

Electronic Thesis and Dissertation Repository

---

8-24-2021 2:00 PM

## A Generative-Discriminative Approach to Human Brain Mapping

Deepanshu Wadhwa, *The University of Western Ontario*

Supervisor: Diedrichsen, Jörn, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Master of Science degree in  
Computer Science

© Deepanshu Wadhwa 2021

Follow this and additional works at: <https://ir.lib.uwo.ca/etd>



Part of the [Artificial Intelligence and Robotics Commons](#), [Computational Neuroscience Commons](#),  
and the [Data Science Commons](#)

---

### Recommended Citation

Wadhwa, Deepanshu, "A Generative-Discriminative Approach to Human Brain Mapping" (2021). *Electronic Thesis and Dissertation Repository*. 8092.

<https://ir.lib.uwo.ca/etd/8092>

This Dissertation/Thesis is brought to you for free and open access by Scholarship@Western. It has been accepted for inclusion in Electronic Thesis and Dissertation Repository by an authorized administrator of Scholarship@Western. For more information, please contact [wlsadmin@uwo.ca](mailto:wlsadmin@uwo.ca).

# Abstract

During everyday behaviours, the brain shows complex spatial patterns of activity. These activity maps are very replicable within an individual, but vary significantly across individuals, even though they are evoked by the same behaviour. It is unknown how differences in these spatial patterns relate to differences in behavior or function. More fundamentally, the structural, developmental, and genetic factors that determine the spatial organisation of these brain maps in each individual are unclear. Here we propose a new quantitative approach for uncovering the basic principles by which functional brain maps are organized. We propose to take an generative-discriminative approach to human brain mapping, with the fundamental idea that, if we understand the underlying principles that organises brain activity maps, we should be able to generate sets of artificial maps that are indistinguishable from real ones. Different generative models are tested by a series of adversarial, classifier models, ranging from linear classifiers based on specific marginal statistics, to a full convolutional neural network. We apply our new framework to a collection of measured finger activity maps measures with fMRI in the human sensori-motor cortex (N=50). To account for characteristics of the brain maps that depend on the specific measurement process (spatial resolution of fMRI, signal-to-noise, etc) we supplemented the generative process with a measurement model. Initial results clearly demonstrate that the matching of simple marginal statistics (covariance and smoothness of activity pattern) is sufficient to fool the human eye, but not a more systematic machine learning approach. The proposed evaluation framework therefore opens up a pathway for discovering specific characteristics of brain activity maps that are important to explain function or individual differences.

**Keywords:** Computational Neuroscience, Brain Mapping, Deep Learning, Generative Modelling, Finger Movement

# Summary for lay audience

Our brains are structured differently. Each individual has a unique brain structure that is shaped by multiple factors, including genetic and developmental factors. Additionally, different parts of our brain work differently and are structured in a way that helps them specialize in their specific function/s. The functioning of the brain, and the resultant activity in different regions of the brain can be recorded and measured using a variety of imaging techniques, including fMRI. These patterns of activity (activity maps) vary widely across individuals and are important in understanding the structure and the functioning of the brain.

My thesis introduces a novel approach towards brain mapping and is aimed at determining the principles that characterize the organisation of the brain activity maps across individuals. Discovering these underlying principles is important, as they have the potential to relate to individual differences in function, or to characterize patient groups with dysfunction in specific domains. This research specifically focuses on understanding the organization of activity patterns in the sensorimotor cortex of the brain.

For understanding these principles, we developed generative models to generate artificial maps based on a dataset of real maps. The generative models were designed to test various hypotheses about the structure of the brain activity patterns. For evaluating the generative models, we developed a comprehensive evaluation framework that is adaptable to any generative model under test. The evaluation framework provides insights into the points of strengths and weaknesses in a generative model. The specific patterns of strengths and weaknesses then guide the improvement of the corresponding generative model/s. The generative model and the evaluation framework hence act as adversaries where the generative model tries to fool the evaluation framework by generating fake activity maps.

This research hence contributes to the broader scientific in multiple ways

– as the proof of insufficiency of known organizational principles as the sole organizational factors in the activity maps, as the proponent of a new generative-discriminative approach towards brain mapping; as a descriptor of a modelling and measurement simulation process for generating artificial brain maps; and as a source of curated data for further research and development.

# Acknowledgement

Venturing into new, unexplored avenues of research is never an easy task. However, the presence of an amazing supervisor to guide you through the path certainly makes the journey a lot less intimidating. I would like to thank my supervisor Dr. Jörn Diedrichsen for his constant support, guidance, and motivation throughout this research. He not only helped me navigate through the challenges of academia, but through his unwavering support and valuable insights, he helped me grow intellectually. I learnt invaluable lessons and skills from him, for which I will forever be indebted. I thank him for believing in me and for providing me an opportunity to learn from him.

I am also grateful to the amazing team at Diedrichsen Lab, and specially Spencer Arbuckle, for all the help and support throughout the research. My interactions with the extremely talented members at the lab helped me delve deeper into the field of computational neuroscience, and truly appreciate the complexities and the charm of this field.

Further, I am thankful to the Department of Computer Science, and the Brain and Mind Institute team for providing me the resources and a conducive environment to carry out the research.

Finally, I am deeply grateful to my family and my friends who despite being away from me, served as a bedrock of emotional support as I made my way through the research.

# Contents

|   |             |
|---|-------------|
| <b>Abstract</b>   | <b>ii</b>   |
| <b>Summary for lay audience</b>   | <b>iii</b>  |
| <b>Acknowledgement</b>  | <b>v</b>    |
| <b>Table of Contents</b>  | <b>vi</b>   |
| <b>List of Tables</b>   | <b>viii</b> |
| <b>List of Figures</b>  | <b>ix</b>   |
| <b>1 Introduction</b>   | <b>1</b>    |
| <b>2 Methodology</b>  | <b>4</b>    |
| 2.1 Overview . . . . .  | 4           |
| 2.2 Measurement and Dataset . . . . .   | 5           |
| 2.3 Simulated measurement process . . . . .                                   | 8           |
| 2.4 Generative process . . . . .  | 11          |
| 2.4.1 Baseline Model . . . . .  | 11          |
| 2.4.2 Usage Model . . . . .   | 13          |
| 2.5 Evaluation process . . . . .  | 13          |
| 2.5.1 Convolutional Neural Network over finger maps (5-<br>map CNN) . . . . . | 14          |
| 2.5.2 Convolutional Neural Network over mean subject maps                     | 16          |
| 2.5.3 Logistic Regression over component strengths . . . . .                  | 18          |
| 2.5.4 Neural Network over component strengths . . . . .                       | 19          |
| 2.5.5 Logistic Regression over the finger covariance matrix .                 | 19          |
| 2.5.6 Neural Network over the finger covariance matrix . . .                  | 20          |

|          |  |           |
|----------|--|-----------|
| <b>3</b> | <b>Results</b>   | <b>21</b> |
| 3.1      | General characteristics of the generated maps are well matched to real data . . . . .        | 21        |
| 3.2      | General statistical characteristics are not sufficient to generate artificial maps . . . . . | 24        |
| 3.3      | Failure of generative models is already visible in the mean maps                             | 25        |
| 3.4      | Non-Gaussianity of real activity data is one contributing factor                             | 26        |
| <b>4</b> | <b>Discussion</b>  | <b>28</b> |
|          | <b>Symbol Table</b>  | <b>32</b> |
|          | <b>Bibliography</b>  | <b>33</b> |
|          | <b>Curriculum Vitae</b>  | <b>35</b> |

# List of Tables

- 2.1 **Experimental characteristics for samples in the dataset.**  
Sub: Number of Subjects, vox: isometric voxel size in mm, TR: time of repetition / sampling duration in s, FS: Field Strength in T, Runs: Number of Run per experiment, TFR: Number of trials per Run, D: Length (Duration) of a single trial in s.  
9
- 2.2 **Properties of the layers in the Convolution Neural Network.** Layer: layer number, Filter: number of filters in the layer, Strides: Number of strides used by kernels, Batch-Norm: whether Batch Normalization used, LeakyReLU: whether LeakyReLU activation was used, Padding: type of padding used (valid padding/ same padding)  
15
- 3.1 **Test statistics for t-tests over component strengths.**  
The table shows the results from the t-test carried out on the component strengths calculated from the measured maps from the usage model vs the corresponding component strength from the real measured maps.  
23

# List of Figures

|     |  |    |
|-----|--|----|
| 2.1 | <b>Overview of generative-discriminative approach.</b> The true activity maps in the M1 region are measured in a volumetric space using fMRI, mapped to a reconstructed cortical surface, and finally to a regular surface resulting in measured maps, $x_i$ . V-N encapsulates the volume to reconstructed cortical surface transformation, and N-S encapsulates the reconstructed surface to regular pixel surface transformation. We estimate the overall mean map $x_{\dots}$ , spatial covariance $\Sigma_p$ , and finger covariance $G$ from the measured maps. To generate artificial maps $c_i$ on the surface, the mean $x_{\dots}$ , subject component $s_i$ , and finger component $f_{i,j}$ are combined in both the models. $s_i$ is sampled from a multivariate normal distribution with covariance $\Sigma_p$ . For sampling $f_{i,j}$ , while the usage model uses $G$ for inter-finger covariance and $\Sigma_p$ for spatial covariance, the baseline model generates independent finger maps with spatial covariance $\Sigma_p$ . The generated activity maps are passed through a simulated measurement process that mimics the measurement process of a real sample. Matrices $(P, V)$ transform the maps to and from the volumetric space. Noise estimate $z_{i,j,k}$ is added to the maps while in the volumetric space to obtain generated measured maps, $y_i$ . . . . . | 6  |
| 2.2 | <b>5-map Convolutional Neural Network.</b> Each convolutional layer consists of a set of filters. The dimension of the output of a given layer is: number of filters x height x width. Layers 1, 2, 3, 4, 5 represent the outputs of the corresponding layers. . . . .   | 16 |

|     |  |    |
|-----|--|----|
| 2.3 | <b>Kernels in a filter..</b> Each filter $R$ in a convolutional layer consists of a set of kernels. In the first layer, each filter is made of 5 different 4x4 kernels (K1, K2, K3, K4, K5) - one kernel for each input channel (finger activity map). The kernels move across their respective activity maps (with padding - gray) and their outputs at each step are added together to form the filter output. . . . . | 17 |
| 3.1 | <b>Real and generated digit maps.</b> Each row shows a sample consisting of activity maps for the 5 fingers. Top 3 rows are real samples, and the bottom three rows show three generated samples from the usage model, generated using the measurement model of the 3 real samples. . . . .  | 22 |
| 3.2 | <b>Classification accuracies (%) of the various classification models for the baseline (gray) and usage (red) model.</b> The classifier from left to right were component strengths with logistic regression, component strengths with neural network, G matrix with logistic regression, G matrix with neural network, subject means, 5-map CNN, 5-map CNN with histogram equalization. . . . .                         | 25 |
| 3.3 | <b>Real maps are slightly non-Gaussian.</b> The quantile-quantile plot of the activation values of the real maps vs the theoretical values from a normal distribution indicates that the values come from a more heavy-tailed distribution. . . . .  | 27 |

# Chapter 1

## Introduction

The brain shows complex spatial patterns of activity. In some regions, the activity patterns follow a clear spatial organization that is easily summarized. For example, the primary visual cortex is organised such that the center of the visual field (the retina) is represented at one end of the region, and the periphery at the other end. Thus, the region has a clear retinotopic organisation (Engel, Glover, and Wandell, 1997), which is relatively invariant across individuals. As a second example, in the primary somatosensory cortex (S1), neurons are broadly arranged according to limb preference, with neurons responding to stimulation to the same limb loosely clustering together. These regions therefore can be described by a somatotopic organisation (Penfield and Rasmussen, 1950). However, there are other brain regions where the activity maps do not appear to adhere to any clear spatial organization. Spatial activity patterns in these brain regions are often fractured and also highly variable across even normal, healthy individuals (Meier et al., 2008). Therefore, spatial variation in brain maps is likely due to random biological variation with very little functional relevance. My thesis is aimed at determining the principles that characterize the organisation of these maps across healthy individuals. Discovering these regularities is important, as these underlying principles, rather than the superficial spatial arrangement of activity patterns, has the potential to relate to individual differences in function, or to characterize patient groups with dysfunction in specific domains.

In this thesis, we focus on finger activity patterns in the primary motor cortex (M1). The production of finger movements relies on the activity of neurons in the hand area of M1, as lesions to M1 result in a loss of dexterous hand control (Liepert et al., 2005). However, how M1 is organized to control

hand movements remains unknown. Traditionally, finger maps were assumed to be organized in a somatotopic fashion in M1, with neurons in the thumb area being activated for thumb movements, and neurons in the index finger area being activated for index finger movements. This somatotopic organization was thought to be genetically defined, and thus quite stable across individuals. However, both neural recordings and functional MRI (fMRI) work have demonstrated that neurons are not organized according to a strict somatotopic principle in M1 (Schieber, 2001) (Ejaz, Hamada, and Diedrichsen, 2015). Instead, the activity patterns for movements of each individual finger were found to be highly overlapping and diffused. Moreover, the spatial activity patterns for movements of the same finger vary substantially across individuals, suggesting that a somatotopic arrangement alone cannot be the sole organizational principle of finger activity maps in M1.

More recently, Graziano and colleagues (Graziano and Aflalo, 2007) proposed that the diffuse and overlapping nature of activity maps in M1 reflect the natural statistics of movement during everyday life. Applying this idea to finger maps yields the prediction that fingers that typically move together during day-to-day activities evoke more similar patterns of neural activity, and vice versa. Using human fMRI, Ejaz and colleagues (Ejaz, Hamada, and Diedrichsen, 2015) demonstrated that the natural statistics of finger movements could explain the diffuse and fractured nature of finger activity maps in M1 and S1. This suggests that there exist organizing principles of neural populations that are not spatially constrained. Here, we investigated whether the natural statistics of everyday hand movements is a sufficient organizing principle or are there other defining regularities governing the organization of finger activity maps.

If the natural statistics of everyday hand movements are sufficient to account for the organization of finger activity maps, we should be able to generate artificial brain maps based on these principles that are indistinguishable from real, measured maps. The generated maps must be able to pass a rigorous evaluation process that is able to detect small differences between the real and measured maps that may not be apparent by pure visual inspection. Our project, therefore, aims to determine the organizing factors behind the finger representations in the motor cortex using a generative-discriminative approach.

In our project, we developed generative models to test the existing theories behind the organization of finger activity maps and further, and to potentially uncover and incorporate yet overlooked characteristics and fill

in the missing gaps in the known principles. The generative models were tested in the discriminative part of the framework using a set of classifiers to evaluate the similarity of the generated finger maps with the real maps. We started with a powerful convolutional neural network (Rawat and Wang, 2017), and then successively stripped the analysis down to increasingly simpler classifier models. We reasoned that a generative model that completely captures all relevant organizing principles of finger activity patterns would be able to fool each level of the entire evaluation framework. Conversely, points of "failure" (i.e., instances when the classifiers could distinguish between real and generated maps) provided insight into what features were not fully captured by the generative models. The failure of a model with a given set of features successfully replicated suggested the insufficiency of that set of features in explaining the sensorimotor organization.

# Chapter 2

## Methodology

### 2.1 Overview

The aim of this thesis was to test our current understanding of the organizing principles that shape the finger activity patterns in sensorimotor cortex. We address this by building generative models that produced sets of spatial activity maps, one map for each of the five fingers. To evaluate how good a generative model was, we measured whether the produced finger maps could be distinguished by different classifier models from real maps measured from human participants (Figure 2.1).

The initial generative model, the baseline model, was based on the idea that a given activity map has three primary constituent components - the mean activity map, deviation of individual subjects from that mean, and idiosyncratic finger-specific patterns around the subject mean, with each component having a characteristic covariance structure. The second generative model, the usage model, took into consideration the presence of an inter-finger covariance structure based on the natural statistics of hand movement. Thus, this generative model encapsulated the hypothesis that the natural statistics of hand movement is a sufficient principle that explains the organisation of finger maps in M1 (Ejaz, Hamada, and Diedrichsen, 2015) (Graziano and Aflalo, 2007).

To evaluate the developed generative models, we compared the generated maps to a library of real finger maps, measured with fMRI. The comparison was done using a set of classifiers, ranging from simple classifiers based on some restricted features of the maps (marginal statistics) to a trained

convolution neural network. The whole process hence took the form of an adversarial process where the two models, the generative model and the classifier, functioned as adversaries (Gershman, 2019). Although the generative model and the classifier models can potentially be linked together by a loss function in future, this was not the case for the model developed here, given the emphasis on trying to determine the factors that characterize the real finger maps.

A generative process that can successfully fool the complete set of such classification models can hence be considered a reliable model for generation of finger maps. In contrast, a non-successful model will fail at some stage during the entire process. The particular pattern of failures will provide insights into the missing or unrealistic aspects in the generated data.

One complication in this framework is that the real activity maps are influenced by the measurement process. For example, the spatial resolution of the fMRI acquisition, the orientation of the slices relative to the folding of the brain, and the signal-to-noise ratio all influence the measured maps to some degree. Since the source of these features is the scanning process and not the brain’s biological process itself, they are not relevant for studying the biological patterns. However, the absence of these features in the generated maps can lead to a classifier successfully differentiating between the real and the generated maps. Therefore, to minimize the influence of these features in the evaluation process, we simulated a measurement process that we applied to the generated activity maps. The specific characteristics of the measurement process applied to an input generated activity map were dependent on a particular subject selected from the library of real maps.

## 2.2 Measurement and Dataset

The neural activity in the sensorimotor cortex, evoked in the participants during isometric finger movements, was recorded for the study. We used functional Magnetic Resonance Imaging (fMRI), a non-invasive imaging technique, to record the activity. fMRI measures the changes in blood flow occurring with brain activity. In the brain, the blood flow is highly regulated based on the oxygen requirements of a particular region. An increase in the activity in a certain region of the brain during a task calls for a higher oxygen requirement. The extraction of oxygen from the local capillaries in response to this demand leads to a sudden drop in the oxygenated blood level and a

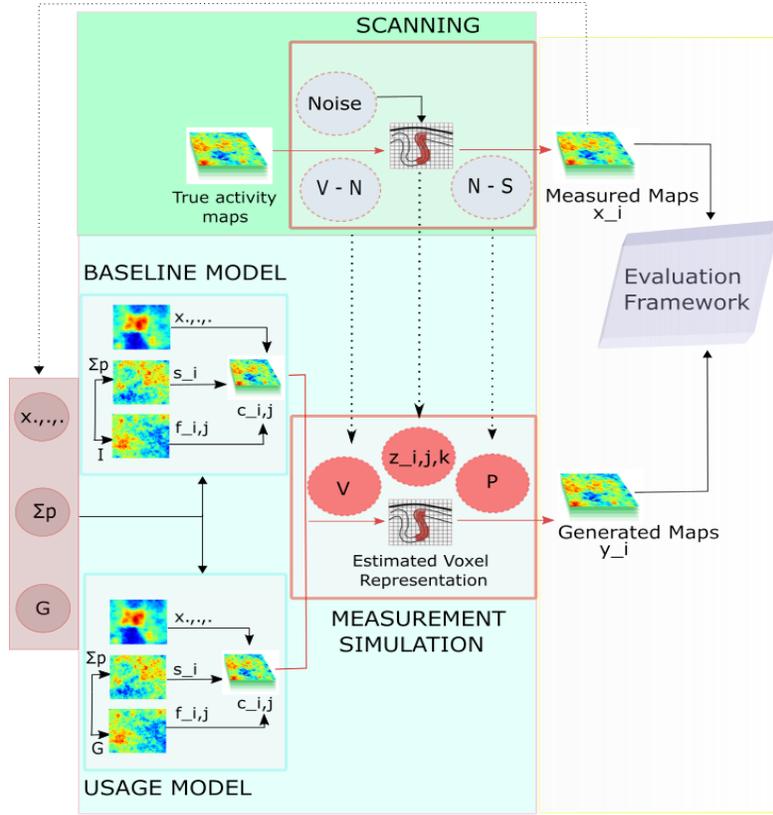


Figure 2.1: **Overview of generative-discriminative approach.** The true activity maps in the M1 region are measured in a volumetric space using fMRI, mapped to a reconstructed cortical surface, and finally to a regular surface resulting in measured maps,  $x_i$ . V-N encapsulates the volume to reconstructed cortical surface transformation, and N-S encapsulates the reconstructed surface to regular pixel surface transformation. We estimate the overall mean map  $x_{i,\dots}$ , spatial covariance  $\Sigma_p$ , and finger covariance  $G$  from the measured maps. To generate artificial maps  $c_i$  on the surface, the mean  $x_{i,\dots}$ , subject component  $s_i$ , and finger component  $f_{i,j}$  are combined in both the models.  $s_i$  is sampled from a multivariate normal distribution with covariance  $\Sigma_p$ . For sampling  $f_{i,j}$ , while the usage model uses  $G$  for inter-finger covariance and  $\Sigma_p$  for spatial covariance, the baseline model generates independent finger maps with spatial covariance  $\Sigma_p$ . The generated activity maps are passed through a simulated measurement process that mimics the measurement process of a real sample. Matrices  $(P, V)$  transform the maps to and from the volumetric space. Noise estimate  $z_{i,j,k}$  is added to the maps while in the volumetric space to obtain generated measured maps,  $y_i$ .

rise in the deoxygenated blood level in the region. This is followed by an increased delivery of oxygenated blood after a lag of 2-6 seconds, a response that clearly over-compensates for the increased demand. The magnetic properties of oxygenated blood are different from deoxygenated blood and this property is exploited to record the brain activity during the rebound using the Blood-Oxygen-Level Dependent (BOLD) response (Glover, 2011).

The fMRI BOLD response is hence, an indirect measure of the neural activity. The activity response in BOLD fMRI is recorded with respect to the resting baseline with a positive activity value signifying a higher neural activity and a negative value signifying a lower neural activity. A single fMRI image is a volumetric digit image, with voxels (i.e., volumetric pixels) as the basic observational units. Depending on the resolution, each voxel can contain hundreds-of-thousands of neurons.

The dataset used for this project is a collection of 50 sets of five finger activity patterns. The data is pooled across several fMRI studies, conducted over a 10 year period on a number of different MR scanners (7T and 3T) (Table 2.1). Each set of maps consists of activity patterns (first-level GLM beta weights) from the sensorimotor cortex of the left hemisphere evoked during individuated finger presses with the right hand on a piano-like keyboard device. In one experiment, participants also made individuated finger presses with the left hand. For this data, the patterns from the right hemisphere were projected to the left hemisphere using a mirror-symmetric surface-based alignment using Freesurfer (workbench FS\_sym surface) (Fischl, 2012). The left and right hemisphere patterns from the same participants were treated as independent samples in the analyses.

For the scanning, the subjects were asked to perform isometric finger movements (Table 2.1) and the scans were recorded over multiple runs per subject. The signal was analyzed using a Generalized Linear Model (GLM) with multiple regressors per run depending on the particular experiment. The regressors were boxcar functions that started at the moment of the first finger press and lasted for a few seconds of the trial duration. These regressors were then convolved with a hemodynamic response function, eventually giving us a temporally invariant value for each measured voxel relative to the resting baseline.

In order to map the activity to specific regions of the cortex, the anatomical structure of each subject's cortex was reconstructed based on T1 weighted anatomical scans using freesurfer (Fischl, 2012). The software approximates the surface between the white matter and the gray matter of the cortex, as

well as the boundary between the gray matter and the surrounding cerebrospinal fluid. Individual activity patterns were then projected on the corresponding cortical surface using this inner and outer surface. All voxels touching the line between the two surfaces were averaged and the averaged values assigned to the surface nodes. To avoid the mixing of signals between M1 and S1 across the central sulcus, any voxels touching the surface at multiple places were removed from further analysis.

Finally, we resampled the finger maps into a regular pixel grid (128x128) using a nearest neighbor approach. This was done to optimize the maps for analysis, and for evaluation using the machine learning based evaluation framework. The maps obtained at this stage were the real measured maps used for further analysis. The scanned finger maps were hence converted into a regularized surface space where later on, the evaluation of the generated maps took place.

Overall, the dataset consisted of  $n = 50$  samples (subjects and hemispheres). Each sample consisted of  $k$  measurements for each of the 5 fingers, measured in a separate imaging run (typically 8-12 runs). Each individual map is denoted by  $x_{i,j,k}$ , the scanned map for sample  $i$ , finger  $j$ , and run  $k$ . The resolution of functional images in each experiment differed. For application of a convolutional neural network, all surface activity patterns were resampled onto a regular 128x128 pixel grid, with  $x_{i,j,k}[a, b]$  representing the activity value of the pixel at  $a^{th}$  row and  $b^{th}$  column.

## 2.3 Simulated measurement process

The simulated measurement process used in the artificial map generation process simulates the effects of the scanning process (spatial sampling, noise), and the subsequent data processing (volume to surface mapping) on the generated activity maps.

Given that the spatial resolution, the signal-to-noise ratio, and the cortical folding differed across different sets of finger maps, we build a simulation of the measurement process for each sample (set of finger maps) in the dataset. The mapping process and the noise-covariance structure used during the simulation was there specific to a particular subject  $i$  from the real data set. To ensure that the classification models would focus on differences in the true maps, rather than differences in the measurement process, we always compare true maps to generated maps that used the measurement process

Table 2.1: **Experimental characteristics for samples in the dataset.**  
 Sub: Number of Subjects, vox: isometric voxel size in mm, TR: time of repetition / sampling duration in s, FS: Field Strength in T, Runs: Number of Run per experiment, TFR: Number of trials per Run, D: Length (Duration) of a single trial in s.

| <b>Description</b>   | <b>Sub</b> | <b>vox</b> | <b>TR</b> | <b>FS</b> | <b>Hem</b> | <b>Run</b> | <b>TFR</b> | <b>D</b> |
|--|------------|------------|-----------|-----------|------------|------------|------------|----------|
| Left and Right hand single finger press (Diedrichsen, Wiestler, and Krakauer, 2013)                        | 6          | 2.3        | 2.72      | 3T        | LR         | 8          | 3          | 8        |
| Right hand single finger flexion and extension presses at multiple forces (Arbuckle, Weiler, et al., 2020) | 9          | 1.5        | 1.5       | 7T        | L          | 8          | 6          | 4.5      |
| Right hand single finger pressing at multiple speeds   | 9          | 2          | 0.7       | 7T        | L          | 8          | 8          | 6        |
| Right hand single finger pressing at multiple speeds (Arbuckle, Yokoi, et al., 2019)                       | 8          | 1.4        | 1         | 7T        | L          | 8          | 8          | 6        |
| Right hand single and multi finger chord presses (Ejaz, Hamada, and Diedrichsen, 2015)                     | 8          | 2.3        | 2.72      | 3T        | L          | 24         | 1          | 10.8     |
| Right hand single finger presses and sensory stimulation (Berlot et al., 2019)                             | 8          | 1.4        | 3         | 7T        | L          | 8          | 3          | 8.2      |

for that sample  $i$ .

The generative model generates the maps in the surface space  $(c_i, j)$ . The simulated measurement process was carried out in three steps: First, to mimic the way that evoked activity on the cortical surface is sampled by the scanner, we sampled the generated surface maps (made up of pixels) into a voxel grid using a pixel-to-voxel transformation matrix  $(P_i)$  for subject  $i$ . The summed weights of each pixel sampled into each voxel was normalized to ensure that voxels which sampled more pixels did not have higher mean activation.

Second, fMRI data have a characteristically low signal-to-noise ratio, which also varies across studies. To match this aspect in our simulated maps, we estimated strength of the signal  $(\omega_i)$  and the standard deviation of the noise  $(\eta_i)$  from the corresponding real sampled maps  $x_i, j, k$ . We also estimated the empirical voxel-to-voxel covariance matrix from the GLM residuals  $(\Sigma_{\epsilon,i})$ . For the generation of each of the measured maps we added a noise map  $(z_i, j, k)$  drawn from a multi-variate normal distribution with the corresponding noise spatial covariance matrix, and was independent across each finger and run per subject.

Third, the maps in voxel space were then sampled back onto the group surface using the subject’s voxel-to-pixel transformation matrix  $(V_i)$  for subject  $i$ . The voxel to pixel transformation matrix represents the relative weighting of each voxel as they are mapped to the pixelated surface grid (i.e. surface space). The process was designed to exactly mimic the process of mapping voxel-based fMRI data onto the surface, and from the surface to the regular 128x128 grid.

Overall, the measurement process that transforms a generated activity map  $c_{i,j}$  into a generated measured map can be summarized as:

$$y_{i,j,k} = P_i(V_i(\omega_i c_j) + z_{i,j,k}) \tag{2.1}$$

where  $c \in C$  is a set of 5 activity maps belonging to the 5 fingers of a generated sample.

The overall measurement process is independent of the generative model, and thus can be readily applied to patterns from any prescribed generative model.

## 2.4 Generative process

The models that generated activity maps were designed to conceptualize specific hypotheses about the principles that underlie the organization of brain activity maps. A generative process that produces activity maps indistinguishable from the real maps can be considered a complete model of the brain activity patterns of interest.

We developed two generative models to reproduce the underlying organizational principles. The first, a baseline statistical model, simply attempts to replicate the mean, variance, and spatial smoothness of the real maps, as well as the variances explained by difference between fingers or individuals. The second, the usage model, then adds a specific co-variance structure across the 5 fingers as another important criterion. This co-variance matrix encapsulates the hypothesis that the organisation of finger representations in sensory motor cortex is solely determined by the statistics of natural movement (Ejaz, Hamada, and Diedrichsen, 2015).

### 2.4.1 Baseline Model

The baseline model is based on the assumption that the activity maps are random and don't have any underlying structure. There is no somatotopic structure present and there is no specific organization of fingers.

To generate actual activity maps, the model needs to specify certain important characteristics, including the spatial distribution of the mean activity, the amount of inter-subject variability, the amount of differences between fingers, and the spatial smoothness of the maps. To achieve this, the generated maps were made from three primary components - the mean activity map  $x_{.,.,.}$ , deviation of individual subjects from that mean (called the subject specific component  $s_i$  for subject  $i$ ), and the idiosyncratic finger-specific patterns around the subject mean (called the finger specific component,  $f_{i,j}$  for subject  $i$ , finger  $j$ ). The generative model for the map for subject  $i$ , finger  $j$  was:

$$c_{i,j} = x_{.,.,.} + \mu s_i + \nu f_{i,j} \quad (2.2)$$

The mean activity map,  $x_{.,.,.}$ , was directly computed from all real maps. The dot-notation indicates that we averaged across samples ( $i$ ), fingers ( $j$ ), and runs ( $k$ ). The mean map was scaled to have a variance across pixels

of 1. The subject specific deviation maps were sampled from a multivariate normal with a mean of 0 and a covariance structure  $\Sigma_p$  across pixels, which determined the spatial smoothness. Finally, the finger-specific patterns were independently sampled from the matrix normal distribution with covariance matrix  $\Sigma_p$  as the column covariance. The subject and finger components were maps were scaled by a factor  $(\mu, \nu)$  reflecting their overall strength.

These subject and the finger scaling factors were estimated from the real samples in proportion to a unity mean activity pattern  $(x_{i,\dots})$  strength.

For estimating the scaling factors for the components, we first computed their strengths from the real measured maps. The strength of the mean activity was computed as the slope  $\zeta_i$  of the regression from the overall mean map to the subject mean:

$$x_{i,\dots} = \zeta_i x_{\dots} + \epsilon_i \quad (2.3)$$

The subject specific component strength,  $\Omega(s_i)$ , also called the residual strength, was measured as the strength of deviation of a subject from the overall mean map scaled by the corresponding slope  $\zeta_i$  calculated above:

$$\Omega(s_i) = \frac{\sum_{a,b=1}^{a,b=128} [(x_{i,\dots}[a,b] - \zeta_i x_{\dots}[a,b])^2]}{16384} \quad (2.4)$$

The finger specific component strength,  $\Omega(f_i)$ , quantifying the strength of the idiosyncratic finger-specific patterns around the subject mean was measured as:

$$\Omega(f_i) = \frac{\sum_{j=1}^5 \sum_{a,b=1}^{a,b=128} [(x_{i,j,\dots} - x_{i,\dots})^2]}{16384 \times 5} \quad (2.5)$$

The computed values of  $\zeta_i, \Omega(s_i), \Omega(f_i)$  were transformed into a matrix with each row of the matrix corresponding to the strengths of components present in a particular real sample. We then applied Singular Value Decomposition on the matrix to obtain the value of the subject and finger scaling factors  $\mu, \nu$ .

To match the smoothness of real maps, the generated maps had a specific pixel-to-pixel covariance matrix ( $\Sigma_p$ ). To estimate  $\Sigma_p$ , we plotted the spatial distance of two pixels on the surface ( $d_{a,b}$ ) against the covariance of their activity values across samples in the de-means dataset  $(x_{i,\dots} - x_{\dots})$ . We then fit the overall relationship with an exponential function by minimizing the squared error between predicted and measure covariance.

$$\Sigma_P[a, b] = ue^{-qd[a,b]^w}$$

The parameters  $u, q, w$  were fitted to the overall data. Their values were hence determined to be 0.127, 0.118, 0.621 respectively. From these parameters we then generated our estimate of  $\Sigma_P$ .

### 2.4.2 Usage Model

The Usage model is based on the idea that the finger activity maps are organised according to the natural statistics of finger movements. (Ejaz, Hamada, and Diedrichsen, 2015). That is, two fingers that often move together in everyday life should also have highly overlapping activity maps. Importantly, this model does not propose a specific somatotopic organisation at all, any common spatial organisation would be an emergent property from the generative process (Graziano and Aflalo, 2007). The maps again, are assumed to be random but they now have a specific covariance structure.

The finger covariance matrix,  $G$ , was estimated as:  $G_{j,l} = cov[x_{i,j}, x_{i,l}]$  from the real samples.

The subject and the finger scaling factors,  $\mu, \nu$ , and the spatial smoothness  $\Sigma_P$  were estimated from the real samples in the same way as in the baseline model.

The finger components,  $f_{i,j}$  (size  $5 \times 16384$ ), for the usage model were hence sampled from a matrix normal distribution with row covariance  $G$  and column covariance  $\Sigma_P$ .

## 2.5 Evaluation process

We used a series of classification models of increasing complexity to test the generated measured maps against the library of real maps. By varying the complexity of the classifiers, we were able to identify what features the generative model failed to capture.

For evaluation, the real and the generated samples were combined to form a dataset  $L$  such that  $L = \{x_1, x_2, \dots, x_n, y_1, y_2, \dots, y_n\}$ . We hence had  $n$  real samples in the dataset, and  $n$  generated samples where  $y_i$  used the simulated measurement process based on  $x_i$ .  $\gamma_i \in \Gamma$  represented the corresponding label (0 for real samples, 1 for generated samples). Each classifier was evaluated by its cross-validated accuracy score (i.e., percent correct classification).

The core framework for generating activity maps, based on the idea that a given activity map has three primary constituent components - the mean activity map, deviation of individual subjects from that mean, and idiosyncratic finger-specific patterns around the subject mean, served as a guideline for developing the evaluation framework.

We used sampling for establishing a reliable threshold for performance evaluation, and for the categorization of the generative model as a success/failure. We hypothesized that if a classification model from the evaluation framework can achieve a higher than chance accuracy on a statistically significant sample size out of the 100 samples (subjects), we can consider the corresponding generative model to have failed. A 95% confidence interval with 5% margin of error gave a statistically significant sample size of 80 (out of a total of 100). A chance accuracy over 80 samples (and correct classification of other 20 samples) by the classifier gives a threshold score of 60%. Hence, we concluded that if any classifier from the evaluation framework can achieve an accuracy of greater than 60% (failure threshold), it would mean a failure of the generative model under consideration.

### **2.5.1 Convolutional Neural Network over finger maps (5-map CNN)**

The most powerful classifier considered here was a convolutional neural network that classified a single sample (set of 5 maps) as either real or generated. The output was a scalar probability of the input sample being real (as opposed to being generated/fake). Success of this model in accurately classifying the real and the generated samples showed the ineffectiveness of the generative model in modelling the overall finger maps correctly. The convolutional neural network is based on the discriminator architecture from the DCGAN model (Radford, Metz, and Chintala, 2015) and uses the following architecture:

The 5-map CNN  $H$ , processed the input through a series of five convolutional layers with LeakyReLU activation (Xu et al., 2015), and outputted the final probability through the output layer with a sigmoid activation function. It was a multi-channel convolutional network where each filter in each layer consisted of multiple kernels such that there was a unique kernel per filter for each of the input channels (each of the 5 finger maps in case of the first layer). The size of a filter in the first layer was hence  $5 \times q \times q$  where there

Table 2.2: **Properties of the layers in the Convolution Neural Network.** Layer: layer number, Filter: number of filters in the layer, Strides: Number of strides used by kernels, BatchNorm: whether Batch Normalization used, LeakyReLU: whether LeakyReLU activation was used, Padding: type of padding used (valid padding/ same padding)

| Layer  | Filters | Strides | BatchNorm | LeakyReLU | Padding |
|--------|---------|---------|-----------|-----------|---------|
| 1      | 32      | 2, 2    | No        | Yes       | Same    |
| 2      | 64      | 2, 2    | Yes       | Yes       | Same    |
| 3      | 128     | 2, 2    | Yes       | Yes       | Same    |
| 4      | 256     | 2, 2    | Yes       | Yes       | Same    |
| 5      | 512     | 2, 2    | Yes       | Yes       | Same    |
| Output | 1       | 1, 1    | No        | No        | Valid   |

were 5 input finger maps (and hence 5 kernels), with a  $q \times q$  sized kernel per finger. There was no sharing of weights between kernels belonging to the same filter. The output of all such kernels was summed to obtain the final output of the filter before the application of a non-linear activation function (LeakyReLU).

The first layer of the network ran a set of 32 filters across each finger in the input sample to generate the input for the next layer. The size of the convolutional kernels in each filter was so chosen such that it downsized the pixel dimensions of the input images by a factor of two at each step (e.g., from 128x128 to 64x64) and was set to be consistent at 4x4 across all layers. The number of filters at each step increased by a factor of two till the last layer where they were compressed down to obtain a 1x1x1 output. These hyperparameters are summarized in Table 2.2.

An activation value of 0.2 was used for the LeakyReLU (Leaky Rectified Linear Unit) function.

The output scalar  $\hat{\gamma}_i$  of the classifier H with parameters  $\theta_h$ , given a 5 x  $p$  (see 4) input  $l_i$ , can be described as:

$$\hat{\gamma}_i = H(l_i; \theta_h) \tag{2.6}$$

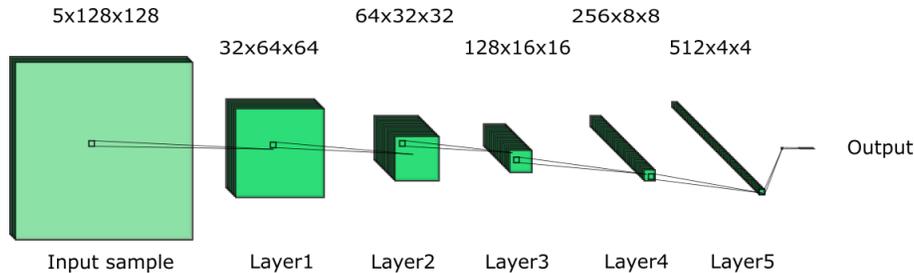


Figure 2.2: **5-map Convolutional Neural Network.** Each convolutional layer consists of a set of filters. The dimension of the output of a given layer is: number of filters x height x width. Layers 1, 2, 3, 4, 5 represent the outputs of the corresponding layers.

The network was trained using a Binary Cross Entropy (BCE) loss as its target loss function for minimization, and Adam optimizer (Kingma and Ba, 2014) as the corresponding optimizer. Given a dataset containing  $n$  samples, the loss function for the classifier can be written as:

$$\psi_H = \frac{-1}{m} \sum_{i=1}^m [\gamma_i \log(H(l_i)) + (1 - \gamma_i) \log(1 - H(l_i))] \quad (2.7)$$

The model was trained using K-fold cross validation (Yadav and Shukla, 2016) with 5 folds. The input and the output data was randomly split into 5 folds with the constraint that a real subject and the corresponding generated subject remained in the same fold. This prevented the model from simply learning the inter-experiment difference between the maps from different experiments (Table 2.1 which is not of interest. The model was then trained on 4 folds for 10 iterations (steps of optimizer) and validated on 1 fold. We recorded the average accuracy across 100 runs with the different validation folds.

## 2.5.2 Convolutional Neural Network over mean subject maps

To determine whether the CNN was able to distinguish real and generated maps based on the mean sample map, or based on the relationship between the 5 finger maps, we also trained a comparable CNN that only took the mean sample map  $l_{i,\dots}$  as an input. The mean sample maps are the maps resulting

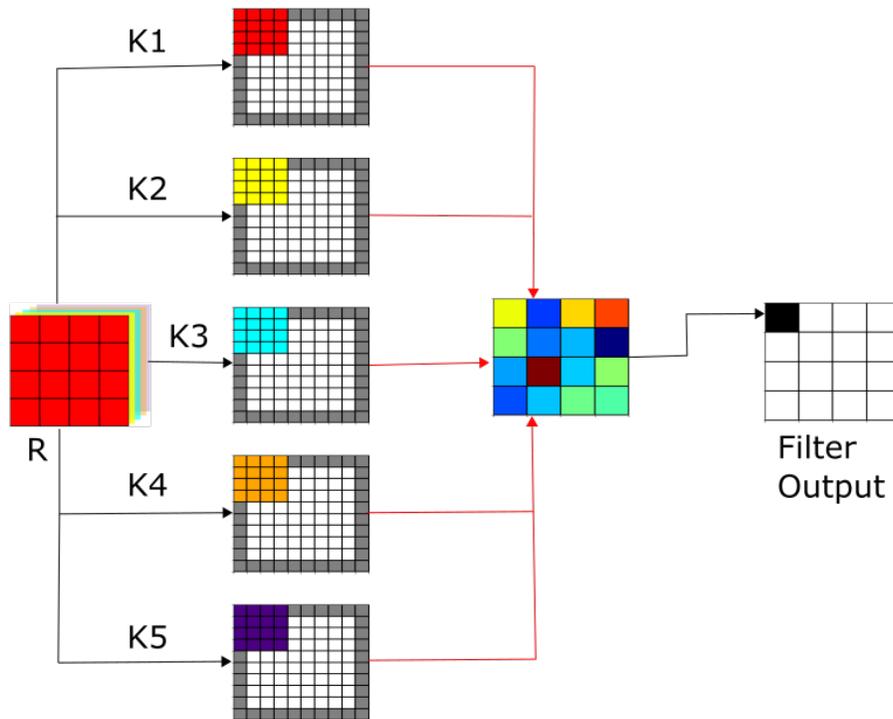


Figure 2.3: **Kernels in a filter.** Each filter  $R$  in a convolutional layer consists of a set of kernels. In the first layer, each filter is made of 5 different 4x4 kernels ( $K1$ ,  $K2$ ,  $K3$ ,  $K4$ ,  $K5$ ) - one kernel for each input channel (finger activity map). The kernels move across their respective activity maps (with padding - gray) and their outputs at each step are added together to form the filter output.

from the mean of all five maps belonging to a particular subject. Hence, this classification network compares the mean sample maps and models the relationship required to differentiate between these maps from the real and the generated datasets. This classification network helps narrow down the problem search process by modelling a simpler relationship between the maps as compared to the use of all five maps at once. Successful classification by this model means the generative model fails to successfully replicate the relationships required for generating mean sample patterns.

The network architecture, as well as the training process, was same as the one used for the 5-map CNN, but with one major difference. The input was a single 128x128 average subject map instead of a set of 5 128x128 maps per subject.

### 2.5.3 Logistic Regression over component strengths

. It is possible that the more complex classifiers described above are successful because the generated maps differ already in some very simple aspects from the real measured maps. To test for this possibility, we computed a set of marginal statistics from the real and generated maps. These statistics were the variances that could be ascribed to different components (see generative process). We then submitted these marginal statistics, rather than the maps, to a logistic regression classifier (Hosmer, Stanley, and Sturdivant, 2013). These components were - the overall mean component, the subject specific component, the finger specific component  $x_{i,j}$ , that contains the pattern specific to a particular finger of a subject; and the noise  $z_{i,j,k}$ . Their respective strengths were denoted as  $\Omega(\bullet)$ .

The overall mean component,  $x_{\dots}$ , reflected the general activity pattern shared across all subjects in the dataset. Its strength was computed as the slope  $\zeta_i$  of the regression from the overall mean map to the subject mean. Equation 2.3 specifies the method used to compute the mean component strength.

The subject specific component reflected the general activity pattern shared by all fingers of a particular subject. The subject component strength was computed from the residuals of the regression against the mean map (Eq. 2.3) by squaring and averaging the residuals across the map (Eq. 2.4).

The finger specific component  $x_{i,j}$ , reflected the pattern specific to a particular finger of a subject. Its strength quantified the strength of the idiosyncratic finger-specific patterns around the subject mean (Eq. 2.5).

Finally, the variance of the noise component was calculated for a specific subjects from the variation of the measurements of each run ( $x_{i,j,k}$  around the mean map for this sample and finger  $x_{i,j}$ ):

$$\Omega(z_i) = \frac{\sum_{j=1}^5 \sum_{k=1}^r \sum_{a,b=1}^{128} [(x_{i,j,k}[a,b] - x_{i,j}[a,b])^2]}{16384 \times r \times j} \quad (2.8)$$

where  $r$  = Number of runs for the subject. We calculated the strengths of these components from the real and the generated samples and compared them using logistic regression. A logistic regression model can tell us whether there a linear separation between real and generated maps was possible in terms of the component strength. We measured the accuracy across 100 iterations using K-fold cross validation. The process for generating the folds was the same as the one used in the 5-map CNN.

#### 2.5.4 Neural Network over component strengths

The previous model determined whether the generated maps can be linearly separated from the real ones based on their component strength. However, the 5-map CNN may learn to separate them based on a non-linear combination of component strengths. The use of a fully connected neural network architecture therefore opens the pathway for exploring non-linear relationships (LeCun, Bengio, and Hinton, 2015) in the component strengths that can help distinguish between the real maps and the generated maps. The architecture used in this case comprised of a network with one hidden layer (five units) and a single unit output with sigmoid activation. Binary cross entropy (BCE) loss was used as the loss function. We recorded the accuracy over 100 iterations. The cross validation process used the same process for generating the folds as the one used for the 5-map CNN. The model was trained for 200 epochs before being evaluated on the left-out fold. we carried out 100 iterations of the training and evaluation process.

#### 2.5.5 Logistic Regression over the finger covariance matrix

An essential characteristic of finger maps is the presence of reliable covariance structure ( $G$ ) between the different finger maps across different individuals (Ejaz, Hamada, and Diedrichsen, 2015). A similar  $G$  matrix measured from

the real and the generated data indicates the correct distribution of inter-finger relationships. The similarity between this matrix observed from the real data,  $G$ , and the generated data,  $\hat{G}$  was quantified using a logistic regression model. The  $G$  matrix was estimated as:

$$G_{i,j} = cov[x_{i,j}, x_{i,j}]. \quad (2.9)$$

The 25 elements of the matrix served as inputs to the logistic regression model. A high accuracy score of this model would signify that the generative model failed to match the inter-finger covariance correctly. The model was tested over 100 iterations using K-fold cross validation with 5 folds.

### **2.5.6 Neural Network over the finger covariance matrix**

The objective of using a Neural Network for the  $G$  matrices is to analyze more complex, non-linear relationships for differentiating between the finger covariance matrices from the real data and the generated data. The network architecture and the loss function were the same as the one used above for the component strength network. The accuracy was again recorded over 100 iteration with K-fold cross validation. The cross-validation process was again the same as the 5-map CNN.

# Chapter 3

## Results

### 3.1 General characteristics of the generated maps are well matched to real data

For each real data set, each consisting of five digit maps, we used the baseline and usage model to generate a number of generated samples (see Figure 3.1). Visual inspection suggested that the generative models were successful in matching the overall characteristics of the real maps including the level of spatial smoothness, inter-individual variability, and general spatial distribution of the patterns.

To evaluate the match between real and generated maps in terms of the overall statistical characteristics quantitatively, we first assessed whether the strength of each component, the mean activity map, deviation of individual subjects from that mean, idiosyncratic finger-specific patterns around the subject mean, and the independent noise for each map, was similar across real and generated maps. For verifying a successful match of these components, we estimated the variance associated with each of these components (section 2.5.3), and then used a logistic regression model to detect replicable differences between real and generated maps. For the maps from the usage model, this resulted in an accuracy of 50.87  $\pm$  0.06%, barely different from a chance performance score of 50%. Additionally, student's t-tests performed on each of the extracted component strengths (from 50 generated samples using the usage model) with the corresponding component strengths from the real data (having 50 real samples) showed non-statistically significant difference in distribution means (Table 3.1). Hence, the component strengths were

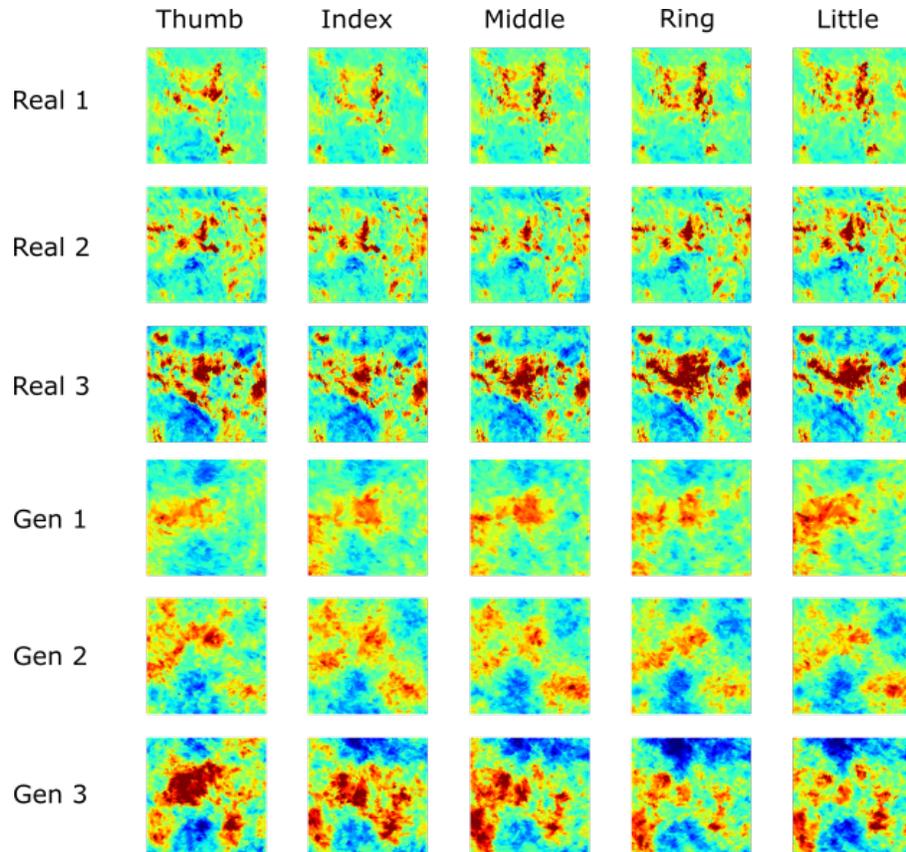


Figure 3.1: **Real and generated digit maps.** Each row shows a sample consisting of activity maps for the 5 fingers. Top 3 rows are real samples, and the bottom three rows show three generated samples from the usage model, generated using the measurement model of the 3 real samples.

Table 3.1: **Test statistics for t-tests over component strengths.** The table shows the results from the t-test carried out on the component strengths calculated from the measured maps from the usage model vs the corresponding component strength from the real measured maps.

| <b>Component Matched (Strength)</b> | <b>t-value</b> | <b>p-value</b> |
|-------------------------------------|----------------|----------------|
| Mean activity map (eqn. 2.3)        | 0.2159         | 0.8294         |
| Subject map (eqn. 2.4)              | 0.6871         | 0.4935         |
| Finger map (eqn. 2.5)               | -1.844         | 0.0680         |
| Noise map (eqn. 2.8)                | 0.0861         | 0.9315         |

successfully matched and were not linearly separable.

Despite the lack of linear separability of the component strengths, there was still a possibility of non-linear separability in the data, such as the dispersion of co-variation of certain component strengths. Therefore, we tested whether the simulated and real patterns were non-linear separable using a neural network with one hidden layer (see Methods). This analysis yielded a accuracy score of 56.4 +- 0.4% for the usage model (see Figure 3.2). The accuracy, while slightly higher than the one observed from logistic regression, was still relatively close to chance accuracy score and under the failure threshold. Thus, while some aspect of the generated maps could be separated non-linearly in the space of simple marginal statistics from the real maps, the deviations were modest.

The central characteristics of the usage model is the covariance of the maps across fingers. This 5x5 covariance matrix ( $\mathbf{G}$ ), is a stable characteristic of finger real maps (Ejaz, Hamada, and Diedrichsen, 2015). For testing the match of these covariance matrices we estimated them from the generated measured maps and the real maps, vectorized them, and compared them using a logistic regression model. An accuracy score of 51.2 +- 0.1% signified the absence of linear separability of the two (real and generated measured maps)  $\mathbf{G}$  matrix distributions. In contrast, the simulated maps from the baseline model, which generated the finger maps to be independently distributed, could be distinguished with an accuracy of 66.3 +- 0.1% from the real maps (Figure 3.2).

A test for non-linear separability of the  $G$  matrix (finger covariance matrix) distributions, similar to the test above with the component strengths, was important at this stage to rule out the presence of any deviant non-linear

relation between the components of the  $G$  matrix. We computed the  $G$  matrices from the generated measured maps and the real maps. Quite evidently, the evaluation framework was able to detect the presence of anomalous finger covariance structures in the generated maps from the baseline model. The logistic regression model classified the  $G$  matrices from the baseline model with a  $66.3 + -0.1\%$  accuracy while same classification with the neural network had a  $89.8 + -0.3\%$  accuracy. At the same time, the real and generated  $G$  matrices from the usage model, when tested with a neural network same as the one used before produced an accuracy of  $51.8 + -0.2\%$  which again, was close to a chance score. Thus, the usage model sufficiently matched the  $G$  matrices. The matching of the  $G$  matrices, alongside the matching of the component strengths mentioned above signalled that the general statistical characteristics between the simulated maps from the usage model were relatively well matched to the real data.

### 3.2 General statistical characteristics are not sufficient to generate artificial maps

Given that the general statistics have been successfully reproduced in the generated maps, we now wanted to test whether there are any important organizing principles in the real measured maps that are not reflected in our generative model. To detect these (so far unknown characteristics) we trained a convolutional neural network model to distinguish between the two datasets (real measured maps and generated measured maps).

The convolutional network achieved a cross-validated accuracy of  $93.4 + -0.8\%$  for the usage model ( $84.6 + -0.6\%$  with histogram equalization), and  $96.5 + -0.3\%$  for the baseline model (Figure 3.2). Clearly, therefore, the convolutional neural network was able to detect some important characteristics of the real maps outside of the general marginal statistics that distinguished them from the generated ones. The high accuracy signified the failure of both - the baseline and the usage generative model. Thus, there is more to the organization in the maps than spatial smoothness and covariance structure. Interestingly, the 5-map CNN had a lower accuracy score when classifying the generated maps from the usage model as compared to the accuracy score when classifying the maps from the baseline model. The general drop in the map classification accuracy when using the usage model pointed to two

areas of success - one, natural statistics of finger movement were verified to be an important representational statistic of finger maps; and two, our evaluation framework was robust enough to detect a deviation from the required representational statistics of the finger maps.

This suggested that the 5-map CNN was sensitive to the changes in the maps arising out of the presence of a finger covariance structure based on the natural statistics of finger movement.

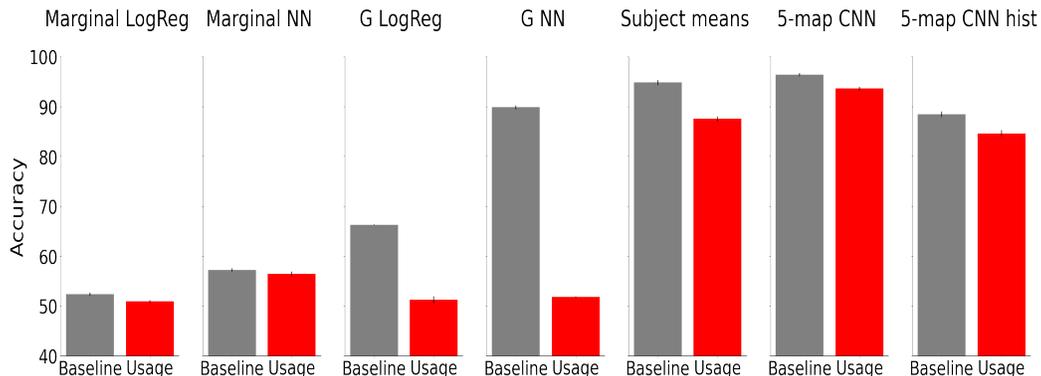


Figure 3.2: **Classification accuracies (%) of the various classification models for the baseline (gray) and usage (red) model.** The classifier from left to right were component strengths with logistic regression, component strengths with neural network, G matrix with logistic regression, G matrix with neural network, subject means, 5-map CNN, 5-map CNN with histogram equalization.

### 3.3 Failure of generative models is already visible in the mean maps

We then wanted to gain some insight into the reasons for the failure of the current generative model. As a first step, we sought to establish whether the successful discrimination was made based on some feature of mean map, or based on the relationship between the maps for the different fingers. To do this, we submitted either all 5 finger maps to the discriminative model (as before) or only the average map across all 5 fingers. High classification accuracy for the latter case would indicate that the main failure of the generative

model lies already in the generation of the mean-sample map.

The convolutional network achieved a cross-validated accuracy of 88.5 +- 0.5% for distinguishing the sample means of real vs the generated data from the usage model (and 94.8 +- 0.5% for the baseline model). The high accuracy score suggested the failure of the statistical model in replicating the necessary organizational properties of the mean sample maps.

### 3.4 Non-Gaussianity of real activity data is one contributing factor

While we attempted to carefully match the mean and variance (both spatially and across components) of the real maps, we did not explicitly match the distribution of the activity values. For simplicity, our generative models uses the normal distribution for the required components. To inspect this validity of this simplification, we inspected a quantile-quantile plot of the real activity maps (see figure 3.3) vs a normal distribution suggested a deviation from the normality assumption. The real maps were found to have a higher kurtosis than the maps from our models.

To determine to what degree the 5-map CNN model learned to distinguish real from the generated maps using this characteristic, we apply histogram normalization to each generated map to exactly match the marginal distribution of the set of corresponding finger maps.

The classification with the 5-map CNN clearly marked a drop in classification accuracy with both generative models. The baseline model was now classified against the real maps with a 88.33 +- 0.9% accuracy while the maps from the usage model were classified with a 84.6 +- 0.6% accuracy.

The drop in the classification accuracies indicates that the normality assumption in the generative model was not justified and contributed to the ability of the classifier to distinguish real from generated maps. As a solution, the activity values should be drawn from a more heavy-tailed distribution. Nonetheless, the analysis also indicated that the network was able to detect some other characteristics of the real maps that were not reflected in the current generative model.

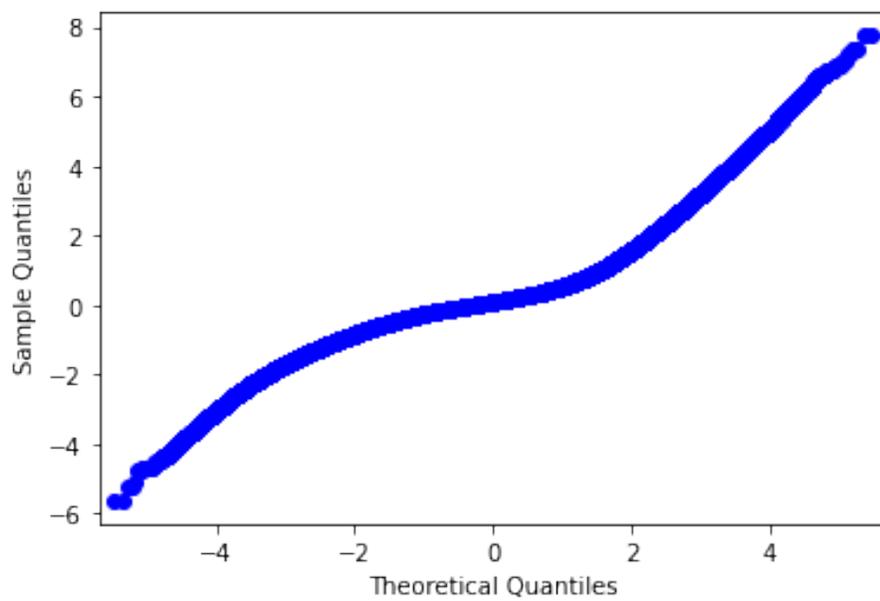


Figure 3.3: **Real maps are slightly non-Gaussian.** The quantile-quantile plot of the activation values of the real maps vs the theoretical values from a normal distribution indicates that the values come from a more heavy-tailed distribution.

# Chapter 4

## Discussion

In our research, we developed a novel generative-discriminative approach to tackle the challenge of human brain mapping. We developed a comprehensive standalone evaluation framework for the evaluation of brain activity maps generated by a given generative model and formulated a development strategy driven by the insights from this framework. We further used this framework to test the hypothesis of the sufficiency of natural statistics of finger movement in explaining the organization of finger activity maps in the M1 region of the brain. This was done using a well-curated data set comprising of various fMRI recordings of different subjects across a range of scanners and experiments.

For studying the sufficiency of natural statistics in explaining the organization of activity maps in the sensorimotor region, we developed a generative model that successfully replicated the known marginal statistics in its output samples (generated activity maps). Additionally, we developed a plug-in measurement simulation approach that simulates the effect of the scanning and mapping process, present in the real samples, on the generated activity maps from any generative model. Our generative model, followed by the measurement simulation, successfully generated sets of measured finger maps (generated samples) that were visually indistinguishable from real measured finger maps. However, we discovered that despite being visually indistinguishable from, and having similar component strengths and spatial covariance structures as the real samples, the samples generated by our model (based on known natural statistics) were easily distinguishable from the real samples. This led us to conclude that the finger maps are not strictly organized. There is a dynamic process that shapes these representations in the

motor cortex that is shaped by our everyday activities. For instance, since the thumb moves alone, its activity maps end up having different properties/location as compared to the other activity maps.

The results from the use of our evaluation framework on the generated samples quantified the successful matching of various component strengths, and the covariance structure of finger maps (the G matrix), in the generated samples. The evaluation framework was unable to differentiate between the real and the generated samples based solely on the strengths of components involved, when evaluated using a logistic regression model, a neural network, and a statistical hypothesis test for deviations in means. Similarly, the model failed to differentiate the finger covariance matrices (G matrices) estimated from real and the generated samples, when evaluating with logistic regression and a neural network.

Yet the evaluation framework was able to successfully underscore the deviation in the generated samples as compared to the real samples when the complete samples (5 maps each per sample from real and generated datasets) were classified using a deep convolutional neural network. The 5-map CNN model was designed to successfully model the deviations in the organizational properties in the real maps vs the generated maps. Its architecture, which allowed it to uniquely process the 5 finger maps (that is, one sample) together, ensured that the features from each map, when pooled together, provided sufficient information to consequently classify the two types of samples (real and generated). This led us to conclude that a mere reproduction of known natural statistics in the generated samples did not make them indistinguishable from the real samples. This meant that there existed other properties which caused deviations in the generated samples from the distribution of real samples. These properties remained evasive under the scope of our current work.

The current thesis reports on the very first attempt to apply this novel framework to functional activity maps.

Additionally, the measurement simulation process developed for transforming the generated activity maps into generated measured maps is an estimation of the real process and introduces some smoothness during mapping to and from the voxel space into the surface space. The evaluation of statistical properties and the subsequent development of new activity maps was done based on measured real maps which had been resampled into the surface space. The measurement and the resampling process ensured that we essentially worked with estimates of real activity maps, and not the real ac-

tivity maps themselves. A possible approach to attenuate these problems can be a high degree of smoothing of the measured activity maps. We can rely on the smoothed activity maps however, this will lead to a loss of information about the fine-grained organization of these maps.

While we used a generative-discriminative framework, we did not employ a full adversarial learning architecture, as has been successful in other areas of machine learning (Goodfellow et al., 2014) (Yi, Walia, and Babyn, 2019) (Creswell et al., 2018). In this work, the generative model is itself a neural network that is trained by the accuracy of the adversarial model. In our case, the usefulness of this approach is severely limited by the available data (50 unique samples (subjects/hemispheres)). and compounded by the high dimensionality (16384 features per finger map). Initial attempts to use a learned generative model showed that the training resulted in the generative model learning specific instances of real maps, rather than extracting structural commonalities across maps. In future work, we can incorporate data from larger available datasets, such as the task-based data from the human connectome project (Barch et al., 2013), for use alongside the individual activity maps.

However, the development of the evaluation framework and the generative-discriminative methodology provides the first steps in a promising avenue of future research. Clearly, the high dimensionality of human brain imaging data sets makes automated feature detection for clinical classification and individual prediction very challenging, even with the effort to build large data with thousands of samples. The use of structured and powerful generative models promises to help with detecting the important principles that reflect the critical features of these maps (Kriegeskorte and Douglas, 2018), and separate these from unimportant biological variation. By using the individual parameters from such generative models, rather than the raw data, the dimensionality of the learning problem could be dramatically improved (Stephan et al., 2017). The evaluation capabilities of our framework not only make it a general, adaptable process for the evaluation of other generative models but also provide the researchers abundant insights into the sources and patterns of failures in the generative model under evaluation.

Thus, despite clear shortcomings and many possible further extensions to pursue, we hope our work can serve the broader research community in multiple ways – as the proof of insufficiency of natural statistics of finger movement as the sole organizational factor in the activity maps in the M1 region; as the proponent of a new generative-discriminative approach towards

brain mapping; as a descriptor of a modelling and measurement simulation process for generating artificial brain maps; and as a source of curated data for further research and development.

# Symbol Table

|  |                       |   |
|--|-----------------------|---|
| $\mathbf{x}_i$                                       |                       | Sample, set of maps for a specific subject                    |
| $\mathbf{x}_{i,j,k}$                                 | $1 \times p$          | Individual measured map of subject $i$ , finger $j$ , run $k$ |
| $\mathbf{x}_{\dots}$ or $\bar{\mathbf{x}}$           | $1 \times p$          | Overall mean  |
| $\mathbf{x}_{i,\dots}$ or $\bar{\mathbf{x}}_i$       | $1 \times p$          | Mean subject map for subj $i$                                 |
| $\mathbf{x}_{i,j,\dots}$ or $\bar{\mathbf{x}}_{i,j}$ | $1 \times p$          | Mean for finger $j$ and sample $i$ across runs.               |
| $\chi$   | $n \times 5 \times p$ | Set of biological true maps                                   |
| $\mathbf{V}_i$                                       | $p \times v$          | Pixel to Voxel transformation matrix for subject $i$          |
| $\mathbf{P}_i$                                       | $v \times p$          | Voxel to Pixel transformation matrix for subject $i$          |
| $\Sigma_{i,\epsilon}$                                | $v \times v$          | Noise covariance matrix for subject $i$                       |
| $\Sigma_{i,p}$                                       | $p \times p$          | Subject covariance matrix for subject $i$                     |
| $\mathbf{D}$   | $p \times p$          | Distance Matrix   |
| $\mathbf{G}$   | $5 \times 5$          | Finger covariance matrix (from real data)                     |
| $\mu$  | 1                     | Subject specific component scaling factor                     |
| $\nu$  | 1                     | Finger specific component scaling factor                      |
| $\omega_i$   | 1                     | Signal scaling factor for subject $i$                         |
| $\eta_i$   | 1                     | Noise scaling factor for subject $i$                          |
| $\mathbf{c}_{i,j}$                                   | $5 \times p$          | Simulated true maps for subject $i$ , finger $j$              |
| $\mathbf{s}_i$                                       | $1 \times p$          | Subject component for subject $i$                             |
| $\mathbf{f}_{i,j}$                                   | $1 \times p$          | Finger component for subject $i$ , finger $j$                 |
| $\mathbf{z}_{i,j,k}$                                 | $1 \times v$          | Noise component for subject $i$ , finger $j$ , run $k$        |
| $\mathbf{y}_{i,j,k}$                                 | $1 \times p$          | Generated measured map for subject $i$ , finger $j$ , run $k$ |
| $\Omega(\bullet)$                                    | 1                     | Strength of a given component                                 |
| $\mathbf{H}(\bullet)$                                | n/a                   | Classifier function   |

where,

$n$  = Number of subjects,

$p$  = Number of pixels,

$v$  = Number of voxels for the current subject (variable)

# Bibliography

- Penfield, W. and Rasmussen, T. (1950). “The cerebral cortex of man; a clinical study of localization of function.” In: *Macmillan*.
- Engel, S., Glover, G. H., and Wandell, B. (1997). “Retinotopic organization in human visual cortex and the spatial precision of functional MRI”. In: *Cerebral cortex (New York, N.Y. : 1991)* 7, pp. 181–192. DOI: <https://doi.org/10.1093/cercor/7.2.181>.
- Schieber, M. H. (2001). “Constraints on somatotopic organization in the primary motor cortex”. In: *Journal of neurophysiology* 86 (5), pp. 2125–2143. ISSN: 0022-3077. DOI: <https://doi.org/10.1152/jn.2001.86.5.2125>.
- Liepert, J. et al. (2005). “Motor Strokes: The Lesion Location Determines Motor Excitability Changes”. In: *Stroke* 36, pp. 2648–2653.
- Graziano, M.S.A. and Aflalo, T.N. (2007). “Mapping Behavioral Repertoire onto the Cortex”. In: *Neuron* 56.2, pp. 239–251. ISSN: 0896-6273. DOI: <https://doi.org/10.1016/j.neuron.2007.09.013>.
- Meier, J.D. et al. (2008). “Complex organization of human primary motor cortex: a high-resolution fMRI study”. In: *J Neurophysiol* 100.4, pp. 1800–1812. ISSN: 0022-3077 (Print) 0022-3077 (Linking). DOI: <https://doi.org/10.1152/jn.90531.2008>.
- Glover, G. H. (2011). “Overview of functional magnetic resonance imaging”. In: *Neurosurgery clinics of North America* 22 (2), pp. 133–vii. ISSN: 1558-1349. DOI: <https://doi.org/10.1016/j.nec.2010.11.001>.
- Fischl, B. (2012). “FreeSurfer”. In: *NeuroImage* 62.2. 20 Years of fMRI, pp. 774–781. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2012.01.021>.
- Barch, D.M. et al. (2013). “Function in the human connectome: task-fMRI and individual differences in behavior”. In: *Neuroimage* 80, pp. 169–189.

- ISSN: 1095-9572 (Electronic) 1053-8119 (Linking). DOI: <https://doi.org/10.1016/j.neuroimage.2013.05.033>.
- Diedrichsen, J., Wiestler, T., and Krakauer, J. W. (2013). “Two distinct ipsilateral cortical representations for individuated finger movements”. In: *Cerebral Cortex* 23 (6), pp. 1362–1377. ISSN: 1460-2199.
- Hosmer, D. W., Stanley, L., and Sturdivant, X.R. (2013). *Applied logistic regression*. Vol. 398. John Wiley & Sons.
- Goodfellow, I. J. et al. (2014). *Generative Adversarial Networks*. arXiv: 1406.2661 [stat.ML].
- Kingma, D. P. and Ba, J. (2014). “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980*.
- Ejaz, N., Hamada, M., and Diedrichsen, J. (2015). “Hand use predicts the structure of representations in sensorimotor cortex”. In: *Nature Neuroscience* 18 (7), pp. 1034–1040. ISSN: 1546-1726. DOI: <https://doi.org/10.1038/nn.4038>.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). “Deep learning”. In: *nature* 521.7553, pp. 436–444.
- Radford, A., Metz, L., and Chintala, S. (2015). “Unsupervised representation learning with deep convolutional generative adversarial networks”. In: *arXiv preprint arXiv:1511.06434*.
- Xu, B. et al. (2015). “Empirical Evaluation of Rectified Activations in Convolutional Network”. In: *CoRR* abs/1505.00853. arXiv: 1505.00853.
- Yadav, S. and Shukla, S. (2016). “Analysis of k-Fold Cross-Validation over Hold-Out Validation on Colossal Datasets for Quality Classification”. In: *2016 IEEE 6th International Conference on Advanced Computing (IACC)*, pp. 78–83. DOI: <https://doi.org/10.1109/iacc.2016.25>.
- Rawat, W. and Wang, Z. (2017). “Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review”. In: *Neural Computation* 29.9, pp. 2352–2449. DOI: [https://doi.org/10.1162/neco\\_a\\_00990](https://doi.org/10.1162/neco_a_00990).
- Stephan, K.E. et al. (2017). “Computational neuroimaging strategies for single patient predictions”. In: *NeuroImage* 145. Individual Subject Prediction, pp. 180–199. ISSN: 1053-8119. DOI: <https://doi.org/10.1016/j.neuroimage.2016.06.038>.
- Creswell, A. et al. (2018). “Generative Adversarial Networks: An Overview”. In: *IEEE Signal Processing Magazine* 35.1, pp. 53–65. DOI: <https://doi.org/10.1109/msp.2017.2765202>.
- Kriegeskorte, N. and Douglas, P.K. (2018). “Cognitive computational neuroscience”. In: *Nature neuroscience* 21.9, pp. 1148–1160.

- Arbuckle, S. A., Yokoi, A., et al. (2019). “Stability of representational geometry across a wide range of fMRI activity levels”. In: *Neuroimage* 186, pp. 155–163. ISSN: 1053-8119.
- Berlot, E. et al. (2019). “Ipsilateral finger representations in the sensorimotor cortex are driven by active movement processes, not passive sensory input”. In: *Journal of Neurophysiology* 121 (2). PMID: 30517048, pp. 418–426. DOI: <https://doi.org/10.1152/jn.00439.2018>.
- Gershman, S.J. (2019). “The Generative Adversarial Brain”. In: *Frontiers in Artificial Intelligence* 2, p. 18. ISSN: 2624-8212. DOI: <https://doi.org/10.3389/frai.2019.00018>.
- Yi, X., Walia, E., and Babyn, P. (2019). “Generative adversarial network in medical imaging: A review”. In: *Medical Image Analysis* 58, p. 101552. ISSN: 1361-8415. DOI: <https://doi.org/10.1016/j.media.2019.101552>.
- Arbuckle, S. A., Weiler, J., et al. (2020). “Structure of Population Activity in Primary Motor Cortex for Single Finger Flexion and Extension”. In: *Journal of Neuroscience* 40 (48), pp. 9210–9223. ISSN: 0270-6474. DOI: <https://doi.org/10.1523/jneurosci.0999-20.2020>.

# Curriculum Vitae

**Name:** Deepanshu Wadhwa

**Post-Secondary Education and Degrees:** Guru Gobind Singh Indraprastha University  
New Delhi, India  
2015 - 2019  
Bachelor of Technology in Information Technology

University of Western Ontario  
London, ON  
2019 - 2021  
MSc Computer Science (candidate)

**Honours and Awards:** Western Graduate Research Scholarship  
2020

**Related Work Experience:** Graduate Research Assistant and Teaching Assistant  
The University of Western Ontario  
2019 - 2021