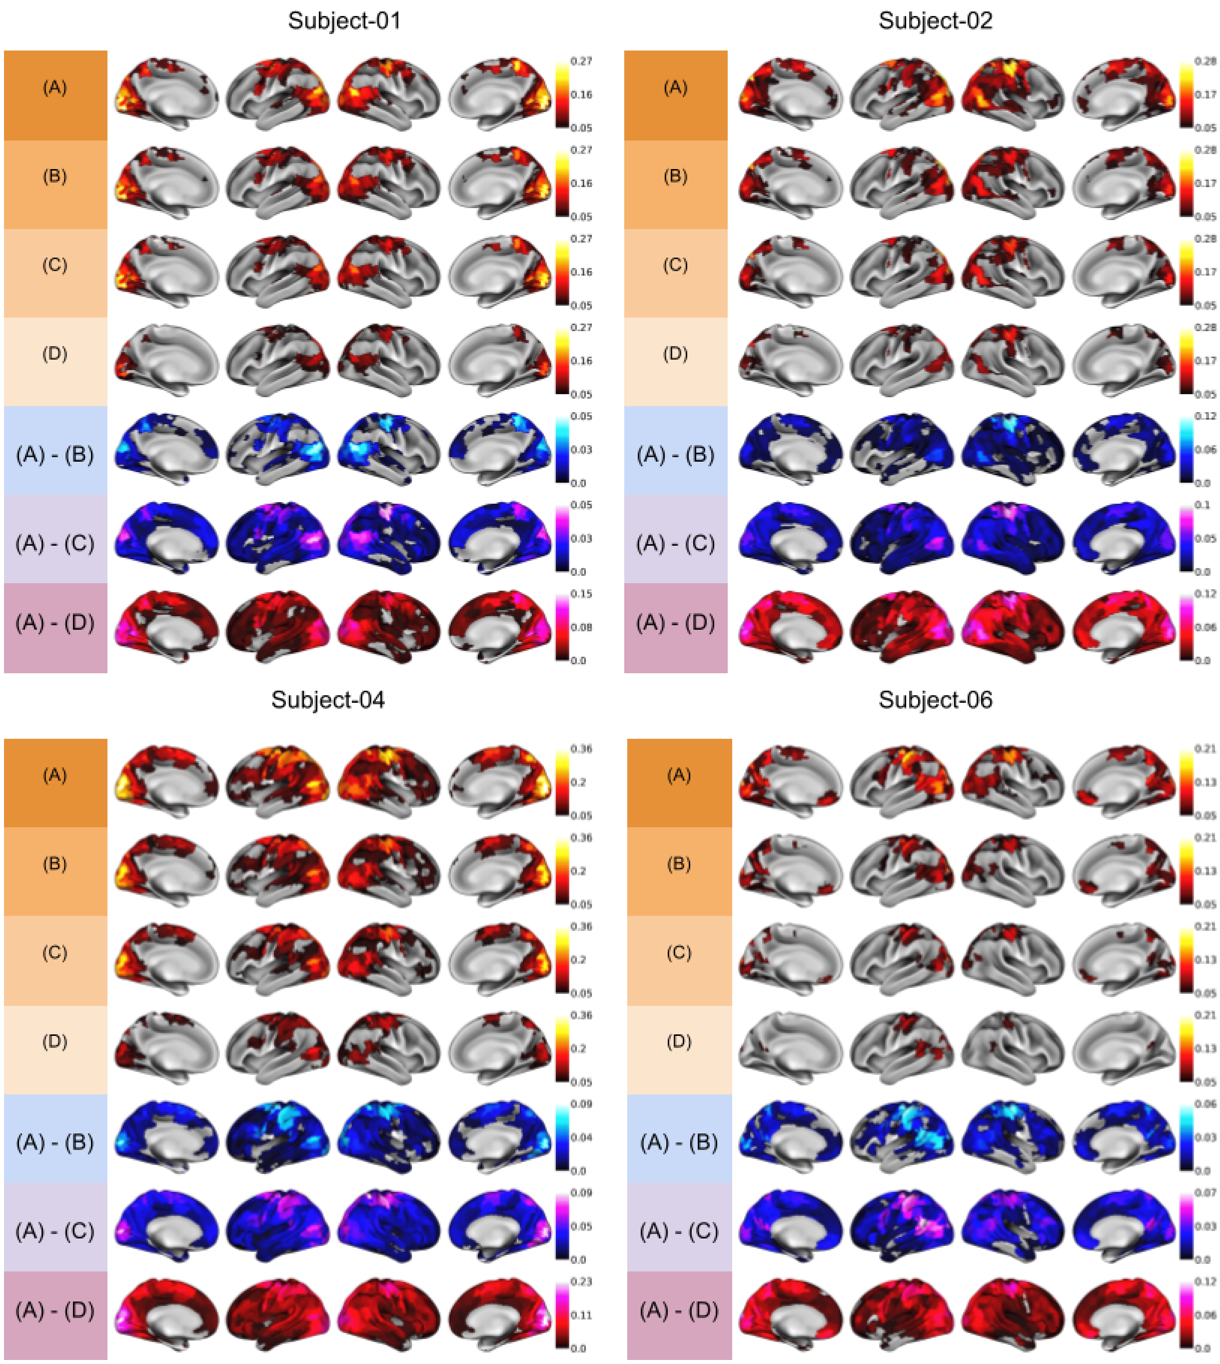

Fig. 0.13. R^2 values on a brain map shown for each of the four subjects. **A**: Trained model (Same subject gameplay). **B**: Untrained model. **C**: Pixel-space model. **D**: Action-space model. **A-B**: R^2(Trained model) - R^2(Untrained model). **A-C**: R^2(Trained model) - R^2(Pixel space model). **A-D**: R^2(Trained model) - R^2(Action space model).

4. Discussion and Conclusion

In this work, we showed that a deep recurrent neural network is able to successfully imitate the actions of a human player in a complex videogame and that the representations learned by this network can be used to perform brain encoding. Those individual imitation models were subject-specific, both in terms of behavior (button presses) and brain encoding. Subject-specific imitation models also performed significantly better than a series of control models. Specifically, we found that the improvement in encoding was particularly marked in motor, somatosensory, and visual cortices.

Our individual imitation models trained with Behaviour cloning (BC) show good quantitative performance (up to 70% accuracy) as well as qualitative face validity when used to play the game autonomously. It is however difficult to conclude whether these models fully imitate individual gameplay. Humans have variable behavior, and it is very likely that a player would not be able to predict their own actions 100% of the time if presented with recorded videos of their gameplay. Study by Chen et. al. [**11**] showed when trained on various Atari games the BC approach achieves accuracies between 60% to 90% depending on the game and the size of training data. Work by Kanervisto et. al.[**33**] showed that the accuracies of the BC models greatly depend on the quality of data. Hence it might not be easy to compare our accuracies to previous works which might have a difference in dataset size, quality, and nature of videogames.

We found that an ANN model trained to imitate the subject serves as a good encoding model of the brain dynamics of that subject with R^2 values up to 35%. The encoding worked particularly well in motor, somatosensory, and visual cortices. Although these brain regions were consistent across subjects, the average and maximum R^2 varied markedly across subjects. These results likely reflect systematic differences in data quality (and noise ceiling) across subjects. Noise in fMRI is complex, but a major source is motion [**59**], and different subjects had markedly different amounts of motion, consistently across scans. We also found that encoding R^2 were poor in subcortical regions across all subjects, and the sequence used in our experiment had low SNR in those regions as is common with multiband accelerated fMRI. In raw image presentation, a noise ceiling can be estimated through repeated stimuli presentation to define an optimal brain response as done in this study by Allen et. al.(2022) [**2**], which was not possible to implement with videogame whose gameplay varied across repetitions and subjects. The only previous study of brain encoding with ANNs in videogames, Cross et. al.[**14**], also achieved maximum encoding R^2 up to 0.38 (average across subjects) and also had best encoding in visual, motor and somatosensory areas. It is not possible to directly compare R^2 between the two studies however, as noise ceiling (and R^2) are directly impacted by parameters such as image resolution, scanning sequence and the amount of spatial smoothing, none of which were harmonized across the two studies.

To understand which regions of the brain were predicted better by which layers of the model, we compared brain encoding predictions from features derived from only early layers vs the last layers of the model. We found that early layers predicted the brain data better in the early visual cortex and the latter layers (which included the LSTM layer) of the model better predicted the higher visual cortex, motor, and somatosensory regions. This can be attributed to the fact that early convolutional layers capture the primitive visual features of the image similar to what happens in the early visual cortex areas, and the latter layers of the model capture the gameplay behavior, which is also exhibited in motor and somatosensory regions of the brain. These layer-wise effects on brain encoding are consistent with results observed in various brain encoding studies like Yamins et. al.[**81**] and Cross et. al.[**14**].

Further, we found that imitation models were subject specific. While imitation models trained on individual gameplay were able to predict actions with an accuracy up to 70% for this individual, imitation models trained on the gameplay of another subject were only able to predict actions accurately up to 30-40% accuracy. This improvement in behavioral accuracy also translated to better brain encoding. Specifically, we found that individual-specific imitation models improved encoding in motor, somatosensory and visual cortices, compared to models trained with different gameplay. These areas are key to processing the input of the task (visual stimuli) and generating outputs (motor command). While imitation models trained on different gameplay know how to play the game, the actions they might are often dissimilar to that of the target subject, as demonstrated by the loss in accuracy for predicting actions. Representations of these subject-specific actions likely explain the ability of subject-specific imitation models to captures more brain variance in motor and somatosensory areas. By contrast, brain encoding in frontoparietal cortices did not show widespread improvement. Frontoparietal cortices are involved in planning, and we expected to show improvement due to subject-specific behavioral imitation models. We still note some improvements in the dorsal attentional network, which is a subset of frontoparietal regions tightly linked with motor control and attention (as the name implies). In the experiment setup, we had asked the subjects to practice the game for about six months before they played the game inside fMRI scanner. They had reached highly consistent gameplay which might have led to high levels of automation, and little effort for planning. This may explain the limited involvement of frontoparietal networks.

Control analysis showed that neither pixel input or the actions of the player were sufficient to capture brain dynamics in a complex task like playing Shinobi. We also see that training an ANN with an imitation model is beneficial to brain encoding, and features derived from an untrained CNN-LSTM network (to which raw input was fed) did not explain any more variance than models trained on raw input game pixel data. These results were observed across all cortical regions. Some previous studies showed that the structure of a CNN network is sufficient to capture a great amount of low-level image statistics prior to any learning [**72**] and such untrained networks can be used to encode the brain in visual tasks [**36**]. Our results show that in complex naturalistic tasks such as videogames, just the stimuli/final behavior/network architecture might not be able to capture the brain dynamics, instead, a deep learning model that learns to perform the task in a way similar to the subject would be able to encode the brain dynamics of that subject better. Cross et. al.[**14**] study on videogame brain encoding, found similar results where raw statistical properties of input/output were not able to explain the variance in brain data well while playing videogames, even if their reference models was not tuned to specific individual gameplay.

One of the limitations of this study is that given the slow temporal resolution of fMRI and the quick nature of video gameplay, it is possible that the advantages of a subject-specific imitation learning model might not have been completely translated to better brain encoding, due to a loss of fast action dynamics. Future work using high temporal resolution imaging techniques, such as EEG or MEG, would thus complement this fMRI study.

To conclude, this work demonstrates the feasibility of studying brain dynamics in naturalistic conditions, using rich dynamic environments. Work done in this study showed that in complex tasks such as playing video games, it is possible to capture individual-specific behavior using imitation learning and this serves as a better model to encode brain data which highlights the potential of combining imitation learning, brain imaging, and videogames to uncover subject-specific relationships between brain and behavior.

Chapter 3

Discussion

In this thesis, we set out to investigate the feasibility of using ANNs to model brain and behavioral data simultaneously, while subjects do naturalistic tasks such as playing a complex videogame. We hypothesized that in such tasks a model that is trained to capture subject-specific behavior should be able to better encode the brain data of the corresponding individual. We recorded brain and behavioral data of subjects playing the videogame “Shinobi III Return of the Ninja Master” inside an fMRI scanner. Using imitation learning on behavioral data, we were able to train an ANN to exhibit a gameplay style similar to each individual subject. Further, we used the ANN representations of this imitation learning model to predict fMRI brain activity using a brain encoding analysis. We saw that these subject-specific imitation models performed significantly better at brain encoding than models trained on the gameplay style of different subjects. Specifically, we found that the improvement in encoding was particularly marked in motor, somatosensory, and visual cortices.

Next, I will briefly discuss some limitations of this work, potential impacts, and outline future research directions that I think should be pursued based on the observations reported here.

In this study, we have only four subjects ($N = 4$), while this could be a limitation if we were to study generalizable group properties, in this study we are building and studying individual-specific models. We found all the results to be consistent across the four subjects. There is a recent trend towards using "deep" neuroimaging data, where there is a large sample size in terms of acquisition length per subject, but a limited number of subjects [**48**]. The trend towards deep fMRI datasets reflects this push to better understand and characterize individual brains, complementary to studies featuring thousands of subjects. Such individual design is common in animal neuroscience research, but more novel in human studies. Group studies typically capture average effects on a group which may not apply to any particular individual in that group. This type of study does not concern itself with generalization at the level of a population but rather captures variations within a subject.

One main limitation of our work, as discussed in the previous chapter, was that given the slow temporal resolution of fMRI and the quick nature of video gameplay, the advantages of a subject-specific imitation learning model might not have been completely translated to better brain encoding. In practice, button presses can be separated by as little as 50 ms, and commonly 150 ms or less. The fMRI signal captures changes in the vasculature of the brain, and those changes only become apparent after seconds following events. The response also aggregates over events, in effect acting as a low pass filter with low cut frequency (typically below 0.08 Hz). It is therefore clear that our brain encoding model cannot capture any fine temporal dynamics due to the intrinsic limitations of the imaging method for this project. On another hand, fMRI has good spatial resolution and covers the whole brain, allowing us to draw detailed maps of brain encoding, but the issue of fast temporal dynamics remains.

EEG and MEG, on another hand, have a millisecond temporal resolution. Early in this Master's project, our plan was to collect both MEG and fMRI data for each participant. Several issues have delayed the collection of the MEG data. Because of COVID-19, the Helium shortage due to Russia-Ukraine war, and other technical difficulties, the MEG platform at UdeM experienced several closures and we were not able to collect the MEG data during the thesis.

To overcome these logistic delays and have high temporal resolution data for the same task, we started collecting EEG data using Wearable Sensing's DSI-24 mobile EEG. Though I was not able to analyze this data, I had hands-on experience with setting up the system and helped acquire brain data on subjects playing a videogame. This was a useful learning experience during my thesis, and this dataset may be useful to apply the same approach as described in this thesis, but this time using EEG data.

After the collection of EEG data, we would have a very unique dataset in which the same subjects play the same videogame in two different modalities: fMRI and EEG. While there have been some previous datasets and works with a multi-model dataset for simple visual cognitive tasks, this would be the first time such rich data would be collected for interactive naturalistic tasks. It would be very interesting to see what additional insights we can uncover from these joint EEG-fMRI networks.

As discussed in the previous chapter we were unable to capture brain dynamics in regions typically involved in planning and learning. We hypothesize that its due to the subjects reaching a high degree of consistency in gameplay due to six months of practice before being scanned. Hence in our next experiment, we have begun scanning subjects as they discover and learn to play the Mario videogame. As suggested in the study by Saxe et. al.(2021)[**63**], one of the main gaps in the current understanding of the brain is its learning principles and the authors suggest deep learning encoding models might help answer it. However, for this, we need to compare and focus on how representations change in ANN and the brain during task training. Comparisons of learning trajectories can distinguish computational principles of learning even without specifying biological implementations. Our next experiment where we are collecting brain and behavior data as subjects learn to play might help us answer this important question of learning in brains.

One more limitation of the setup used in our study might be that it is a side-scroller videogame, which implies that visual inputs are quite similar even with different gameplay, thus minimizing the brain response variability. Hence first-person 3d navigation games like Doom, where depending on the gameplay the game scenes change a lot, might help us uncover subject-specific relationships between brain and behavior more accurately.

The CNeuroMod project's objective is to provide a distinctive neuroimaging resource that contains vast amounts of data at an individual level and covers a wide range of cognitive domains. It is to my knowledge the most extensive individual fMRI dataset currently available, encompassing diverse cognitive domains such as movie watching, video game playing, and functional tasks developed for the Human Connectome Project. And the most exciting feature is this data is that the same subjects are being scanned for tasks in various cognitive domains. We are currently building individual encoding models for all these tasks, and next, we would like to combine these individual models across these multiple domains. Studying the interaction between multiple latent spaces (trained for distinct tasks, on distinct sensory inputs and/or modalities) might help us build a more holistic encoding model of the individual [**75**]. Building such individual artificial models combining brain signals and behavior can not only help us uncover the workings of the brain but also can prove to be very useful in the domain of personalized medicine [**71, 25**].

To conclude, this work demonstrates the feasibility of studying brain dynamics using artificial neural networks in naturalistic conditions, using rich dynamic environments. Work done in this study showed that in complex tasks such as playing video games, it is possible to capture individual-specific behavior using imitation learning and this serves as a better model to encode brain data which highlights the potential of combining imitation learning, brain imaging, and videogames to uncover subject-specific relationships between brain and behavior.

References

[1] Stephen Adams, Tyler Cody, and Peter A Beling. A survey of inverse reinforcement learning. *Artificial Intelligence Review*, 55(6):4307–4346, 2022.

[2] Emily J Allen, Ghislain St-Yves, Yihan Wu, Jesse L Breedlove, Jacob S Prince, Logan T Dowdle, Matthias Nau, Brad Caron, Franco Pestilli, Ian Charest, et al. A massive 7t fmri dataset to bridge cognitive neuroscience and artificial intelligence. *Nature neuroscience*, 25(1):116–126, 2022.

[3] John R Anderson, Daniel Bothell, Jon M Fincham, Abraham R Anderson, Ben Poole, and Yulin Qin. Brain regions engaged by part-and whole-task performance in a video game: a model-based test of the decomposition hypothesis. *Journal of cognitive neuroscience*, 23(12):3983–3997, 2011.

[4] B.B. Avants, C.L. Epstein, M. Grossman, and J.C. Gee. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis*, 12(1):26–41, 2008.

[5] Daphne Bavelier, Rebbeca L Achtman, Merry Mani, and Julia Föcker. Neural bases of selective attention in action video game players. *Vision research*, 61:132–143, 2012.

[6] Yashar Behzadi, Khaled Restom, Joy Liau, and Thomas T. Liu. A component based noise correction method (CompCor) for BOLD and perfusion based fmri. *NeuroImage*, 37(1):90–101, 2007.

[7] Pierre Bellec and Julie A Boyle. Bridging the gap between perception and action: the case for neuroimaging, ai and video games.

[8] Julie Boyle, B Pinsard, A Boukhdhir, S Belleville, S Brambatti, J Chen, J Cohen-Adad, A Cyr, Fuente, P Rainville, and P Bellec. The courtois project on neuronal modelling – first data release. 26th annual meeting of the organization for human brain mapping, 2020. 2020.

[9] Charlotte Caucheteux and Jean-Rémi King. Language processing in brains and deep neural networks: computational convergence and its limits. *BioRxiv*, pages 2020–07, 2020.

[10] Charlotte Caucheteux and Jean-Rémi King. Brains and algorithms partially converge in natural language processing. *Communications biology*, 5(1):1–10, 2022.

[11] Brian Chen, Siddhant Tandon, David Gorsich, Alex Gorodetsky, and Shravan Veerapaneni. Behavioral cloning in atari games using a combined variational autoencoder and predictor model. In *2021 IEEE Congress on Evolutionary Computation (CEC)*, pages 2077–2084. IEEE, 2021.

[12] JCH Christopher. Watkins and peter dayan. *Q-Learning. Machine Learning*, 8(3):279–292, 1992.

[13] Radoslaw Martin Cichy, Aditya Khosla, Dimitrios Pantazis, Antonio Torralba, and Aude Oliva. Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific reports*, 6(1):27755, 2016.

[14] Logan Cross, Jeff Cockburn, Yisong Yue, and John P. O'Doherty. Using deep reinforcement learning to reveal how the brain encodes abstract state-space representations in high-dimensional environments. *Neuron*, 2020.

[15] Logan Cross, Jeff Cockburn, Yisong Yue, and John P O'Doherty. Using deep reinforcement learning to reveal how the brain encodes abstract state-space representations in high-dimensional environments. *Neuron*, 109(4):724–738, 2021.

[16] Thomas Deffieux, Charlie Demene, Mathieu Pernot, and Mickael Tanter. Functional ultrasound neuroimaging: a review of the preclinical and clinical state of the art. *Current opinion in neurobiology*, 50:128–135, 2018.

[17] James J DiCarlo, Davide Zoccolan, and Nicole C Rust. How does the brain solve visual object recognition? *Neuron*, 73(3):415–434, 2012.

[18] Jörn Diedrichsen and Nikolaus Kriegeskorte. Representational models: A common framework for understanding encoding, pattern-component, and representational-similarity analysis. *PLoS computational biology*, 13(4):e1005508, 2017.

[19] Katharina Dobs, Julio Martinez, Alexander JE Kell, and Nancy Kanwisher. Brain-like functional specialization emerges spontaneously in deep neural networks. *Science advances*, 8(11):eabl8913, 2022.

[20] Michael Eickenberg, Alexandre Gramfort, Gaël Varoquaux, and Bertrand Thirion. Seeing it all: Convolutional network layers map the function of the human visual system. *NeuroImage*, 152:184–194, 2017.

[21] Oscar Esteban, Ross Blair, Christopher J. Markiewicz, Shoshana L. Berleant, Craig Moodie, Feilong Ma, Ayse Ilkay Isik, Asier Erramuzpe, Mathias Kent, James D. andGoncalves, Elizabeth DuPre, Kevin R. Sitek, Daniel E. P. Gomez, Daniel J. Lurie, Zhifang Ye, Russell A. Poldrack, and Krzysztof J. Gorgolewski. fmriprep. *Software*, 2018.

[22] Oscar Esteban, Christopher Markiewicz, Ross W Blair, Craig Moodie, Ayse Ilkay Isik, Asier Erramuzpe Aliaga, James Kent, Mathias Goncalves, Elizabeth DuPre, Madeleine Snyder, Hiroyuki Oya, Satrajit Ghosh, Jessey Wright, Joke Durnez, Russell Poldrack, and Krzysztof Jacek Gorgolewski. fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nature Methods*, 2018.

[23] Oscar Esteban, Christopher J Markiewicz, Ross W Blair, Craig A Moodie, A Ilkay Isik, Asier Erramuzpe, James D Kent, Mathias Goncalves, Elizabeth DuPre, Madeleine Snyder, et al. fmriprep: a robust preprocessing pipeline for functional mri. *Nature methods*, 16(1):111–116, 2019.

[24] VS Fonov, AC Evans, RC McKinstry, CR Almli, and DL Collins. Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage*, 47, Supplement 1:S102, 2009.

[25] Tami Freeman. https://physicsworld.com/a/virtual-brain-helps-improve-the-outcome-of-epilepsy-surgery/.

[26] Olivia Guest and Andrea E. Martin. How computational modeling can force theory building in psychological science. *Perspectives on Psychological Science*, 16(4):789–802, 2021. PMID: 33482070.

[27] Matthew J. Hausknecht and Peter Stone. Deep recurrent q-learning for partially observable mdps. *CoRR*, abs/1507.06527, 2015.

[28] Elnaz Jahani Heravi, Hamed H Aghdam, and Domenec Puig. Classification of foods using spatial pyramid convolutional neural network. In *CCIA*, pages 163–168, 2016.

[29] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, nov 1997.

[30] Ha Hong, Daniel LK Yamins, Najib J Majaj, and James J DiCarlo. Explicit information for category-orthogonal object properties increases along the ventral stream. *Nature neuroscience*, 19(4):613, 2016.

[31] Shailee Jain and Alexander G. Huth. Incorporating context into language encoding models for fmri. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, page 6629–6638, Red Hook, NY, USA, 2018. Curran Associates Inc.

[32] Mark Jenkinson, Peter Bannister, Michael Brady, and Stephen Smith. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):825–841, 2002.

[33] Anssi Kanervisto, Joonas Pussinen, and Ville Hautamäki. Benchmarking end-to-end behavioural cloning on video games. In *2020 IEEE conference on games (CoG)*, pages 558–565. IEEE, 2020.

[34] Kendrick N Kay, Thomas Naselaris, Ryan J Prenger, and Jack L Gallant. Identifying natural images from human brain activity. *Nature*, 452(7185):352–355, 2008.

[35] Alexander J.E. Kell, Daniel L.K. Yamins, Erica N. Shook, Sam V. Norman-Haignere, and Josh H. McDermott. A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron*, 98(3):630–644.e16, May 2018.

[36] Gwangsu Kim, Jaeson Jang, Seungdae Baek, Min Song, and Se-Bum Paik. Visual number sense in untrained deep neural networks. *Science Advances*, 7(1):eabd6127, 2021.

[37] Diederik P Kingma and J Adam Ba. A method for stochastic optimization. arxiv 2014. *arXiv preprint arXiv:1412.6980*, 434, 2019.

[38] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'12, page 1097–1105, Red Hook, NY, USA, 2012. Curran Associates Inc.

[39] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[40] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

[41] Chien Heng Liao, Keith J Worsley, J-B Poline, John AD Aston, Gary H Duncan, and Alan C Evans. Estimating the delay of the fmri response. *NeuroImage*, 16(3):593–606, 2002.

[42] Martin A Lindquist and Tor D Wager. Principles of functional magnetic resonance imaging. *Handbook of neuroimaging data analysis*, pages 3–48, 2014.

[43] William Lotter, Gabriel Kreiman, and David Cox. A neural network trained to predict future video frames mimics critical properties of biological neuronal responses and perception. *arXiv preprint arXiv:1805.10734*, 2018.

[44] Christian K Machens, Michael S Wehr, and Anthony M Zador. Linearity of cortical receptive fields measured with natural sounds. *Journal of Neuroscience*, 24(5):1089–1100, 2004.

[45] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016.

[46] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

[47] Yalda Mohsenzadeh, Caitlin Mullin, Dimitrios Pantazis, and Aude Oliva. Emergence of topographical correspondences between deep neural network and human ventral visual cortex, 2019.

[48] Thomas Naselaris, Emily Allen, and Kendrick Kay. Extensive sampling for complete models of individual brains. *Current Opinion in Behavioral Sciences*, 40:45–51, 2021.

[49] Thomas Naselaris, Kendrick N Kay, Shinji Nishimoto, and Jack L Gallant. Encoding and decoding in fmri. *Neuroimage*, 56(2):400–410, 2011.

[50] Alex Nichol, Vicki Pfau, Christopher Hesse, Oleg Klimov, and John Schulman. Gotta learn fast: A new benchmark for generalization in rl. *arXiv preprint arXiv:1804.03720*, 2018.

[51] Yuval Nir, Lior Fisch, Roy Mukamel, Hagar Gelbard-Sagiv, Amos Arieli, Itzhak Fried, and Rafael Malach. Coupling between neuronal firing rate, gamma lfp, and bold fmri is related to interneuronal correlations. *Current biology*, 17(15):1275–1285, 2007.

[52] Chigozie Nwankpa, Winifred Ijomah, Anthony Gachagan, and Stephen Marshall. Activation functions: Comparison of trends in practice and research for deep learning. *CoRR*, abs/1811.03378, 2018.

[53] Seiji Ogawa and Tso-Ming Lee. Magnetic resonance imaging of blood vessels at high fields: in vivo and in vitro measurements and image simulation. *Magnetic resonance in medicine*, 16(1):9–18, 1990.

[54] C. Olah. “understanding lstm networks,” net: http://colah. github. io/posts/2015-08-understandinglstms. 2015.

[55] Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997.

[56] Marc Palaus, Elena M Marron, Raquel Viejo-Sobera, and Diego Redolar-Ripoll. Neural basis of video gaming: A systematic review. *Frontiers in human neuroscience*, page 248, 2017.

[57] Jonathan Peirce, Jeremy R Gray, Sol Simpson, Michael MacAskill, Richard Höchenberger, Hiroyuki Sogo, Erik Kastman, and Jonas Kristoffer Lindeløv. Psychopy2: Experiments in behavior made easy. *Behavior research methods*, 51(1):195–203, 2019.

[58] Dean A. Pomerleau. Alvinn: An autonomous land vehicle in a neural network. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems 1*, pages 305–313. Morgan-Kaufmann, 1989.

[59] Jonathan D Power, Bradley L Schlaggar, and Steven E Petersen. Recent progress and outstanding issues in motion correction in resting state fmri. *Neuroimage*, 105:536–551, 2015.

[60] Arie Rachmad Syulistyo, Dhebys Suryani, and Pramana Yoga Saputra. Sibi (sistem isyarat bahasa indonesia) translation using convolutional neural network (cnn). *IOP Conference Series: Materials Science and Engineering*, 732:012082, 01 2020.

[61] Blake A Richards, Timothy P Lillicrap, Philippe Beaudoin, Yoshua Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, Rui Ponte Costa, Archy de Berker, Surya Ganguli, et al. A deep learning framework for neuroscience. *Nature neuroscience*, 22(11):1761–1770, 2019.

[62] Stéphane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. No-regret reductions for imitation learning and structured prediction. *CoRR*, abs/1011.0686, 2010.

[63] Andrew Saxe, Stephanie Nelli, and Christopher Summerfield. If deep learning is the answer, what is the question? *Nature Reviews Neuroscience*, 22(1):55–67, 2021.

[64] Martin Schrimpf, Idan Blank, Greta Tuckute, Carina Kauf, Eghbal A Hosseini, Nancy Kanwisher, Joshua Tenenbaum, and Evelina Fedorenko. The neural architecture of language: Integrative reverse-engineering converges on a model for predictive processing. *BioRxiv*, 2020.

[65] Martin Schrimpf, Jonas Kubilius, Ha Hong, Najib J Majaj, Rishi Rajalingham, Elias B Issa, Kohitij Kar, Pouya Bashivan, Jonathan Prescott-Roy, Franziska Geiger, et al. Brain-score: Which artificial neural network for object recognition is most brain-like? *BioRxiv*, page 407007, 2018.

[66] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[67] Dan Schwartz, Mariya Toneva, and Leila Wehbe. Inducing brain-relevant bias in natural language processing models. *Advances in neural information processing systems*, 32, 2019.

[68] Kawin Setsompop, Julien Cohen-Adad, Borjan A Gagoski, Tommi Raij, Anastasia Yendiki, Boris Keil, Van J Wedeen, and Lawrence L Wald. Improving diffusion mri using simultaneous multi-slice echo planar imaging. *Neuroimage*, 63(1):569–580, 2012.

[69] Leslie N. Smith. A disciplined approach to neural network hyper-parameters: Part 1 - learning rate, batch size, momentum, and weight decay. *CoRR*, abs/1803.09820, 2018.

[70] Thomas C. Sprague, Geoffrey M. Boynton, and John T. Serences. The importance of considering model choices when interpreting results in computational neuroimaging. *eNeuro*, 6(6), 2019.

[71] Anup Tuladhar, Jasmine A Moore, Zahinoor Ismail, and Nils D Forkert. Modeling neurodegeneration in silico with deep learning. *Frontiers in Neuroinformatics*, page 56, 2021.

[72] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018.

[73] Sebastian Urchs, Jonathan Armoza, Clara Moreau, Yassine Benhajali, Jolène St-Aubin, Pierre Orban, and Pierre Bellec. Mist: A multi-resolution parcellation of functional brain networks. *MNI Open Research*, 1:3, 2019.

[74] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. *CoRR*, abs/1509.06461, 2015.

[75] Rufin VanRullen and Ryota Kanai. Deep learning and the global workspace theory. *Trends in Neurosciences*, 44(9):692–704, 2021.

[76] Aria Wang, Michael Tarr, and Leila Wehbe. Neural taskonomy: Inferring the similarity of task-derived representations from brain activity. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, pages 15501–15511. Curran Associates, Inc., 2019.

[77] Panqu Wang, Vicente Malave, and Ben Cipollini. Encoding voxels with deep learning. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 35(48):15769–15771, Dec 2015. 26631460[pmid].

[78] Panqu Wang, Vicente Malave, and Ben Cipollini. Encoding voxels with deep learning. *Journal of Neuroscience*, 35(48):15769–15771, 2015.

[79] Lilian Weng. https://lilianweng.github.io/posts/2018-02-19-rl-overview/.

[80] Michael C-K Wu, Stephen V David, and Jack L Gallant. Complete functional characterization of sensory neurons by system identification. *Annu. Rev. Neurosci.*, 29:477–505, 2006.

[81] Daniel LK Yamins and James J DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nature neuroscience*, 19(3):356–365, 2016.

[82] Daniel LK Yamins, Ha Hong, Charles F Cadieu, Ethan A Solomon, Darren Seibert, and James J DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the national academy of sciences*, 111(23):8619–8624, 2014.

[83] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? *arXiv preprint arXiv:1411.1792*, 2014.

[84] Dong Yu and Michael L Seltzer. Improved bottleneck features using pretrained deep neural networks. In *Twelfth annual conference of the international speech communication association*, 2011.

[85] Anthony Zador, Blake Richards, Bence Ölveczky, Sean Escola, Yoshua Bengio, Kwabena Boahen, Matthew Botvinick, Dmitri Chklovskii, Anne Churchland, Claudia Clopath, et al. Toward next-generation artificial intelligence: catalyzing the neuroai revolution. *arXiv preprint arXiv:2210.08340*, 2022.

[86] Tong Zhang, Bin Yu, et al. Boosting with early stopping: Convergence and consistency. *The Annals of Statistics*, 33(4):1538–1579, 2005.

[87] Boyuan Zheng, Sunny Verma, Jianlong Zhou, Ivor W Tsang, and Fang Chen. Imitation learning: Progress, taxonomies and challenges. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–16, 2022.