



The
University
Of
Sheffield.

Understanding factors driving the origins and plasticity of somatosensory representations.

Laura Rose Edmondson

A thesis submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

The University of Sheffield
Faculty of Science
Department of Psychology

March, 2022

Acknowledgements

First, I would like to give my greatest gratitude to my supervisor, Dr Hannes Saal, for his continuous support, encouragement and enthusiasm throughout my PhD.

I would like to thank our collaborators at the University of Oxford and UCL— Dr Daan Wesselink and Prof. Tamar Makin— for their assistance with the final chapter on plasticity, providing the fMRI data and being receptive to our ideas and opinions. Furthermore, I would like to thank Dr Alejandro Jimenez-Rodriguez for his assistance with the derivations in chapter 3.

I would like to thank Active Touch Lab members for their support and for providing helpful feedback for conferences and presentations. Furthermore, I thank my friends and colleagues of the F01 PhD office past and present- Anil, Blayze, Dan, Eduardo, Felipe, Giulia, John, Manal, Matt, Yue & Zeke. You kept me going with all the coffee, snacks and chat!

I also wish to thank my parents and sister for their unconditional support. Finally, I would like to thank Adriel for always motivating me, his patience and understanding, especially during the write-up.

Abstract

Somatosensory cortex contains a topographic representation of the body surface, whereby neighbouring regions on the body are represented by adjacent cortical regions. Furthermore, some regions are magnified, such that their representations are larger than would be expected by their physical size. What determines the setup and plasticity of these representations? We first investigate the possible contributions of the peripheral afferent densities and the statistics of tactile input, of which both tend to be greater in over-magnified regions. We consider that the brain has limited capacity resources—bottlenecks, which constrains representation of input regions. Building on previous work in efficient coding, we use linear second-order models to maximise information. We show that the optimal representation depends crucially on the width of the bottleneck; however, exact patterns of region over and under-representation differs depending on the combination of receptor density and activation. We test the model’s predictions using published empirical measurements of both factors and resulting cortical sizes for a highly touch-specialised organism, the star-nosed mole. We demonstrate the importance of usage statistics in determining allocations in this case. Second, we investigate a popular unsupervised self-organising model, which can produce topographic maps of inputs. We find that this model largely disregards the density of receptors and instead follows the input statistics. Furthermore, this model was unable to reproduce the inverse receptive field relationship, a key feature of cortical sensory maps, where magnified regions tend to have smaller receptive fields. Finally, we consider other contributors to somatosensory representation by modelling empirically measured effects of a short-term plasticity protocol. We found the canonical hand representation could be reproduced by implementing more divergent feedforward connectivity and explicit lateral connections. Furthermore, we suggest a role for homeostatic control of cortical activation, which was able to reproduce the global activation decrease across the hand map under digit anaesthetic. Overall, we demonstrate that somatosensory representations are dependent and differentially affected by the input statistics and density of receptors. Other factors such as network connectivity and resource capacity are also crucial for forming and maintaining representations.

Contents

1	Introduction	1
1.1	Aim and motivations	1
1.2	Outline	2
1.3	Declaration of work contributed by others	4
2	Introduction: Sense of touch	5
2.1	Somatosensory periphery	5
2.1.1	Tactile afferents of the skin	5
2.2	Hand use and statistics of tactile interactions	8
2.3	Tactile information ascension to cortex	10
2.3.1	Dorsal column medial lemniscus pathway	10
2.3.2	Thalamus	11
2.4	Somatosensory cortex	11
2.4.1	Initial somatotopic map formation	13
2.4.2	Properties of subregions of S1	15
2.4.3	Somatosensory cortical layers	17
2.4.4	Investigating hand and digit representations in S1	17
2.5	Tactile perception	21
2.6	Modelling the somatosensory system	22
3	Nonlinear scaling of resource allocation in sensory bottlenecks	24
3.1	Introduction	24
3.1.1	Theories of efficient coding	24
3.1.2	Noise and efficient coding	26
3.1.3	Approximations with Fisher information	27
3.1.4	Bottlenecks	28
3.1.5	Aims of this chapter	29
3.2	Methods	30
3.2.1	Problem overview	30
3.2.2	Combined whitening and dimensionality reduction	31

3.2.3	Whitening with multiple input regions	33
3.2.4	Analytic calculation of eigenvalues	35
3.2.5	Allocation in the bottleneck	36
3.3	Results	39
3.3.1	Heterogeneous density	39
3.3.2	Heterogeneous statistics	40
3.3.3	Interplay between stimulus statistics and receptor density	42
3.3.4	Resource limits determine the extent of changes under plastic manipulation	45
3.4	Discussion	47
3.4.1	Comparison with previous approaches	47
3.4.2	Bottlenecks in tactile processing	48
3.4.3	Implications for sensory processing	50
3.4.4	Limitations and future work	53
4	Efficient coding in the Star-Nosed Mole	55
4.1	Introduction	55
4.1.1	The star-nosed mole	55
4.1.2	Anatomy of the star	56
4.1.3	Behavioural use of the star	57
4.1.4	Somatosensory representations of the star	57
4.1.5	Aim of this chapter	58
4.2	Methods	58
4.3	Results	61
4.4	Discussion	64
5	Self-organising models of somatosensory representations	68
5.1	Introduction	68
5.1.1	Somatosensory models	68
5.1.2	Aims of this chapter	73
5.2	Methods	73
5.2.1	Input sets	73
5.2.2	Kohonen SOM mapping algorithm and parameters	74
5.2.3	Selected SOM parameters	77
5.2.4	Simulations	78
5.3	Results	79
5.3.1	Effects of varying receptor densities in the Kohonen SOM	79
5.3.2	Problems of implementing density via changing the stimulus statistics	80
5.3.3	Manipulating the statistics ratio and receptive field sizes	82
5.4	Discussion	87

6	Plasticity of the hand representation	91
6.1	Introduction	91
6.1.1	Plasticity in the somatosensory system	92
6.1.2	Mechanisms of somatosensory plasticity	95
6.1.3	Roles of the types of plasticity	96
6.1.4	Aims of this chapter	97
6.2	Methods	98
6.2.1	Empirical fMRI data: baseline and post-lidocaine	98
6.2.2	Computational modelling methods	99
6.2.3	Comparing before and after block: Representational similarity analysis	105
6.3	Results	107
6.3.1	Baseline fMRI results	107
6.3.2	Toy model results	107
6.3.3	fMRI block results	111
6.3.4	Full model: Somatosensory responses to lidocaine block	111
6.4	Discussion	115
6.4.1	All-to-all map connectivity is required for activation of all clusters	115
6.4.2	Global and local changes are required to reproduce the effects of digit-block	116
6.4.3	Limitations and future work	118
7	General discussion	120
7.1	Overview	120
7.2	Considerations for future work	122
7.3	Implications for robotics and sensing applications	123
7.4	Considerations for neuroscience: clinical applications and body augmentation	123
	Appendices	153
A	Chapter 3: Non-linear scaling of resource allocation.	154
A.1	Differences between 1D and 2D regions	154
A.2	Expression as the total amount of information	157
A.3	Full derivation of problem	158
A.4	Ordering in the 2D square case	162
A.5	Allocation for multiple regions in 2D	163
A.6	Allocation for the 1D case	164
B	Chapter 4: Star-Nosed Mole	165
B.1	Star-nosed mole prey model	165
B.2	Comparison of bottleneck best fits	167
B.3	Choice of error measure for calculating the bottleneck	167

Chapter 1

Introduction

1.1 Aim and motivations

A prominent feature of primary sensory cortices is representational maps. In the somatosensory cortex, an entire topographic representation of the body exists, such that adjacent cortical regions represent nearby locations on the body surface. Furthermore, there is also magnification, where some body regions, such as the hands and lips, have a much larger representation in the map than a proportional allocation to body part size. Two possible reasons for this have been proposed. Firstly, this could simply reflect the densities of the afferents across the body— afferent numbers are much greater in the hands than other regions of the body. Second, it could be due to the typical usage of body parts. For example, the hand is essential in everyday interactions and manipulations of objects. Most previous work has focused on visual sensing, investigating only the statistics of incoming stimuli via the study of natural images. Therefore, a key contribution of this work is the consideration of receptor density, which provides an upper limit on the amount of incoming information. This consideration is highly relevant to touch processing, given previous suggestions that this could be driving cortical magnification.

Furthermore, the scale of magnification could be determined by the amount of resources a system has. Given this, we also investigate how limited capacity bottlenecks may constrain representations. We explore this using ideas from efficient coding theory in chapters three and four. In chapter five, we investigate the effects of manipulating the receptor density and stimulus statistics using the Kohonen SOM algorithm, which has been used as a model of somatosensory development and plasticity, based on self-organisation and Hebbian learning. In the final chapter, we investigate other contributors to the somatosensory map, including the architecture of the network and homeostatic plasticity, which emerge as important factors to reproduce the effects of tactile input manipulation.

The full contributions from each chapter are detailed below, alongside the resulting publications.

1.2 Outline

Chapter 2: Background

Here we provide a broad overview of the thesis context, detailing the somatosensory system from the periphery to cortex. First, we describe the set-up of peripheral sensing, focusing on the hand, followed by a description of the pathway between the hand and somatosensory cortex. We then discuss tactile information processing in somatosensory cortex and the development and features of the body map, such as topography and cortical magnification. We draw on research from both humans and non-human primates.

Chapter 3: Nonlinear scaling of resource allocation in sensory bottlenecks

Chapter 3 investigates the contribution of input receptor densities and stimulation statistics to the representation under limited capacity restrictions, utilising ideas from efficient coding. For both factors, we demonstrate that the width of the bottleneck, or amount of retained information, is crucial for determining the relative allocation of resources to regions.

Contributions:

- Application of ideas from efficient coding theory to tactile sensing, particularly redundancy reduction via decorrelation methods. We develop a normative linear model to understand optimal representations of regions under a sensory bottleneck.
- We provide an analytic solution to the problem, assuming exponential covariance matrices for multiple 1D or 2D regions varying in their receptor densities and activation.
- We demonstrate that the density of receptors and statistics leads to different optimal allocations of bottleneck resources between regions.
- We find that for this model the bottleneck width is crucial in determining whether region representations are expanded or contracted relative to their density.

Related publications:

- **Edmondson, L. R.**, Rodriguez, A. J., & Saal, H. P. (2019). Nonlinear scaling of resource allocation in sensory bottlenecks. *In Advances in Neural Information Processing Systems* (pp. 7545-7554).
- **Edmondson, L. R.**, Rodriguez, A. J., & Saal, H. P. (2021). Expansion and contraction of resource allocation in sensory bottlenecks. *BioRxiv*.

Chapter 4: Efficient coding in the Star-Nosed Mole

This chapter applies the efficient coding model of chapter 3 to predict cortical somatosensory allocations of the sensing ‘ray’ organs in the star-nosed mole. The contact statistics, receptor densities and cortical allocations have been measured empirically for the mole, allowing comparisons

between the model and measured allocations.

Contributions:

- Application of the model developed in chapter 3 to a model organism, the star-nose mole. We show that an efficient coding scheme could explain cortical somatosensory representations.
- We demonstrate that the model can produce a good fit to the empirical cortical data. We find that the ray usage statistics are important in determining the allocations, in line with previous work.

Related publications:

- **Edmondson, L. R.**, Rodriguez, A. J., & Saal, H. P. (2021). Expansion and contraction of resource allocation in sensory bottlenecks. *BioRxiv*.

Chapter 5: Self-organising models of somatosensory representations

This chapter tests a previous model of somatosensory representations and plasticity-the Kohonen SOM. This model produces maps of model input regions with topographic representations of the space. Previous studies have mainly used this model to investigate plasticity. Here we test how the modelled maps are affected by varying density and statistics, which have not previously been systematically investigated.

Contributions:

- Calculate map allocations when changing the input densities in the Kohonen SOM. We find little effect of the input region density.
- We test a previous assumption that the change in density can be modelled as a change in the statistics, and demonstrate how this deviates from explicitly manipulating the receptor densities.
- The statistics of input closely match the allocation of map area to each region. However, we show that the inverse magnification rule in the Kohonen SOM does not depend on the statistics of contact but on the size of the stimuli.

Chapter 6: Malleability of somatosensory cortex

In Chapter 6, we investigate the possible mechanisms for the changes in somatosensory cortical representation of the hand after a local anaesthetic nerve block to the index finger. We show how a simple mechanistic model implementing a homeostatic based plasticity mechanism can reproduce the change in fMRI measured cortical representation under local nerve block.

Contributions:

- We first create a toy model to demonstrate that distributed feedforward and lateral connectivity is required to reproduce the canonical hand representation.

- We simulate realistic tactile inputs from the hand before and after digit block. Using these inputs, we recreate fMRI cortical responses in our model. We find that without additional cortical changes, the fMRI responses under block cannot be reproduced.
- We demonstrate that homeostatic based mechanisms could reproduce the hand representation changes under local nerve block.

Related publications:

- Wesselink, D. B., Sanders, Z.-B., **Edmondson, L. R.**, Dempsey-Jones, H., Kieliba, P., Kikkert, S., Themistocleous, A. C., Emir, U., Diedrichsen, J., Saal, H. P., and Makin, T. R. (2020). Malleability of the cortical hand map following a finger nerve block. *BioRxiv*.

Chapter 7: General discussion

We conclude the thesis by summarising the overall findings of each chapter. We describe future avenues of research and broader considerations and applications of the findings.

1.3 Declaration of work contributed by others

Chapter 3: Derivations in section A.3 of the appendix were contributed by Jimenez-Rodrigo, A.

Chapter 6: Collection and analysis of fMRI data, psychophysics and MR spectroscopy was completed by: Wesselink, D. B., Sanders, Z.-B., Dempsey-Jones, H., Kieliba, P., Kikkert, S., Themistocleous, A. C., Emir, U., Diedrichsen, J., and Makin, T. R.

Chapter 2

Introduction: Sense of touch

Our sense of touch conveys highly dynamic and varied information, contributing to a rich perceptual experience of our environment– from feeling the shape and texture of a grasped object to the warmth of a loved one’s touch. Our experiences are highly influenced by the active nature of touch, with the stimulation we receive to our skin largely dependent on how we interact with the environment. Feedback from over 230,000 mechanoreceptors enables precise manoeuvring and manipulation of objects. Without this, everyday tasks can become arduous if not impossible, for example, being able to pick up a glass without applying too much force to break it, but enough, so it does not slip from our grasp. Furthermore, tactile sensing has its own unique set of complexities different from other senses.

2.1 Somatosensory periphery

2.1.1 Tactile afferents of the skin

The skin covers the surface of our body and is comprised of three layers, the outermost epidermis, followed by the dermis and subcutaneous layer (see Fig. 2.1). Within these layers, there are several classes of mechanoreceptive ($A\beta$ & C) alongside free nerve endings. Of the mechanoreceptors, low threshold mechanoreceptive C afferents respond to non-painful light, and ‘affective’ touch (Pawling et al., 2017; Löken et al., 2009). This thesis will focus on myelinated $A\beta$, which are associated with discriminative touch and active object interactions.

There are four types of $A\beta$ cutaneous mechanoreceptors, classified into either rapidly adapting (RA, alternatively Fast adapting, FA) or slowly adapting (SA), which each convey different information (Johnson, 2000). Rapidly adapting type I and II afferents terminate in Meissner’s and Pacinian corpuscles and slowly adapting types I and II in Merkel’s Discs and Ruffini endings, respectively. The afferents and their locations in the skin are shown in Fig. 2.1. RA afferents respond predominately to dynamic movement of the skin at the onset and offset of contact or stimulation, with no response to sustained indentation. SA afferents respond to both the onset and offset of stimuli and

sustained contact. In addition, afferents are responsive to certain stimulus qualities, such as the frequency and type of indentation. It was previously thought that separate afferent classes are responsible for specific perceptual qualities, such as texture, shape identification, or motion. However, during naturalistic tactile interactions, receptor classes are activated simultaneously. Therefore it is likely that a combination of their responses underlies tactile experiences (Saal and Bensmaia, 2014).

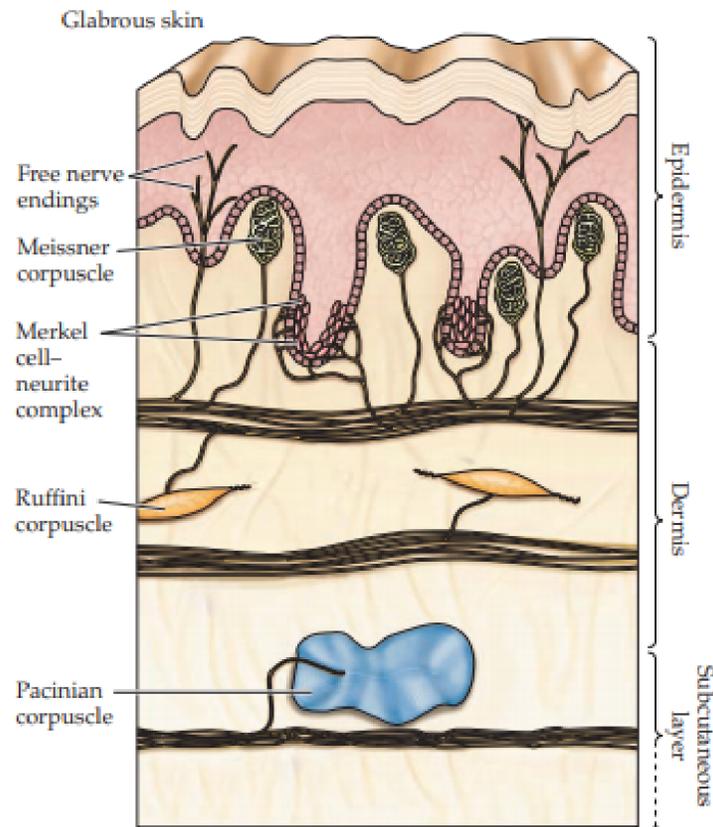


Figure 2.1: Cross-section of the glabrous (non-hairy) skin, demonstrating the layers of the dermis and placement of mechanoreceptors and other afferents. Reproduced from Purves et al. (2018).

The afferents are distributed throughout the body in different densities, with the largest being over the hands and lips (see Corniani and Saal (2020) for an extensive review). The densities may be correlated with behavioural relevance, as these regions serve essential functions. For example, the hands are heavily involved in tactile interactions and manipulation of objects in the environment; the mouth and lips are involved in language production and eating, which is crucial for survival. Here we will focus on the hand as it has been studied most extensively out of all body parts.

Three major nerves innervate the hand—the radial, median and ulnar. The radial nerve innervates mostly the skin dorsum; the median innervates the glabrous (non-hairy) skin and dorsum tips of

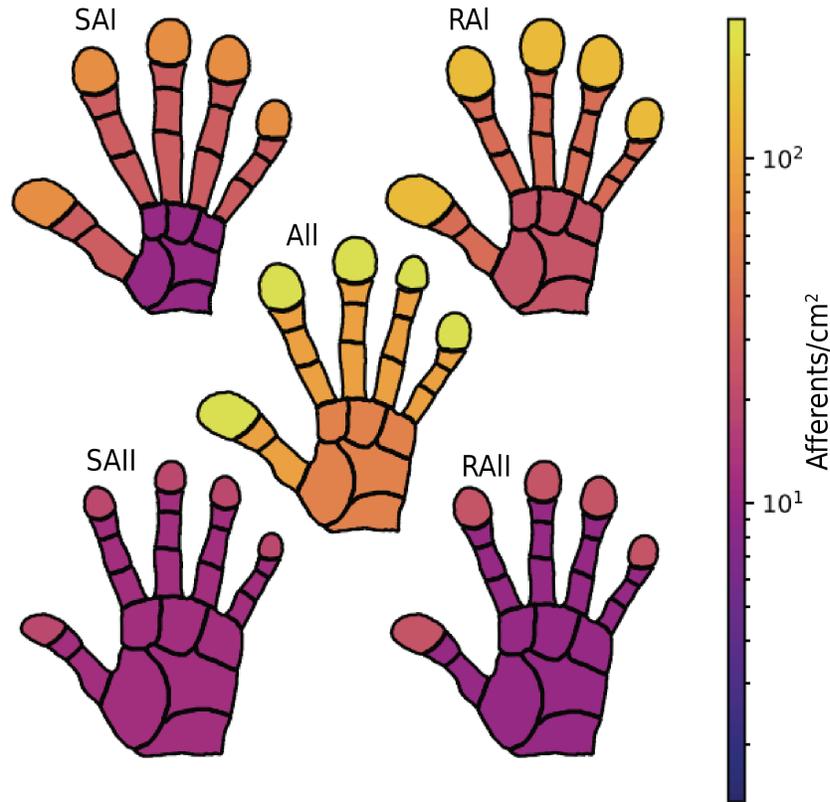


Figure 2.2: Diversity of afferent densities across the hand. Typical densities of the four afferent classes are shown. Colourbar denotes number of afferents per cm^2 . Larger sizes and brighter yellow colours indicate higher densities. Both SAI and RAI are denser in the fingertips than RAIi and SAIi, which are more evenly distributed across the hand surface. Densities within the medial and proximal phalanges are similar. Figure reproduced from Corniani and Saal (2020).

the thumb, index, middle and ring fingers. The ulnar innervates the dorsum and glabrous regions of the little and part of the ring fingers.

The densities of afferents are not equal across the body, with the glabrous skin of the hand containing 17,000 afferent fibres (Johansson and Vallbo, 1979). The density of afferents is greatest in the fingertips than any other part of the hand (Johansson and Vallbo, 1979), which is mostly driven by the increase in more populous RAI and SAI units (see Fig. 2.2). In contrast, RAIi and SAIi afferents are more evenly distributed throughout the hand. Overall, there are slightly more rapidly adapting fibres in the hand (43% RAI, around 13% RAIi) than slowly adapting (25% SAI, 19% SAIi) (Vallbo and Johansson, 1984). Between digits, the innervation densities of different phalanges are relatively comparable (Paré et al., 2002).

2.2 Hand use and statistics of tactile interactions

The pattern of tactile afferent activation is determined by both active uses of our hands and passive stimulation. Considering active touch, how we use our hands to grasp and manipulate objects determines where and what kind of stimulation the afferents receive. Therefore, the statistics of hand-use should be considered. Nevertheless, due to the difficulties of collecting tactile data, few studies have attempted to quantify hand-use. This contrasts to other senses, where statistics can be directly calculated from the stimuli due to mostly passive information flow. For example, in vision, natural images provide examples of typical patterns or structures in the environment that our retina receives (Kersten, 1987). In audition, recordings of environmental sound can be analysed (Nelken et al., 1999; Moerel et al., 2012). In contrast, quantifying natural touch is more complex—touch is an active sense, and the skin is a 2D deformable surface providing 3D contact information, with complex skin mechanics determining the exact pattern of stimulation (Manfredi et al., 2012; Sripathi et al., 2006). Although the possible space of hand movement and digit contact combinations is large, interactions tend to be stereotyped (Feix et al., 2016). The position and shape of the hand depend on the geometry of the object and the forces required to manipulate it (Cutkosky and Wright, 1986; Cutkosky, 1989; Taylor and Schwarz, 1955). As a result, several grasp taxonomies have been developed, enabling broad categorisation of similar movements. Initial work segmented these into two categories of power grasps and precision grips (Cutkosky, 1989). Power grasps are used when force needs to be applied to objects and can involve large areas of the hand making contact with the object, such as the palm and all the digits. Two power grasp examples are shown in Fig. 2.3A, top row. Precision grips are used for the manipulation of small objects and rely heavily on the fingertips and certain digits such as the thumb and index finger (see Fig. 2.3A, bottom row). Further refinement of the Cutkosky (1989) classification resulted in the well-accepted GRASP taxonomy, which consists of 33 independent grasps (Feix et al., 2009, 2016). In this, a further class of ‘intermediate’ grasps that use elements of both power and precision was identified, see Fig. 2.3A, middle row. Aside from grasping, exploratory hand contacts are made with objects to explore their features, such as texture or curvature. Lederman and Klatzky (1987) classified these movements into six types which are shown in Fig. 2.3B. Some examples include lateral motion, where fingertips are moved across the object to examine its texture (Fig. 2.3B, top left) and contour following, where edges of the objects are distinguished for shape classification (Fig. 2.3B, bottom right). One notable feature of exploratory actions is their heavy involvement of the digit tips.

Collecting naturalistic hand usage data is difficult due to the active nature of touch. As a result, little is known about the exact statistics of contact for different skin surfaces of the hand. Gonzalez et al. (2014) calculated the frequency of hand region contacts based on videos of natural hand usage in housekeepers and factory machining workers (Zheng et al., 2011; Bullock et al., 2013). The grasping types used in the videos were classified according to both the GRASP taxonomy (Feix et al., 2009) and exploratory movements of Lederman and Klatzky (1987). The typical frequencies

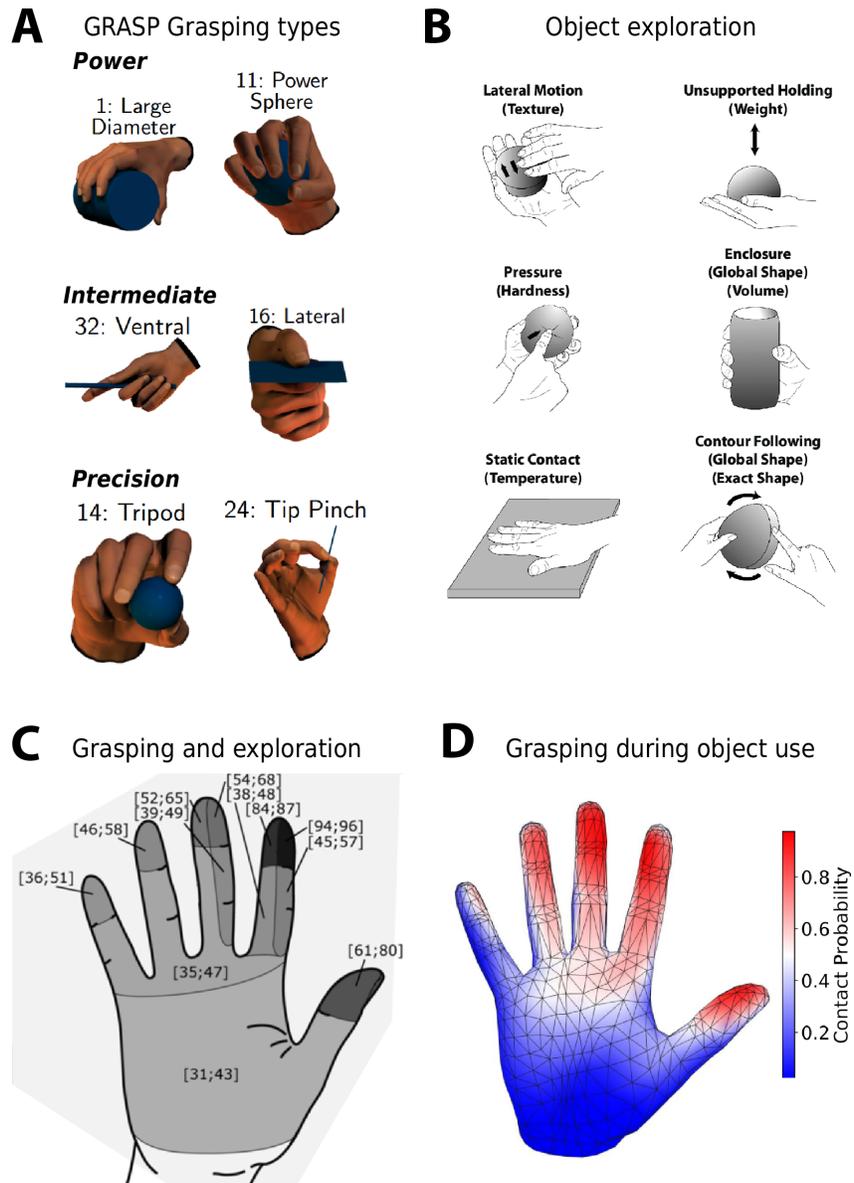


Figure 2.3: Hand interactions and skin contact probabilities. **A.** Examples of the three different classes of grasp types from the GRASP taxonomy. Power grasps can involve contacts with large hand areas, such as all digits and the palm. Precision grips involve greater fingertip contact for fine manipulation of smaller objects. Intermediate grasps involve a combination of both, alongside lateral types of movement. Reproduced from Feix et al. (2016) © 2016 IEEE. **B.** Six hand movement types when exploring objects from the classification of Lederman and Klatzky (2009). **C.** Typical contact across hand regions from analysis of different grasps and exploratory movements in everyday environments. Numbers refer to the percentage of use for each region; darker colours indicate greater use. These are given as a range of possible usage from different combinations of exploration and grasping contacts. Reproduced from Gonzalez et al. (2014). **D.** Contact probabilities from grasping interactions with 25 objects where the goal was to use the object. Reproduced from Brahmhatt et al. (2020).

of these movements were used to infer the amount of contact for the digits and palm of the hand. These are shown in Fig. 2.3C. Gonzalez et al. (2014) noted a correlation of usage with the density of afferents in the fingers. However, some fingers, such as the thumb and index finger, were found to have more contact than other digits and the palm (see darker shaded regions in Fig. 2.3C), corresponding to the larger magnification of the thumb and index finger in somatosensory cortex (Martuzzi et al., 2014; Penfield and Boldrey, 1937; Overduin and Servos, 2004), this will be discussed further in section 2.4.4. Other studies of specific interactions, for example, touching of the face during driving, also found most contacts were made with thumbs or fingertips (Ralph et al., 2021). Finally, improvements in tactile data collection with the development of higher spatial resolution tactile gloves (Sundaram et al., 2019), and use of machine learning algorithms in decoding and reconstructing hand contacts from videos (Grady et al., 2021; Cao et al., 2020; Brahmabhatt et al., 2020; Shan et al., 2020) have led to the creation of bench-marking datasets for tactile interactions (Brahmbhatt et al., 2019). However, these studies are mostly focused on the development of accurate recording methods and, therefore, analysis of collected data has been limited. From recordings of grasping with 25 different objects, Brahmabhatt et al. (2020) reconstructed the corresponding contacts from different regions of the hand. This demonstrated similar results to that of Gonzalez et al. (2014), with greater contact of the thumb, index and middle fingers, weighted towards the fingertips (see Fig. 2.3D).

2.3 Tactile information ascension to cortex

Information from tactile afferents ascends to somatosensory cortex with other peripheral information, e.g. muscle receptor activity, via the spinal cord. Although tactile signals are sent to many locations in the brain, such as motor cortex (Asanuma et al., 1980; Fang et al., 2006) and the cerebellum (Boillat et al., 2020), here we focus on ascension to the somatosensory cortex via the thalamus.

2.3.1 Dorsal column medial lemniscus pathway

The dorsal column medial lemniscus pathway carries tactile information from peripheral body parts to somatosensory cortex. The full pathway is depicted in Fig. 2.4. From the hand, the three main nerves (radial, median and ulnar) connect to the spinal cord roots at the brachial plexus. They ascend ipsilaterally to the cuneate nucleus at the level of the medulla. Within the dorsal columns, fibres are organised in a topographic manner (Smith and Deacon, 1984). The lower body is carried in a bundle known as the fasciculus gracilis, and the upper body in the cuneate tract. These go to different subdivisions, the gracile nucleus and the cuneate nucleus, respectively. The first-order neurons (afferents) synapse onto second-order neurons in the dorsal column nuclei (DCN) of the medulla of the lower level of the brainstem. It is thought that little processing occurs at this level, and instead, it mostly relays information (Witham and Baker, 2011). However, there is some evidence of cuneate nucleus activation being modulated by descending cortical inputs,

predominately excitatory effects, but also some inhibitory via inter-neurons (Aguilar et al., 2003). These may have some role in enhancing ascending tactile information and modulation of activity during movement.

Initial somatotopic arrangement in the DCN was shown in cats (Kruger et al., 1961). This demonstrated a largely stereotyped map of the body representation, which has been confirmed over several mammal species, demonstrating a mediolateral organisation progression of the tail to upper leg in the gracile nucleus, and forelimb to neck in the cuneate (Florence et al. (1988); Qi and Kaas (2006), see also Loutit et al. (2021) for a full review of DCN representations and discussion of differences between species). More recently, the cuneate has been examined using electrophysiological methods, a challenging task given the small size of brainstem structures and accessibility (Suresh et al., 2017), which enabled precise identification of the nuclei locations and boundaries. Using cutaneous vibratory stimuli, a somatotopic representation of the body was confirmed in rhesus macaques.

2.3.2 Thalamus

The ventral posterior nucleus (VPN) of the thalamus is the principal relay for information from cutaneous mechanoreceptors to somatosensory cortex¹. The VPN is split into the medial VP (VPM), which receives information from the face, and the lateral VP (VPL), from all other body parts. The VP has a complete topographically ordered representation of the contralateral body mapping (Padberg et al., 2009). The hand has a mediolateral ordered digit representation within the medial VPL, where each digit is separated by a septa boundary (Qi et al., 2011). The digits occupy a large portion of the VPL in humans, and the lips and tongue, a large area of the VPM (Lenz et al., 1988). Most connections from the VP are to a subregion of somatosensory cortex—area 3b, but also project to areas 3a, 1, 2 and 5. Most connections to 3a are deep inputs from the anterior VPL (Friedman and Jones, 1981).

Considering the processing of tactile inputs, most neurons are activated by RAI and SA inputs, segregated into small clusters of each (Lenz et al., 1988). There is little evidence on the extent of somatosensory processing in the thalamus, and it is currently thought that thalamic responses are similar to those of primary afferents. Therefore, information from the thalamus may simply be relayed to cortex, without consideration to cognitive task demands (Camarillo et al., 2012). This suggests that the thalamus is not engaged in working memory, decision making or modulated by attention.

2.4 Somatosensory cortex

The somatosensory cortex is the main sensory processing region for tactile information from the body. The cortical surface is divided into five layers, with most thalamic information entering at layer IV (Hubel and Wiesel, 1962). Penfield and Boldrey (1937) initially mapped the somatosensory

¹In humans, the VPN is termed the Ventral Caudalis (VC).

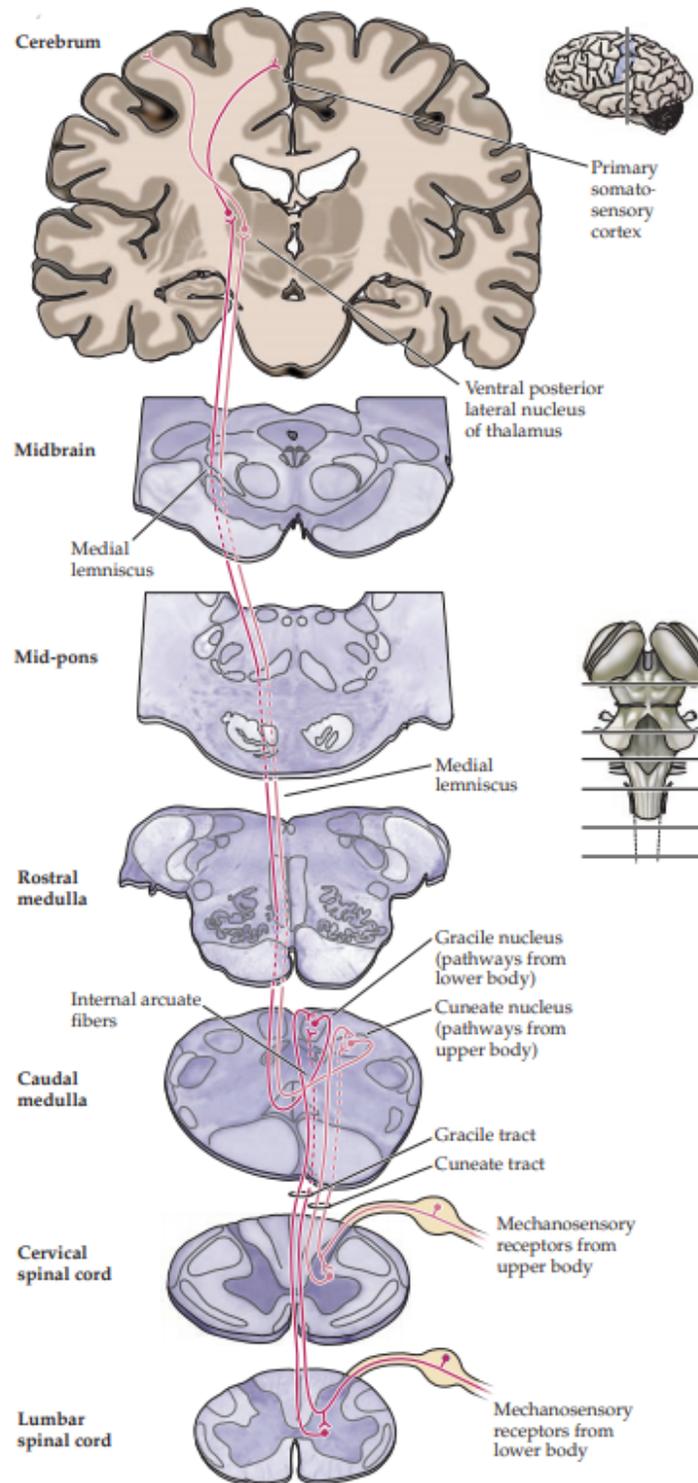


Figure 2.4: Ascension of cutaneous tactile sensory information from mechanoreceptors to somatosensory cortex. Right inset shows the brainstem, horizontal lines denote the hierarchy of regions depicted on left. Adapted from Purves et al. (2018).

cortex in humans and found a complete representation of the body surface, arranged topographically, such that neighbouring body parts are generally mapped adjacent to each other. The feet are most medial in the midline, and the face most lateral. This representation is referred to as a map of the body, and the textbook depiction of this arrangement is Penfield's famous homunculus (Penfield and Rasmussen, 1950), shown in Fig. 2.5. This arrangement of the body has been replicated in non-human primates (NHPs) (Nelson et al., 1980; Krubitzer et al., 2004; Krubitzer and Kaas, 1990; Sur et al., 1980), and in humans (Sanchez Panchuelo et al., 2018; Saadon-Grosman et al., 2020; Roux et al., 2018; Sakai et al., 1995).

Alongside ordered mapping of the body parts, another property of somatosensory representation is cortical magnification. This phenomenon is where regions of a sensory input surface are magnified in cortex beyond their proportional body size. This magnification property is seen across cortical sensory regions, such as in vision, where the fovea of the retina is magnified in the primary visual cortex (Wässle et al., 1990; Connolly and Van Essen, 1984), and in audition, where behaviorally relevant sound frequencies are magnified (Recanzone et al., 1993).

In somatosensory cortex, the glabrous skin is only 5% of the total skin surface but occupies around 30% of the representation in primates (Lehnert et al. (2021) based on data from Sur et al. (1982)). Within this, the glabrous surface of the hands are greatly magnified (Penfield and Boldrey, 1937; Sur et al., 1980). The exact cause of somatosensory magnification is unknown; however, it has been proposed to be related to either the peripheral innervation density of the skin (Catani, 2017) or hand use (Merzenich et al., 1984). For example, the hands tend to have a larger representation and a higher innervation density. Similarly, whiskers of rodents are organised in barrels in somatosensory cortex, with proportional numbers of neurons to whisker innervation (Lee and Woolsey, 1975). Alternatively, it may be beneficial for the sensory system to allocate more of its representational resources to body parts that are more frequently used. For example, the magnification of the hands and lips may be due to their higher usage from interaction with objects (Gonzalez et al., 2014), eating and communication.

2.4.1 Initial somatotopic map formation

Several factors drive the emergence of the adult somatosensory map. First, the broad topography demonstrated by Penfield and Boldrey (1937) is highly stable across individuals and likely set up from genetic signalling processes during early development. Although most evidence for this is from visual cortex, processes are likely similar across sensory cortices. Early work on the retinal to primary visual cortex mapping demonstrated that initial topography is set up from molecular signalling during fetal development. These labels from eph/ephrin ligands and receptors specify map positioning (Flanagan, 2006; McLaughlin and O'Leary, 2005). Gradients lead to a set of positions for growth targets, guiding axonal projections and ensuring local neighbourhood relationships (e.g., topographic ordering is obeyed). For example, it has been found that disrupting these molecular gradients in mice can lead to distorted somatosensory map sizing (Vanderhaeghen et al., 2000).

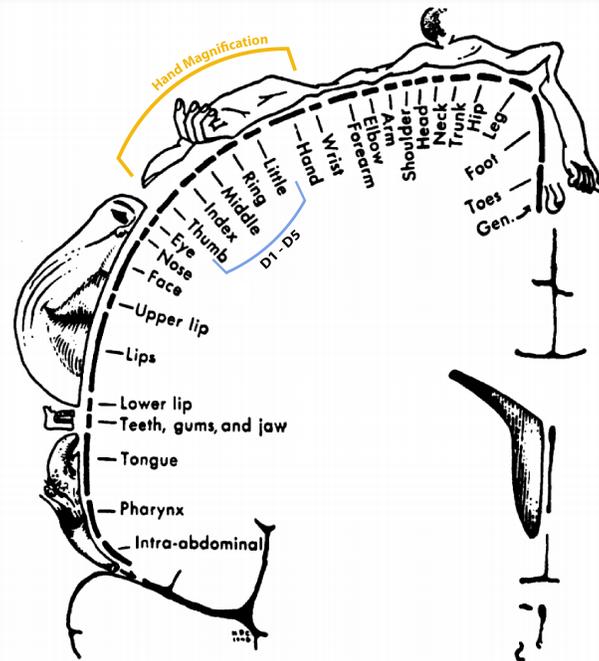


Figure 2.5: Somatosensory homunculus demonstrates topographic ordering of body parts and magnification of areas such as the hands and face. Feet are represented in the medial midline, and the face and intraoral structures most laterally. Adapted from Penfield and Rasmussen (1950).

Following birth, in early postnatal development, there is a restricted stage of increased large-scale plasticity—the critical period—in which experience is vital to ensure the development of adult-like representations. As with molecular gradients, most investigation on critical periods has focused on visual cortex. Notable is the work of Wiesel and Hubel (1963) who demonstrated that ocular dominance columns in cats only developed for one eye when there was monocular deprivation for two to three months from birth. However, if there was some visual experience before suturing, ocular dominance columns developed normally.

There is evidence that the coarse body representation is already formed at birth for the somatosensory cortex. In studies of cats, body maps were found to have an adult-like construction early in development (Rubel, 1971). In NHPs, some differentiation has been found in the presentation of maps at birth, dependent on the species. For new world monkeys, such as squirrel monkeys and marmosets, maps were found at birth (Krubitzer and Kaas, 1988). In contrast, maps in old-world macaque monkeys were only present around one month after, with a lack of neuronal response at birth to tactile cutaneous stimulation (Krubitzer and Kaas, 1988). However, more recently, fMRI studies have revealed that coarser maps may be present at birth, which then undergo rapid development within the first days of life. Arcaro et al. (2019) found that by two to seven months, the representations of individual digits could be identified, coinciding with the development of motor skills. While a coarse representation existed, the responses were undeveloped in more immature animals at the earlier months of life, which corresponds to Krubitzer and Kaas (1988). Differences

noted between old and new world monkeys are also reflected in behaviours, with new world monkeys requiring less assistance from their mother in climbing unaided (Shively and Mitchell, 1986; Krubitzer and Kaas, 1988).

In line with these results, EEG studies with human infants have demonstrated the presence of only coarse somatotopy within the first year of life (Meltzoff et al., 2019; Saby et al., 2015; Rigato et al., 2014). A refinement process of these early maps was seen from around six months, with adult-like responses appearing two years from birth. EEG studies can demonstrate coarser representation but not more minor spatial detail. Thus these studies can only tell us that a map exists, but not precise locations of body parts. 3T fMRI has recently been used to study preterm infants (Dall’Orso et al., 2018). This demonstrated that a topographic body map exists even in early development, with correct positioning of key body parts such as the mouth and upper and lower extremities. Dall’Orso et al. (2018) did not investigate the finer representations, such as individual fingers, so we cannot infer whether finer-grained representations exist at this stage or not in humans. Nevertheless, these studies show that, as with non-human primates, general topography exists early on in human infancy.

It has been widely demonstrated that plasticity continues into adulthood, but this may be different from that of critical periods. For example, early development may concern the fixing of thalamocortical pathways, whereas adult plasticity may enable changes in cortical lateral connections (Marik and Gilbert, 2017). Plasticity on a finer scale can continue into and throughout adulthood and will be discussed in chapter 6. This may drive individual differences in somatotopy, which have been documented frequently in animal models and humans.

2.4.2 Properties of subregions of S1

Somatosensory cortex can be segmented into subregions 1, 2, 3a and 3b (Brodmann, 1909). These areas have been identified in primates using electrophysiological methods (Kaas et al., 1979; Kaas, 1983; Pons et al., 1985; Merzenich et al., 1987; Nelson and Chen, 2008) and humans using fMRI (Sanchez-Panchuelo et al., 2010; Martuzzi et al., 2014; Kurth et al., 2005; Francis et al., 2000; Besle et al., 2014). Area 3a is located in or near the fundus of the central sulcus, with area 3b following, then 1 and 2, in a rostral to caudal arrangement. There is a general hierarchy of processing, with areas 1 and 2 further downstream, integrating information from primary regions 3a and 3b (Iwamura et al., 1993).

This thesis will focus predominantly on area 3b representations; however, we will briefly discuss 3b’s differences to areas 3a, 1 and 2. Several notable reviews cover the broad properties of processing in these areas, such as Delhaye et al. (2018). An in-depth discussion of somatosensory cortex responses to tactile information of the hand will follow, as this region is most studied.

Area 3b

Area 3b is considered the first processor of most cutaneous information, usually referred to as the ‘primary’ somatosensory cortex. It takes input from mainly cutaneous receptors via the thalamus (Kaas et al., 1979). In monkeys, 77% of neurons in 3b respond to cutaneous and 20.9% respond to deep (proprioceptive) stimuli (Iwamura et al., 1983). As a result, area 3b has the most evident somatotopic gradient of organisation across the subregions, with easily identifiable representations of body parts, such as individual fingers.

Area 3a

Neurons in area 3a respond predominately to proprioceptive signals such as joint and muscle movements. There are several reviews of 3a such as Jones and Porter (1980); Delhaye et al. (2018), and a detailed study in marmosets by Krubitzer et al. (2004). See also Lutz and Bensmaia (2021) for a recent review on proprioception. Area 3a is connected to neighbouring motor cortex and supplementary motor areas (Darian-Smith et al., 1993; Huerta and Pons, 1990; Huffman and Krubitzer, 2001), as well as area 2 (Huffman and Krubitzer, 2001). The receptive fields of afferents in this area can be broader than that of 3b, with responses to individual digits up to areas across the whole hand (Krubitzer et al., 2004).

Area 1

Neurons in area 1 respond mostly to cutaneous stimulation (approximately 90%, see Delhaye et al. (2018)). These seem to be driven by the region’s hierarchical relationship with area 3b, as lesions to 3b abolish area 1 activation (Garraghty et al., 1990). Area 1 is involved in spatial integration (Iwamura et al., 1993; Sur et al., 1980), and texture discrimination, as lesions to area 1 lead to impairment (Randolph and Semmes, 1974; Carlson, 1981). Area 1 is generally less somatotopic than 3b with representations having smaller inter-digit distances, and larger overlaps (Stringer et al., 2014; Nelson and Chen, 2008; Friedman et al., 2008)). Receptive fields are less selective, with some spanning multiple digits (Ashaber et al., 2014; Iwamura et al., 1983).

Area 2

As with Area 1, somatotopy in area 2 is less distinct (Pons et al., 1985), with larger RFs (some spanning multiple digits) (Hyvärinen and Poranen, 1978). Hand representation neurons in area 2 are responsive to both cutaneous and proprioceptive stimuli along a gradient, with cutaneous closer to area 1 and proprioceptive more caudal (Iwamura et al., 1993; Pons et al., 1985). They are 45% responsive to cutaneous stimuli and 55% to proprioceptive stimuli (Seelke et al., 2012). Area 2 is likely involved in shape and size processing as neurons show responses to complex shapes and tuning for curvature (response of multiple orientations presented simultaneously). Furthermore, lesions to the region lead to a lack of shape discrimination ability (Carlson, 1981). As area 2 integrates both proprioceptive and cutaneous information, it has also been implicated in stereognosis perception

(sensing of 3D structure in shapes) (Yau et al., 2016).

2.4.3 Somatosensory cortical layers

Similar to other neocortical regions, primate somatosensory cortex is divided into six layers numbered I-VI (Qi et al., 2008). In area 3, layer IV receives most input from upstream thalamic targets of tactile peripheral inputs (the VPL and VPM), followed by deep parts of layer III (Jones, 1975; Garraghty et al., 1989). In areas 1 and 2, thalamic projections target layer III (Jones, 1975). In these middle layers, topographic maps of tactile inputs are most prominent, with typically stronger selectivity of certain tactile inputs and larger response amplitudes to stimulation (Yu et al., 2019; Lee et al., 2021). Recorded receptive fields of neurons in these layers also tend to be smaller than surrounding layers in both area 1 and 3b (Sur et al., 1985). Similarly in rodents, the cortical barrel field has the most defined whisker representations in the input layer IV (Woolsey and Van der Loos, 1970).

Beyond layer III and IV, other layers are involved in secondary processing of the layer IV inputs. More superficial layers I-III have intracortical projections to and from other brain regions, aiding the integration of tactile information with other cortical networks (Kuehn and Pleger, 2020). Layers V and VI have stronger feedback connections to subcortical regions, such as the thalamus (Harding-Forrester and Feldman, 2018).

Studying the differences between the cortical layers in humans is challenging due to the low cortical thickness (around 2mm in area 3b). Furthermore, the fine-scale of digit representations mean they may occupy only a few voxels each, especially with lower spatial resolution imaging (Yu et al., 2019; Schluppeck et al., 2018). Recent advances with high field fMRI and the development of layer-specific imaging techniques have enabled some initial work to further understand differences between the layers. For example, more superficial layers are involved in the processing of predictable tactile inputs (Yu et al., 2019).

2.4.4 Investigating hand and digit representations in S1

Most tactile sensing research has focused on the representation of the hand. This is likely due to its behavioural importance and corresponding magnification in cortex, enabling easy localisation and mapping with electrophysiological methods and lower-field fMRI.

In area 3b, the hand is organised with the thumb most lateral, next to the face representation, then an ordering according to the digits, index, middle and ring, to the little finger which is most medial. Typically these are labelled D1-D5, where D1 represents the thumb and D5 the little digit. Example topography across S1 is shown in Fig. 2.6A. Broad digit ordering is highly reproducible across studies (Schweisfurth et al., 2011; Sanchez-Panchuelo et al., 2012; Schweisfurth et al., 2014; Besle et al., 2014; Stringer et al., 2014; Overduin and Servos, 2004; Cunningham et al., 2013; Schweizer et al., 2008). Gradients exist between the subregions with ‘mirror reversals’ along the

boundary of each subregion, such that the representation from digit tips to digit base flips along the border. Between 3a and 3b, the representations of the digit tips are aligned; between 3b and area 1, the digit bases are aligned, see Fig. 2.6B (Merzenich et al., 1987; Sanchez-Panchuelo et al., 2012; Sánchez-Panchuelo et al., 2014; Blankenburg et al., 2003). Variations between and within subregions have been noted and are further described below. These are typically assessed through measures such as cortical sizes of each digit’s representation, overlaps between these representations, their distances, and receptive fields of the neurons within them.

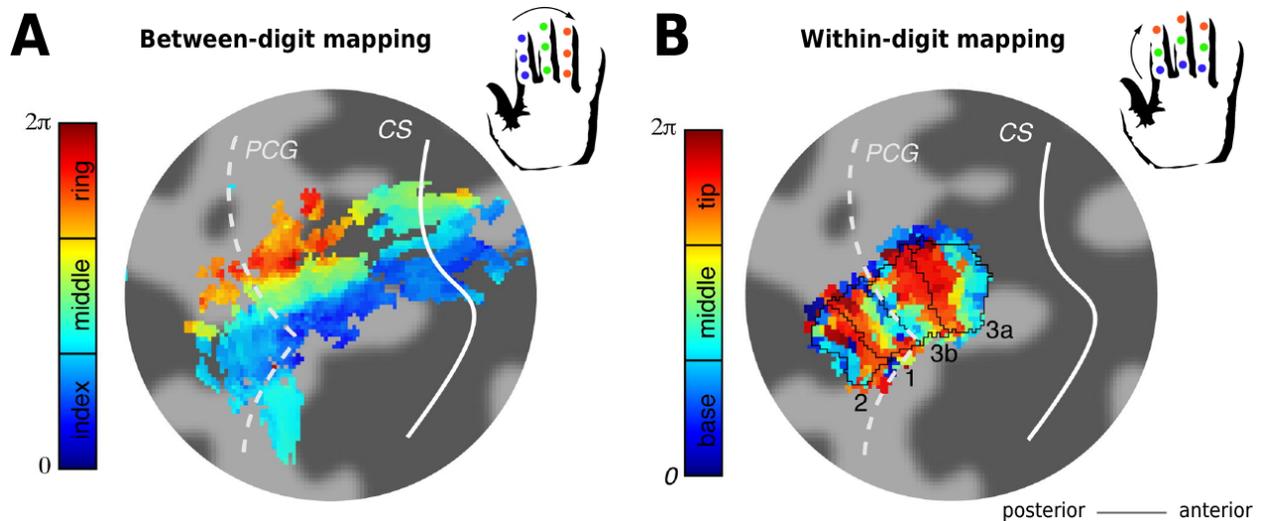


Figure 2.6: 7T fMRI mapping of between and within-digit representation. **A.** Representations of each digit, D2-D4 across S1 (areas 3a, 3b, 1 and 2). Across subregions, digits are topographically arranged. A strong gradient emerges lateral to medial (D2-D4) within the region corresponding to 3b, located right of the PCG (Post central gyrus). **B.** Within-digit mapping for each subregion of S1. Black outlines denote the subregions. ‘Mirror’ reversal occur along subregion borders, for example, the digit tips are aligned for region 3a and 3b. See also the large representation of digit tips (red) compared to other phalanges. Digit representations are assessed with a travelling wave design, where digits are stimulated sequentially. This can either be between digits D2-D4 (see arrow on hand inset in A) or within each digit’s phalanges (see arrow on hand inset in B). As the travelling wave is a periodic stimulation paradigm, the shift between stimulation of each digit or phalange is denoted in terms of a circular phase (where $0 - 2\pi$ radians denotes the full circle— stimulation cycle). The full phase circle is segmented by the number of digits or phalanges (here, 3). The colorbar denotes the phase of the fMRI signal in relation to when each digit or phalange is stimulated. More details of this paradigm can be found in Sanchez-Panchuelo et al. (2010). Figure adapted from Sánchez-Panchuelo et al. (2014).

Gradient of cortical digit sizes

Cortical sizes of digit representations are typically assessed through passive cutaneous stimulation of the digit whilst recording the extent of S1 responses. However, assessing sizes quantitatively across studies can be challenging due to differences in stimuli presented, processing of data, and spatial resolution of the imaging method. Here we will discuss general trends between cortical digit representation sizes.

There is a general gradient of cortical digit sizes in humans, with D1 being the largest and D5 the smallest, with some variation for the digits in-between. Following from Penfield and Bol-

drey (1937)'s initial mapping, many studies show the largest area is devoted to the thumb (D1) (Hämäläinen et al., 2000; Maldjian et al., 1999; Overduin and Servos, 2004; Martuzzi et al., 2014; Schweisfurth et al., 2018). In general S1, the thumb is around twice as large as the other digit representations (Martuzzi et al., 2014; Penfield and Boldrey, 1937). The index finger has also has a notably larger representation size (Maldjian et al., 1999; Overduin and Servos, 2004; Martuzzi et al., 2014) (however Krause et al. (2001) found a similar representation size to D3). This may be due to the index digit and thumb having a prominent role in tactile interactions, for example, in precision grips (see section 2.2). D4 is typically smaller than the first three fingers (Overduin and Servos, 2004), and D5 has the smallest area, which is relatively consistent between studies (Sanchez-Panchuelo et al., 2010). This could be due to its physical size or that it tends to be used less than other digits (Gonzalez et al., 2014).

Less overlap of activation in anterior somatosensory cortex

Somatosensory cortex activation is not entirely localised to a specific digit's region of interest (ROI); instead, activation can be more widespread across the hand representation. This is assessed through measuring overlaps between cortical activation, typically of neighbouring digits. There are several reasons for this more widespread activity. First, thalamic input is more widespread than to a single topographic region, meaning cortical neurons could have input from multiple digits, particularly neighbouring digits (Padberg et al., 2009). Second, tactile afferents are not only activated on individual digits, but activation is spread much further over the hand beyond the location of stimulation, which is picked up by PC afferents (Manfredi et al., 2012; Shao et al., 2020). This would lead to a response in neighbouring digit representations in area 3b (see more on this in chapter 6). Overlaps could also result from lateral connectivity between the neurons, which may extend across the whole hand region. Between regions, there is more overlap in posterior regions 1 and 2 than area 3b (Deuchert et al., 2002; Kurth et al., 2005; Krause et al., 2001; Besle et al., 2014). For example, areas suspected to be 1 and 2 show responses of up to all five-digit tips in any ROI, whereas other regions typically show an overlap of around three digits.

Distances reflect somatotopic organisation

In line with the somatotopic ordering, the largest distance is between the thumb and little digit compared to other digit pairs (van Westen et al., 2004; Gelnar et al., 1998). The ordering within D1-D5 representation is consistent across many studies in humans (Maldjian et al., 1999; Kurth et al., 2000; Schweizer et al., 2008; Sanchez-Panchuelo et al., 2010; Stringer et al., 2011; McGlone et al., 2002), and with NHPs (Nelson and Chen, 2008). The Euclidean distance typically increases relative to D1, indicating their strict somatotopic relationship. For area 3b, the distance decreases between D1 and D5, such that there is a smaller distance between D4 and D5 than D2 and D3 (Duncan and Boynton, 2007; van Westen et al., 2004; Schweisfurth et al., 2018; Martuzzi et al., 2014).

Between regions, in those more anterior such as 3a, 3b and 1, the distance separation tends to be larger (Deuchert et al., 2002). Separation is also greater for area 3b compared to area 1 (Kurth

et al., 2000; Nelson and Chen, 2008), with one study reporting a 1.6x increase (Stringer et al., 2011). A common feature in both 3b and area 1 is that the separation between D1 and D2 is greater than all other digits. This corresponds with the larger spatial extent occupied by the thumb in 3b (Nelson and Chen, 2008), as discussed above.

Within individual digits the fingertips are more magnified.

Higher tesla fMRI has enabled more fine-grained mapping of representations, including within each individual digit. This has enabled comparison between the fingertip, medial and proximal phalanges (Sánchez-Panchuelo et al., 2014). There is some discrepancy between studies on the extent of within-finger mapping. For example, some have found a generally ordered representation, with rostral to caudal ordered gradients running between the digit tip and each phalange (medial then proximal) and digit base at the top of the palm (Sánchez-Panchuelo et al., 2014; Blankenburg et al., 2003). However, other studies have demonstrated more mixed results for the ordering of the phalanges, which could reflect individual differences in hand use for these digits (Schweisfurth et al., 2014).

The sizes of each digit are reflected in their typical within-region distances. Mean distances from tip to base were 6.7mm for D2 and 5mm for D5, suggesting a larger overall representation of D2 (Schweisfurth et al., 2011). Typically the distal fingertip has a much larger cortical representation compared to other phalanges (Schweisfurth et al., 2014; Sur et al., 1980; Overduin and Servos, 2004). This could be due to the higher density of afferents (see section 2.1.1), or that typical digit usage is weighted towards the tips (Gonzalez et al., 2014).

Two studies found consistent within-digit ordering of the little digit's phalanges (D5) across participants but not for the index (D2) (Schweisfurth et al., 2011, 2014). This may be due to hand use of D5 being more stable across participants than other digits. For example, D4 and D5 may have a more supporting role in grasping and manipulation, and therefore more similar use.

Causes of hand magnification

Why is the hand magnified in somatosensory cortex? Magnification has been shown to occur in many sensory regions. Two causes of this have been proposed, firstly, the density of afferents is much greater in areas that are more magnified, for example, the digit tips of the hand (Sur et al., 1980), and the fovea of the retina (Wässle et al., 1990; Cowey and Rolls, 1974). Therefore the magnification could simply reflect the innervation densities (Catani, 2017). Nevertheless, in touch, each digit can also be magnified in a way that does not correspond to its afferent innervation, and innervation densities and cortical area size are only weakly correlated (Corniani and Saal, 2020). As a result of this mismatch, others suggest an important role of hand use (Merzenich et al., 1984). Studies from NHPs have found that manipulating the hand inputs can change the size of digit representations, such that those that are used more frequently have a larger representation (Wang et al., 1995; Dutta et al., 2014; Merzenich and Jenkins, 1993). See also Vidyasagar et al. (2014) in humans, and chapter 6 on plasticity. These studies suggest hand use has a role in determining cortical mapping of the hand representation.

Hand use could explain individual differences in magnification and mapping

Individual differences can include differences in somatotopic ordering, for example, reversals of digits (Schweizer et al., 2008; Stringer et al., 2011; Cunningham et al., 2013; Gelnar et al., 1998), representational sizes, overlaps and distances. Fine map details may be refined through experience, reflecting hand-use throughout life and changing through plasticity to match these (Merzenich et al., 1987). Nevertheless, there are some inconsistencies between maps for NHPs, particularly in the overlap of representations. This could be due to broad spatial and temporal input similarity, which may be more common in NHPs than humans due to differences in hand use. For example, humans have opposable thumbs with high tactile acuity, meaning it is possible to use precision grips when manipulating objects (Napier, 1956). Input statistics may be affected by several factors, such as the size of hands or dexterous and fine motor abilities, that will lead to differences in how objects are grasped and manipulated. There are some species-level (Kaas, 1983) and individual differences between mappings in somatosensory cortex. For example, spider monkeys use prehensile tail movements in the manipulation of small objects, rather than thumbs (Jones and Lederman, 2006). Individual differences in humans have also been found within maps of the hand. For example, in area 3b, there can be a missing representation of the little digit, twisting of the finger representations, and variations in digit overlap of between 3-34% (Schweizer et al., 2008).

Receptive fields and the inverse cortical magnification relationship

Cortical receptive fields (RFs) describe features or areas of an input space where a neuron elicits its maximal response. For example, in touch, RFs are typically used to describe the extent of the skin surface to which a neuron is responsive. RFs have been characterised more thoroughly in NHP studies than in humans as they are typically measured using electrophysiological recordings.

The size of RFs varies both across the hand and between S1 subregions. For example, in areas 3a, 1 and 2 receptive fields are typically larger than 3b (Sur et al., 1985; Merzenich et al., 1978; Hyvärinen and Poranen, 1978), and can span multiple digits (Iwamura et al., 1983, 1985). In area 3b, they are mostly restricted to single-digit pads or individual fingers (Xerri et al., 1999). Digit tips typically have smaller RFs than the rest of the hand, in line with the inverse magnification rule—larger magnification of a region leads to smaller corresponding RFs of that region (Sur et al., 1984, 1980; Sanchez Panchuelo et al., 2008; Nelson et al., 1980). This is particularly notable in area 3b (Sur et al., 1980; Xerri et al., 1999; Pons et al., 1987; Merzenich et al., 1978). The inverse magnification relationship may be a universal feature of topographic maps, as it has also been identified in other sensory regions, for example, in visual cortex (Harvey and Dumoulin, 2011; Connolly and Van Essen, 1984; Chaplin et al., 2013; Wässle et al., 1990; Perry et al., 1984).

2.5 Tactile perception

Tactile acuity is the ability to discriminate between different stimuli placed on the skin. A higher acuity means that stimuli placed at a smaller distance apart can be distinguished. Tactile acuity is typically greatest on the fingertips (Mancini et al., 2014), where there is a higher density of receptors,

making use of the ability to perceive small discrepancies on objects such as fine structure. Acuity varies between digits, being highest for the index and decaying towards the ring (Vega-Bermudez and Johnson, 2001; Grant et al., 2006) and little digits (Duncan and Boynton, 2007). Acuity can be measured using several techniques. For two-point discrimination, stimuli are placed at a measured distance, and participants are asked to judge the number of stimuli perceived. Alternative methods involved the detection of fine hairs moved along the skin, or tasks detecting the positioning of raised bumps, or orientation of grating domes. Spatial acuity is mainly driven by SA1 afferents, which have small receptive fields (Peters et al., 2009), and a corresponding high spatial resolution. The spacing of SA1 afferents and two-point discrimination thresholds are also highly correlated (Corniani and Saal, 2020). When innervation is low, afferents have an increased distance to their neighbours, such that they are unable to detect a stimulus placed between them.

Tactile perceptual abilities are also limited by fingertip size. SA1 Merkel cell numbers were found to be stable across fingertip size, meaning that they are more densely packed in smaller fingers (Peters et al., 2009). This may cause notable differences between acuities in people of different ages and sexes. For example, children and females tend to have lower discrimination thresholds (Peters and Goldreich, 2013; Peters et al., 2009), as the size limits the lower bound for tactile acuity. Therefore, the smaller the fingertip size, the greater the tactile acuity ability.

Acuity has also been related to cortical magnification in both touch and vision (Duncan and Boynton, 2003, 2007). Correlations between larger cortical area size and increased tactile acuity have been found across S1 subregions (Duncan and Boynton, 2007; Härtner et al., 2021). This may be due to larger, more magnified regions typically having smaller receptive fields, allowing for the detection of stimuli within a smaller skin region. Although receptive field sizes do not scale exactly with cortical magnification, they follow closely, with a $-2/3$ power rule in somatosensory cortex (Sur et al., 1980). Similarly, increased representation size induced by rTMS (repetitive transcranial magnetic stimulation) has been found to improve tactile perception for the index digit (Tegenthoff et al., 2005).

Spatial acuity may also be trained; for example, blind braille readers learn to distinguish fine bumps for reading and have correspondingly lower orientation detection thresholds (Van Boven et al., 2000). However, as mentioned above, training of tactile discrimination may be limited by the size of the fingertip. This provides a minimum for the acuity due to the differences in spacing of the receptors (Wong et al., 2013). Overall, this suggests some relationship between the density of receptors and acuity.

2.6 Modelling the somatosensory system

In comparison to visual sensing, relatively few advances have been made in modelling somatosensory cortex development, organisation and function (Goodhill, 2007). Previous modelling of somatosensory cortex has mostly been motivated to understand cortical plasticity in the brain due to exper-

iments in NHPs that demonstrated extensive reorganisation of body representations (Merzenich et al., 1984). In comparison, research on visual sensing has extensively modelled the pathways and cortical representations. One of these is the theory of efficient coding, which uses normative models to formalise the system’s goals in terms of a cost function. This can then be optimised under specific constraints or assumptions. In chapter 3 we use linear normative models previously applied in vision to understand the contribution of receptor density and input statistics when the amount of information is constrained by a bottleneck. In chapter 4, we then apply this to the star-nosed mole, to test whether the model can predict empirically measured cortical sizes.

In contrast to these abstract normative models of representations, in chapter 5 we utilise the Kohonen-SOM model, which can reproduce somatosensory map features such as topography. We investigate changing the receptor densities and statistics in this model, and the resulting magnification and receptive fields.

In the final chapter, we model changes to digit representations under a short term input manipulation that was measured with high-resolution fMRI. Here, we develop a mechanistic model to investigate the structure of connections to and within somatosensory regions and additional cortical mechanisms that demonstrate plastic changes. This modelling was enabled by recent advances in peripheral models of touch, which enable responses of mechanoreceptors to be accurately modelled in response to stimuli (Saal et al., 2017). Using this we could differentiate whether somatosensory activation changes following the digit block emerged due to widespread activation across the hand or cortical changes.

Chapter 3

Nonlinear scaling of resource allocation in sensory bottlenecks

3.1 Introduction

Our sensory systems provide a constant stream of information about our environment, which the brain must optimally process and utilise. This information is received through a set of receptors which can vary in their densities between input regions. This is a prominent feature in both touch and vision, for example, the greater densities of mechanoreceptors on the digit tips (Johansson and Vallbo, 1979) and photoreceptors in the fovea (Goodchild et al., 1996; Wells-Gray et al., 2016), respectively. Furthermore, the input statistics can also vary, such as the frequency of receptor activation by environmental stimuli. How should different input regions be represented, given a limited capacity of neural resources available?

In this chapter, we investigate the problem of resource allocation considering both receptor density and input statistics. We develop a theory using an abstraction of this problem, where the cost function can be precisely defined and solved analytically. Our ideas are based on previous works utilising ‘efficient coding’ theory, where extensive work has characterised optimal processing in the visual system. However, this has primarily focused on the statistical structure of inputs, while optimal representation for different densities of input receptors has received little prior attention. Given this, we choose a simple form of efficient coding for our model— maximising information via redundancy reduction. The following section provides an overview of redundancy reduction techniques and related efficient coding works.

3.1.1 Theories of efficient coding

Efficient coding theory, based on information-theoretic ideas of Shannon (1948), provides a framework for how the brain and nervous system should seek to maximise incoming information about the

environment. The term "efficient coding" incorporates several lines of investigation (Chalk et al., 2018), including redundancy reduction via decorrelation (Barlow, 1961; Attneave, 1954; Atick and Redlich, 1990; Atick, 1992), maximising neural and metabolic efficiency (Clarke and Sokoloff, 1994; Ames, 2000), sparse coding techniques (Olshausen et al., 1996; Olshausen and Field, 1997, 2004) and predictive coding (Palmer et al., 2015; Singer et al., 2018). The goals of efficient coding can therefore vary between determining what information is relevant and what can be encoded given the constraints of the system (Chalk et al., 2018). The definitions and assumptions of these different ideas are further outlined below.

Redundancy reduction via decorrelation

Barlow (1961) first proposed efficient coding as a theoretical model of neuronal sensory encoding. For systems to make the most efficient use of available resources, they should utilise their total coding capacity. The neural code should seek to faithfully represent the information in the input signal, for example, by matching its statistics. Through this, the aim is to maximise the mutual information between the input and the encoding. In most cases, the channel capacity is lower than the input dimensionality, such that all information cannot be transmitted— an information bottleneck (Atick and Redlich, 1990). One example of this is the optic nerve, where there are many fewer retinal ganglion cells than photoreceptors in the retina (Wells-Gray et al., 2016). In this case, we can take advantage of the high redundancy in sensory signals due to spatial interdependence on the sensory input sheet and across the timing of incoming signals (Attneave, 1954). Removing this redundancy makes it possible to transmit to a lower capacity output while minimising loss of information content.

The visual system may solve the efficient coding problem hierarchically, where redundancy is reduced at each stage. This results in a set of successive linear transfers for each stage in the visual processing stream. The typical sensory input to the visual system is itself redundant, with images of natural scenes having an amplitude that drops off at a factor of $1/f$ with increasing spatial frequency (Field, 1987) (power spectrum falls off with $1/f^2$ (Graham et al., 2006)). This means that neighbouring pixels are highly correlated in natural images, accounting for around 40% of the redundancy (Graham and Field, 2010; Kersten, 1987). This redundancy can be as high as 74% for images of faces (Kersten, 1987). The signal can be decorrelated using whitening transforms (further detailed in section 3.2.2), which produce visual filters that emulate the response properties and centre-surround receptive field structure of retinal ganglion cells (in this case, Zero-phase whitening or ZCA (Atick, 1992; Atick and Redlich, 1992; Bell and Sejnowski, 1997)).

Independent component analysis

Decorrelation techniques remove redundancy from the signal as far as second-order statistics (covariance); however, the signals in natural scenes have some non-gaussian structures which cannot be estimated using whitening alone (Simoncelli and Olshausen, 2001). These can be extracted using

techniques such as Independent Component Analysis (ICA) (Hyvärinen et al., 2009), which remove higher-order dependencies between neurons. ICA demixes a signal of combined source signals and recovers the individual sources. For natural images, it returns Gabor-like filters, which are oriented and localized, reflecting simple cell "edge detection" receptive fields (RFs) seen in V1 (Bell and Sejnowski, 1997; Simoncelli and Olshausen, 2001). Although ICA considers higher-order statistics, some dependencies remain that can only be removed by considering non-linear forms of processing (Simoncelli and Schwartz, 1999).

Sparse coding

When a neural population is a sparse code, most neurons within the population are inactive during a stimulus presentation, such that the overall activity is low. Work by Olshausen and colleagues (Olshausen et al., 1996; Olshausen and Field, 1997, 2004) suggests that finding sparse linear codes for the input space (in this case, natural scenes) creates receptive fields comparable to those found in the primary visual cortex. Preserving information but maximising sparsity results in a set of basis functions with coefficients that can explain all natural images with as much statistical independence as possible. The sparsity of the representation is maximised by setting as many basis function coefficients to zero as possible. The resulting model RFs reflect those of V1 simple cells, producing Gabor filters (these are also similar to those that can be produced through ICA (Bell and Sejnowski, 1997)).

3.1.2 Noise and efficient coding

Barlow's initial hypothesis was further developed by Atick and Redlich (1990) to account for system noise. As with redundancy, noise reduces the mutual information between the input and output, adding uncertainty to the observation of the input. Noise occurs throughout the sensory system; however, noise levels can vary and can be attributed to the inputs (sensory noise) or neurons (neural noise). Noise corrupts signal transmission, meaning that channel capacity is reduced, and decoding of signal sources is more difficult. This difference is described as the signal to noise ratio (SNR). The mutual information of the response R to the input stimulus X , is $I(R; X) = H(R) - H(R|X)$, where $H(R)$ is the entropy of the responses, and $H(R|X)$ is the entropy for the noise. In a noiseless system, the mutual information is $I(R; X) = H(R)$, where the optimal coding scheme is decorrelation. In models with input (sensory) and/or output (neural) noise, studies have demonstrated variations in the receptive fields, such as different shapes and overlap of retinal ganglion cells (Karklin and Simoncelli, 2011; Doi and Lewicki, 2014). In high noise scenarios, there is a benefit to some redundancy existing in the signal (Barlow, 2001), such as having many neurons responding to the same stimulus, as this helps remove any uncertainty in the code given the noise. Atick and Redlich (1990) demonstrated that varying the SNR leads to differences in the neural code, where neurons will minimise redundancy via whitening at no and low noise. For higher noise levels, the system increases the amplitude of correlations to increase the SNR, effectively averaging the noise, to better distinguish the signal.

Doi and Lewicki (2014) examined optimal, efficient coding regimes for the periphery and central vision under different noise levels. Their model analytically characterised several elements of sensory systems, such as neural and sensory noise, sensory degradation (due to optical blur) and a limited number of neurons. Considering these factors, the model minimised the reconstruction error between the stimulus and the neural encoding. Overall an optimal linear code performed better than whitening for all noise levels and different convergence ratios (representing 1:1 for fovea and 16:1 for the periphery). However, the advantage for the optimal code was higher for the fovea than periphery and under higher levels of noise. Adding spatial locality constraints, but not sparseness (such as sparse weights and responses), could reproduce centre-surround receptive fields similar to those found in the retina.

3.1.3 Approximations with Fisher information

Other methods to those previously presented involve maximising the mutual information directly to calculate optimal arrangements of neuron tuning curves to variations in input statistics across a region. However, due to the high dimensionality of input signals, it is difficult to calculate mutual information analytically beyond simple problems (Wei and Stocker, 2016), as all possible sensory stimuli observations would need to be considered for each neuron (Yarrow et al., 2012). Fisher information can be used to calculate a lower bound on the mutual information; therefore, it can be used as a proxy of the full calculation (Brunel and Nadal, 1998)¹. Fisher information has been used most extensively to study the optimal arrangement of neuron tuning curves over spaces with non-uniform input probabilities (Ganguli and Simoncelli, 2010, 2014, 2016; Yerxa et al., 2020). Instead of assuming that even-sized receptive fields uniformly tile the input space (homogeneous populations), the optimal population should expand the number of receptive fields over high probability density regions, at the detriment to low probability. This has been demonstrated under resource constraints, including a cap on the total number of neurons or total spiking (Ganguli and Simoncelli, 2010, 2014). Representing the input space involves allocating many narrower tuning curves over the most probable part of the input distribution leading to warping of the homogeneous population. Under this code, stimuli that occur more often are allocated more of the processing resources. This also has implications for discriminability. The parts of the region where stimuli more commonly occur have lower discrimination thresholds (discriminability is inversely proportional to stimulus probabilities). Evaluations of this theory on different acoustic (acoustic and modulation frequencies) and visual attributes (orientation, spatial frequency and retinal speed) demonstrated that model predictions fitted well to empirical data of tuning widths, preferred stimuli and discrimination thresholds (Ganguli and Simoncelli, 2016). Recent work has extended these ideas for multiple stimulus features (Yerxa et al., 2020). Neural regions such as the primary visual cortex (V1) encode multiple input dimensions, such as stimuli orientations and spatial frequencies. Here, probability densities can vary along the different dimensions. Similar to work in one-dimensional inputs, a

¹Note that some studies have since suggested that this may be only true under certain assumptions, such as of the noise, and in others, it may be calculating an upper bound instead (Yarrow et al., 2012).

2D heterogeneous population of neurons can be warped such that the neural density is greatest for those regions of multiple dimensional input space where the joint probabilities are highest. Conversely, the neuronal density is lower where input probabilities are lower for both dimensions.

To summarise, several different approaches to efficient coding could be used to investigate sensory processing. Most works have considered this in visual sensing, given the availability of natural images. These methods have reproduced features of the visual processing pathway, such as receptive field characteristics at each stage, beginning with simple decorrelation models. In addition, techniques such as Fisher Information have been used to understand the allocation of receptive fields under varying region input statistics. Next, we will discuss considerations of the type and width of the bottleneck- which determines the amount of resources available for encoding information.

3.1.4 Bottlenecks

A further factor in the optimal reconstruction of the input space is the capacity of neural resources available to process sensory information. The brain has a certain limit on the amount of information that can be represented, which means it may not be possible to represent every input receptor in a 1:1 mapping. This relates to efficient coding theories, whereby the brain may optimise its resources to encode the most informative stimuli. In processing the incoming sensory signals, the brain encounters limitations on processing capacity through both the number of available neurons and metabolic constraints (Ames, 2000; Clarke and Sokoloff, 1994). For example, in vision, the optic nerves form a bottleneck between the retina ganglion cells and lateral geniculate nucleus (LGN), reducing the information flow travelling onto the primary visual cortex (V1) threefold (Zeharia et al., 2015). However, once reaching the cortex, the number of neurons increases again, and the representation becomes over-complete; in V1, there are approximately 1000 times more neurons than the optic nerves (Barlow, 1980; Graham and Field, 2010). The topography seen in primary sensory regions may also be a result of neural constraints (see Wilson and Bednar (2015) for a review of this). Placing related regions in close proximity reduces the required wiring costs and distances for related information to travel. If maps are topological due to the wiring length requirement, topology places a further constraint on the organisation of representations.

Under an informational bottleneck, the brain should seek to maximize the information in the channel, performing data compression with minimal loss. Representing and transmitting information under bottlenecks has had renewed interest, especially in deep learning (Lindsey et al., 2019; Shwartz-Ziv and Tishby, 2017; Tishby and Zaslavsky, 2015). The information bottleneck (IB) principle approach of Tishby (Tishby et al., 2000; Tishby and Zaslavsky, 2015) emphasizes the relevance of information in a target variable. This approach aims to find a representation that preserves information in the output variable Y , while compressing information available in the input, X . Here, information in X that is useful for predicting a target variable Y is found and considered the optimal representation. This conceptualization differs from those previously described, as it focuses on meaning and relevance rather than transmitting or reproducing the original signal ac-

curately. Deep neural nets, which have shown similarities to visual system processing on specific tasks (Yamins et al., 2014) may be of use for identifying relevance under IB theory, nevertheless whether deep-nets show information maximization under this principle has been debated (Saxe et al., 2019). Previous studies employing the IB method have also mostly focused on the temporal aspects of input representation (Palmer et al., 2015).

In the present study, we will focus on accurate *spatial* reconstruction, which is important for tactile processing. Here we consider resource bottlenecks, where the amount of resources (neurons) to represent stimuli are reduced, such as in the optic nerve. Resource bottlenecks have been investigated in convolutional neural nets (CNNs) by Lindsey et al. (2019). This work unified previous research on both retinal and cortical V1 efficient coding. By limiting the output neurons at the retina, both centre-surround and oriented RFs emerged. Furthermore, the exact amount of resources allocated to the downstream processing (depth of the network, number of layers) affected the retinal representations, with shallow networks leading to the retina extracting relevant features (retina performs nonlinear computations). On the other hand, if the network is deep, the retina should encode as much information about the stimulus as possible to be further processed by downstream regions (linear computations). Lindsey et al. (2019) suggested that RFs of different species may be explained by the size of the bottleneck and the amount of resources provided for downstream processing, which emphasises that constraints placed by the architecture of the visual system shape representations. Nonlinear retinal computations are performed by those species with the tightest bottlenecks, whereas larger bottlenecks have similar responses to a linear encoding.

Compared to the statistics of the input space, the redundancy in densely arranged input receptors has been less well characterised. Signals from densely packed regions, such as the fovea, are likely to show greater correlation than those that are more sparsely distributed. Most earlier work has assumed density of input receptors as constant (Atick, 1992; Doi et al., 2012). One study compared encoding at different retinal eccentricities, suggesting different coding methods depending on the retina or the periphery. As previously described, Doi and Lewicki (2014) investigated optimal sensory codes under different constraints comparing the fovea and periphery. The optimal, efficient representation was calculated for bottlenecks at the known sizes for vision (based on the convergence of cones to RGCs at different eccentricities), 1:1 for the fovea and 16:1 for the periphery. Although these optimal code reconstruction errors were investigated for different convergence ratios, this study did not compare the trade-off in representing high and low-density populations of receptors under a shared set of limited resources.

3.1.5 Aims of this chapter

What might be the expected representation of an input space in early sensory processing regions, given differences in the density of receptors and statistics of activation? This chapter aims to provide a theoretical account of the influences of these factors on representation and magnification given the resource constraint of a sensory bottleneck. We investigate this through the efficient cod-

ing framework, utilising the method of redundancy reduction via decorrelation, which was discussed in section 3.1.1. This method was chosen as it provides a simple formulation for maximising information that can be solved analytically. Furthermore, the contribution of the density of receptors has not been extensively studied in efficient coding models.

We calculate the analytic solution for resource allocations in 1D and 2D for two model input regions (and an extension to multiple regions) under varying bottleneck widths. We compare varying density and input statistic ratios and how these impact the allocation. Furthermore, we provide several predictions for plasticity, considering how allocation changes under input statistic manipulations. Whilst we focus here on the theoretical model, in the following chapter (Ch.4) we apply this method to investigate efficient coding applied to the specialised somatosensory system of the star-nosed mole.

This work was completed in collaboration with Dr A. Jimenez-Rodriguez, who contributed the derivations of the eigenvalues for the analytical solution, particularly section A.3 in the appendix. Other work, including the solution for allocation in 2D, multiple region allocations, numeric validation, all analysis and interpretation of allocations under the manipulation of density and statistics ratio, was completed by the author, L. Edmondson.

3.2 Methods

3.2.1 Problem overview

This chapter examines the effects of both non-uniform receptor densities and receptor activation on representations when there is a capacity restriction. The full set-up of the model is shown in Fig. 3.1, where an example comparison between two digits, the index and the middle, is made. First, on the sensory input sheet, receptors can be non-uniform in their densities. In panel *A*, the densities over the digits vary and are greater on the digit tips than the medial and proximal phalanges. Second, the input statistics can also change, such that one digit may be contacted more than another. For example, the index digit may have an overall higher activation compared to the middle digit (see darker blue shading in panel *A*). This results in differences in the receptor activation covariances between the two regions (panel *B*). Therefore the index digit will have a higher covariance than the middle, modelled as a scaling of the covariance function. Here we assume that the covariance decays exponentially, such that receptors that are closer together have higher covariance than those further apart. When receptors are more densely tiled, they co-vary more with their neighbours (see higher density orange points below curves versus lower density yellow in panel *B*).

There may be a limited amount of resources in sensory systems to represent each peripheral input region. Therefore, we ask how should these resources be allocated given the varying receptor density and statistics between regions? Figure 3.1C demonstrates a hypothetical allocation for our example. In this case, the index digit may have greater resources assigned given its larger activation. Furthermore, given the higher density of receptors, the digit tips may also have a larger

representation than both the medial and proximal phalanges.

3.2.2 Combined whitening and dimensionality reduction

We perform decorrelation using a second-order statistical linear model. We assume that the sensory receptors are arranged on either a one or two-dimensional sensory sheet. Here we detail the analytical solution for 2D; however, 1D follows the same process, and the equations are included in appendix A.

For a region with a total of n input receptors and output space of size m , a 1:1 mapping would occur when $m = n$. In this case, the allocation would simply be proportional to the number of input receptors in each region. However, when a limit is placed on the total amount of information, we restrict the total number of outputs to be less than the number of input receptors, $m < n$, such that there is a bottleneck. Here, we are interested in allocations under these bottleneck cases. We next show how to calculate the typical full solution for the allocation of output units to a single region, and how to determine the RFs of each of these units under desired constraints. We then show how our problem only depends on the eigenvalues of this solution, leading to a simpler analytic specification that can solve allocation to multiple regions (section 3.2.3) under a set of limited shared resources.

We represent the inputs from the region as a matrix \mathbf{X} of dimensions $n \times z$, where z represents the number of sensory input patterns. The aim is therefore to find an $m \times n$ dimensional matrix \mathbf{W} . Decorrelation is performed by multiplying the inputs by the whitening matrix, \mathbf{W} , such that $\mathbf{W}\mathbf{X}$ is uncorrelated:

$$\mathbf{X}^T \mathbf{W}^T \mathbf{W} \mathbf{X} = \mathbf{I}. \quad (3.1)$$

Here we can set $\mathbf{W} = \mathbf{\Sigma}^{-\frac{1}{2}}$, where $\mathbf{\Sigma} = \mathbf{X}^T \mathbf{X}$. Using this we can express the solution in terms of the eigenvectors, $\mathbf{\Phi}$, and the eigenvalues $\mathbf{\Lambda}$ (as a diagonal matrix) of the covariance $\mathbf{\Sigma}$:

$$\mathbf{W} = \mathbf{P} \mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{\Phi}^T. \quad (3.2)$$

The matrix \mathbf{W} can be considered a set of non-unique whitening filters, which depend on the choice of orthogonal matrix \mathbf{P} . The receptive fields of each unit become the principal components in the case that $\mathbf{P} = \mathbf{I}$. However, principal components can span the whole surface and are not local. To obtain some spatial constraint to the receptive fields, ZCA (Zero-phase component analysis) can be performed, where \mathbf{P} is instead set to the eigenvectors, $\mathbf{\Phi}$, of the input covariance $\mathbf{\Sigma}$.

We aim to consider cases where there is a bottleneck ($m < n$), such that not all information can be retained. In order to calculate optimal allocation under this constraint, an Orthogonal Procrustes problem can be solved (Doi and Lewicki, 2014) to find the m -dimensional matrix \mathbf{P}^* (where m is

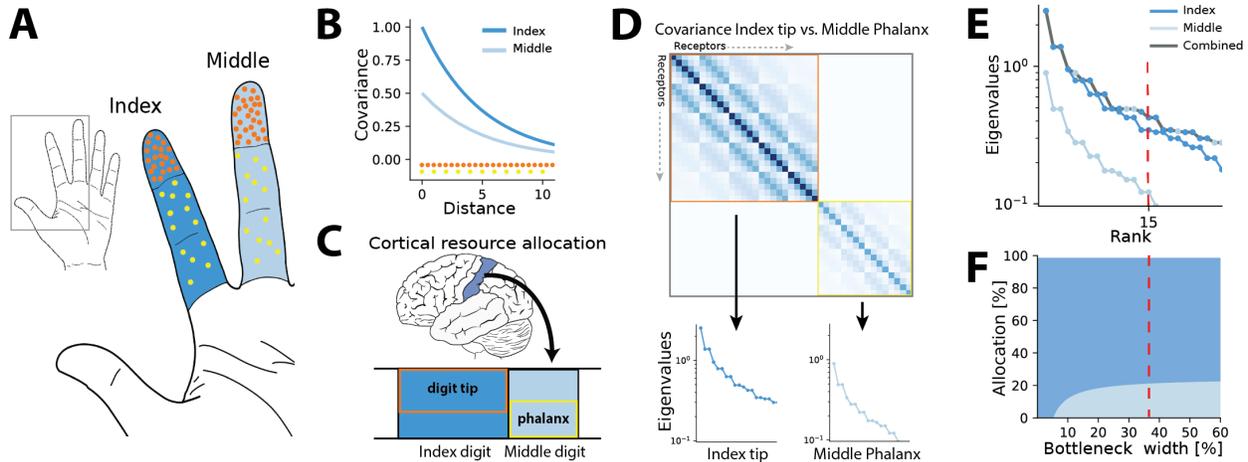


Figure 3.1: The problem of resource allocation. **A.** How should resources be allocated given differences in receptor densities and activation? Here we show an example for the hand, between the index and middle digits. Across each digit, the densities of receptors differ. Here the digit tips have a higher density (orange points) than the medial and proximal phalanges (yellow points). The digits also have different usage, determined by the typical contact statistics with objects in the environment. In this example, there is higher activation on the index finger (darker blue colour) than the middle digit (lighter blue colour). **B.** The co-activation of receptors within a region is governed by their covariance function. The covariance's decay speed determines how much receptors covary given their distance. Below each curve demonstrates two different densities, higher density in orange and lower density in yellow. When points are more densely packed, direct neighbours covary more. The covariance can also be scaled, denoting if a region is more active than another. Here the covariance for the index digit is double that of the middle digit (darker blue line versus lighter blue lower). **C.** Given these two factors, how should a limited set of resources be allocated between the two regions? Here allocations may not be proportionally allocated to each input region. Instead, certain factors may influence which region has a greater allocation of neural resources. In this example, the index digit has a greater allocation influenced by higher activation, whereas the middle digit has a lower proportion of the resources. Within the index finger representation, the digit tip is more densely innervated, leading to a larger representation of the digit tip than the phalanges. **D.** We consider the covariance of receptors in each region as a block matrix. Each value represents the covariance of one receptor to another. Along the diagonals, the covariance is greater as these signify covariances between the neighbouring receptors. Note that this block covariance is for two square 2D regions, hence the pattern of covariances. Two regions are compared, the index digit tip, which has a greater density (hence the larger block of the block matrix) and also higher activation (leading to darker blue colours in the covariance). The two regions are denoted by the orange and yellow boxes for index tip and middle phalanx, respectively. Between the regions, the blocks on the off-diagonal are equal to zero. Given the restriction of the limited resources, we aim to find the most optimal allocation of the regions. Here we utilise efficient coding ideas of decorrelation, which is calculated through eigendecomposition of the covariance matrix. Each resulting eigenvalue indicates which region an output neuron's receptive field will be placed in; for example, eigenvalues from the index region indicate that an output neuron will be allocated to that region. Arrows from the block covariance show corresponding eigenvalues for each region. Eigenvalues for the index tip (higher density and activation) have a larger magnitude for the initial eigenvalues than the lower density, lower activation middle phalanx. **E.** When the covariance matrix is a block matrix, its eigenvalues are also a block matrix. This means that eigenvalues can be independently calculated and then sorted by their magnitude into a combined set (shown by the grey line). We can select a subset of these from their rank order to create a bottleneck or limited resources. The top eigenvalues are selected from this combined set, up to a limit determined by the bottleneck width (here shown by the red dotted line). **F.** At each bottleneck, the proportional allocation between the regions can be calculated. Here, the largest eigenvalues were mostly from the index region; therefore, its allocation is greater at the corresponding bottleneck (see red dashed line).

the size of the bottleneck). \mathbf{P}^* minimises the error of reconstruction between the input set and a set of local receptive fields (whitening filters), \mathbf{W}_{opt} . In Doi and Lewicki (2014) whitening filters were set to be uniformly tiled Gaussian bumps, where their width is dependent on the dimension of the bottleneck. The matrix \mathbf{P}^* is found by:

$$\mathbf{P}^* = \min_{\mathbf{P}} \left\| \mathbf{W}_{opt} - \mathbf{P}\mathbf{\Lambda}^{-\frac{1}{2}}\mathbf{\Phi}^T \right\|_F^2, \quad (3.3)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. We retain the m largest eigenvalues $\mathbf{\Lambda}$ and corresponding components (eigenvectors) $\mathbf{\Phi}$, up to to the size of the bottleneck. The precise RF structure is partly dependent on the choice of \mathbf{P} . Considering only a single region, as in the above cost function, the number of eigenvalues determines the total number of output units, and therefore the total number of receptive fields. Each of these eigenvalues can therefore be thought of as corresponding to a neuron that has an RF over the input space. To add additional constraints on the precise RF structure of each of these neurons (such as their shape and local location within the region), we would need to fully solve Eq.3.3 which considers these additional constraints through the specification of \mathbf{P} . If the bottleneck is larger (larger m), more eigenvalues are retained, corresponding to more neurons with RFs over that region. In the case of a single region, it is clear that all resources would be assigned to that region. Therefore, previous works using a single input region were predominately focused on the RF structures of the neurons. In this study, we consider how a single set of resources (units) is allocated between multiple regions. For example, a single pool of neurons where each would have a corresponding RF falling on one of the regions. Therefore, we focus on the global location (which region) of each unit's RF, rather than the RF structures of each of these allocated units. In this case, we do not need to solve for the precise structure, and our solution is only dependent on the eigenvalues of the covariance matrix $\mathbf{\Sigma}$ of all the possible regions. When a bottleneck is imposed, we select the m largest of these eigenvalues, and our problem is now determining which region the eigenvalues originated from. This process and the analytical solution is described fully in the following section.

3.2.3 Whitening with multiple input regions

The above solution can be extended to determine where an output neuron would place its receptive field between multiple different input regions. Here we will calculate the allocations for two input regions; however, the solution for any number of regions is presented in the appendix A.

As our problem depends on the covariance matrix $\mathbf{\Sigma}$, we will first discuss its structure. Each region can differ in their receptor densities and/or activation, which impacts the correlation and decay of the covariance matrix. Here we assume that the covariances decay exponentially, which allows us to calculate an analytic solution to the allocation. We model the two regions as independent of each other, such that the covariance between receptors in different regions is zero. This approximation means that the covariance is a block matrix. The blocks along the diagonal are the covariances

of each region, and those on the off-diagonal are zero. When regions are relatively large and correlations decay to zero over relatively short distances, it is suitable to assume that the two regions are independent. In this case, the correlations will not extend far enough to enter neighbouring connected regions. For further discussion of this and comparison between the block and non-block covariances, see Edmondson et al. (2019). Furthermore, we assume square regions for our main derivation of two regions; however, the result is similar for rectangular regions and only diverges for extremely long narrow regions. Regions in sensory systems tend to be separated, for example, the digits are separated, and borders between regions are narrow. For the digits between the tip and medial phalanx, there is only one border, and this is along the short side of the phalanx. The only case where the assumption may not be valid is for very small regions which are highly connected and with long-ranging correlations, such that the neighbouring region would also be highly correlated when a region is contacted.

Considering again the example of two regions in Fig. 3.1, the block matrix for a comparison between the index tip and a single phalanx of the middle digit is shown in panel *D*, top. The covariance block for the index tip (orange) is larger as there are more receptors in this region. Furthermore, in our example, this region is also more active (see scaled covariance in panel *B*); thus, the covariances are larger for neighbouring receptors. In the figure, darker colours indicate larger covariances.

For two input regions the block diagonal covariance is denoted as:

$$\Sigma = \begin{bmatrix} \Sigma^{(R1)} & \mathbf{0} \\ \mathbf{0} & \Sigma^{(R2)} \end{bmatrix} \quad (3.4)$$

Our method does not fully solve for the precise location or sizes of the receptive fields within a region (full whitening²), instead we ask in which of the regions each neuron's receptive fields will fall. In the previous section, we demonstrated that we only need to know the eigenvalues of the covariance matrix to calculate this. When the covariance is a block diagonal matrix, the eigenvalues and eigenvectors are also of block diagonal form. By Cauchy Interlacing theorem, this means that the same solution can be obtained by calculating the eigenvalues and eigenvectors of each region alone (as they are independent), as can be obtained by calculating them together. This approximation enables an analytic solution to be calculated, as shown in the next section. The following eigenvalues and eigenvectors can then be written as:

$$\Lambda = \begin{bmatrix} \Lambda^{(R1)} & \mathbf{0} \\ \mathbf{0} & \Lambda^{(R2)} \end{bmatrix} \quad \text{and} \quad \Phi = \begin{bmatrix} \Phi^{(R1)} & \mathbf{0} \\ \mathbf{0} & \Phi^{(R2)} \end{bmatrix} \quad (3.5)$$

If an eigenvalue is from $\Lambda^{(R1)}$, this would suggest that the neuron should place its receptive field over R_1 .

²Where \mathbf{P}^* is found, as discussed in section 3.2.3.

Fig. 3.1D demonstrates the corresponding eigenvalues for each diagonal block (each region) of the covariance matrix. The combined set of eigenvalues from both regions is demonstrated in Fig. 3.1E, where the ordering is shown in dark grey.

If the total number of output neurons available to represent the inputs is reduced, a bottleneck is imposed on the inputs. Here we implemented limited resources by restricting the width of the bottleneck. Similar to dimensionality reduction techniques, we retain only the largest m eigenvalues from the full combined set of both regions, where m is the size of the bottleneck. To calculate the final allocation, we, therefore, count the number of eigenvalues from each region with rank lower than the bottleneck size. In Fig. 3.1E we show an example imposed bottleneck in red. The corresponding allocations for each region over increasing bottleneck widths are shown in Fig. 3.1F. The allocation is calculated as the percentage of output units (eigenvalues) associated with each region. At the example bottleneck (shown in red), there were more eigenvalues of higher magnitude in the index digit region (darker blue) than the middle digit (lighter blue).

In the next section, we use our assumption of the exponential covariance decay and a block matrix structure to demonstrate that a direct analytic calculation of the allocation for two regions can be found.

3.2.4 Analytic calculation of eigenvalues

We assume that the covariance between receptors has an exponential decay, allowing an analytic solution. The covariance matrix can then be calculated based on the distance between receptors (see Fig. 3.1C). Below we outline the allocation steps; the complete derivations can be found in appendix A.

We first demonstrate the calculation of a baseline region, R_1 . The covariance is:

$$\Sigma_{ij}^{(R1)} = e^{-\sigma|x_i-x_j|}, \quad (3.6)$$

where x_i and x_j are the coordinate locations of i th and j th receptors, and σ is the decay constant, or the inverse lengthscale. A demonstration of the effect of changing σ on the covariance decay can be seen in Fig. 3.2.

Eigenvalues can be calculated analytically by expressing the problem in the continuous domain. When expressed as an integral homogenous equation, we can calculate the eigenvalue-eigenvector problem for R_1 :

$$\lambda_k \phi_k(x) = \int_0^L e^{-\sigma|x-y|} \phi_k(y) dy, \quad (3.7)$$

where the problem is expressed on a bounded domain of length L , which corresponds to the region

size for one of the dimensions of the input. The k th eigenfunction is denoted by $\phi_k(x)$, and the corresponding eigenvalues λ_k .

The solutions of this problem are related to the Laplacian operator (see appendix A), such that:

$$\lambda_k = \frac{2\sigma}{\mu_k + \sigma^2}, \quad (3.8)$$

where μ_k are the eigenvalues of the Laplacian operator.

Using Dirichlet boundary conditions, where the problem is bounded at both sides for each dimension, the general solution for the Laplacian eigenvalue problem is (Strauss, 2007):

$$\mu_k = \mu_{l,m} = \frac{l^2\pi^2}{L_1^2} + \frac{m^2\pi^2}{L_2^2}, l, m = 1, 2, .. \quad (3.9)$$

where L_1 and L_2 are the domain lengths for each dimension of the rectangle.

For a second region, we consider that the receptor densities and activation may vary from the first region. The covariance for R_2 is therefore notated with considerations that the activation ratio a , and ratio of receptor densities by d may differ:

$$\Sigma_{ij}^{(R2)} = ae^{-d\sigma|x_i-x_j|}. \quad (3.10)$$

The effect of the activation ratio, a is a scaling of the covariance matrix (also shown in Fig. 3.1B). The density ratio d alters the spatial extent of the covariance of the placed receptors, therefore implicitly accounting for the different receptor densities.

For the eigenvalues R_2 , the calculation is as before, but now considering the receptor density as an implicit stretching of the space (density ratio d), the region length L is adjusted:

$$\lambda_k\phi_k(x) = \int_0^{L\sqrt{d}} ae^{-d\sigma|x-y|}\phi_k(y)dy \quad (3.11)$$

For a square region, the total density is d , calculated as $(\sqrt{d})^2$, with each axis of as \sqrt{d} . As the density ratio is given for 2D, for each dimension the size of the space will be stretched by \sqrt{d} .

3.2.5 Allocation in the bottleneck

We now demonstrate how the above solution relates to the bottleneck. Given some limited capacity for processing, how should the resources be allocated amongst the two regions? Here we consider the allocations for two regions, following from above, however the allocation for multiple regions will be given in the appendix, A.5.

We assume square regions in 2D. As both lengths are the same, we use a single variable L . Following (3.8), the eigenvalues for each region R_1 and R_2 are calculated as:

$$R_1 : \lambda_{l,m}^{(R1)} = \frac{2\sigma}{\pi^2 L^{-2}(l^2 + m^2) + \sigma^2} \quad (3.12)$$

$$R_2 : \lambda_{n,o}^{(R2)} = \frac{2\sigma a\sqrt{d}}{\pi^2 L^{-2}(n^2 + o^2) + \sigma^2} \quad (3.13)$$

where l, m and $n, o \in \mathbb{N}$ enumerate eigenvalues for region R_1 and R_2 , respectively.

When the bottleneck width is varied, the allocation of neurons to the regions depends on the ordering of the eigenvalues from each region. Therefore, we must determine the sorted eigenvalue ranks for each pair (l, m) . Unlike the 1D case, there is no natural ordering of the eigenvalues for two dimensions. Instead, the ordering is approximated by counting the number of lattice points enclosed into the quarter circle with a radius of $p = l^2 + m^2$ (see appendix A.4 for full details). We call this function $N(p)$ and set $p^{(R1)} = l^2 + m^2$ and $p^{(R2)} = n^2 + o^2$. The number of neurons allocated to R_1 is therefore defined as a function of the number allocated to R_2 . Setting $\lambda^{(R1)} = \lambda^{(R2)}$ and solving for $p^{(R2)}$ yields:

$$p^{(R2)} = a\sqrt{d}p^{(R1)} + \frac{L^2\sigma^2 a\sqrt{d} - L^2\sigma^2}{\pi^2}. \quad (3.14)$$

The convergence point for this function depends on the activation and density ratios. As we allocate more neurons to region R_1 , the ratio $\frac{N(p^{(R1)})}{N(p^{(R2)})}$ between regions simplifies to $\lim_{R_1 \rightarrow \infty} \frac{N(p^{(R1)})}{N(p^{(R2)})} = a\sqrt{d}$. The fraction allocated to each region therefore converges to $\frac{1}{1+a\sqrt{d}}$ and $\frac{a\sqrt{d}}{1+a\sqrt{d}}$ for R_1 and R_2 respectively.

In the following sections, we present resulting allocations for two-dimensional sheets, the latter of which is more relevant for real sensory systems such as visual and tactile processing. The results for 1D regions are presented in appendix A. Using the solution presented above, we calculate efficient allocations when varying both the density and activation ratios. The bottleneck restriction is based on the number of output neurons available. However, see appendix A for a demonstration where the bottleneck is defined as the amount of information being transmitted. As in the example (Fig. 3.1), we solve the problem with two sensory regions assuming they are equal in size, but the receptor densities and activation can vary between them. For more regions, we provide the solution for calculating the eigenvalues for each region. A further example of solving multiple regions can be found in the following chapter 4.

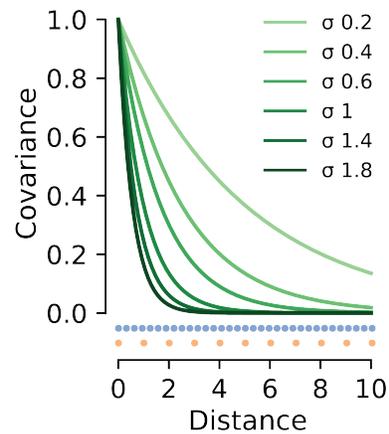


Figure 3.2: Covariance decay with varying values of σ . Green curves denote a change of the σ parameter in the exponential equation. Darker colours indicate larger values. Varying densities are shown underneath the curves, with the orange points denoting the spacing of points in 1D. Blue points show increased density of the second region, here 1:3. When σ is smaller, the decay is slower over the region, meaning receptors are correlated at wider distances. Increasing the value of σ leads to smaller covariance with distance, such that neighbouring receptors have less similar responses.

3.3 Results

We first demonstrate how allocations change for two regions depending on their density ratio (heterogeneous density) or activation ratio (heterogeneous activation) alone, while the other factor is held constant. Following this, we show the combined effect for both density and activation. Finally, we investigate how bottleneck widths may affect the extent of plasticity.

3.3.1 Heterogeneous density

First, we examined the effect of manipulating the receptor density ratio whilst holding the activation constant. Our key finding is that the allocation between the two regions is dependent on the width of the bottleneck. For different density ratios, allocations are not proportional between the two regions. Instead, the higher density region could both expand and contract its allocation relative to the proportional input density mapping. In general, for narrower bottlenecks (typically smaller than 10% of the total input neurons for both regions), the higher density region is expanded (over-represented compared to the proportional density). This expansion can be an exclusive representation of only the higher density region or a greater allocation than proportional. This can be seen for the example in Fig. 3.3A, where the blue (higher density region) is expanded for small bottlenecks. Here expansion is any blue area below the red dashed line, which indicates an allocation proportional to the density.

What is the cause of this expansion? When the density ratio is increased, there is a multiplicative scaling of the eigenvalues for the higher density region, meaning that there are more eigenvalues from the higher density region to be allocated at narrower bottlenecks than the lower density region. An example of this scaling in a simplified 1D example can be found in Fig. 3.5A.

For intermediate-sized bottlenecks, the lower density region is expanded (or alternatively higher density contraction as shown in Fig. 3.3A). The negative exponential form of the covariance function leads the allocation to converge to a fixed ratio of allocation between the regions when approaching wider bottlenecks (see yellow dashed line in Fig. 3.3A). Tendency to converge to the limit is seen for any heterogeneous density ratio (see Fig. 3.3B). This convergence occurs towards wider bottleneck widths, according to $1/(1+\sqrt{d})$, where d is the density ratio.

As the bottleneck further increases, all information arising from lower density regions has been captured (all eigenvalues from the lower density region have been allocated). Therefore, the rest of the bottleneck is only allocated to the higher density region. For the example in Fig. 3.3A, this leads to a slow increase in the higher density region's allocation. When there is no longer a bottleneck, the allocation is proportional to the density. This effect of allocation changes emerges for any heterogeneous density ratio between the two regions. Figure 3.3B demonstrates the effect of increasing density ratios, each showing a non-linear allocation shape.

Spatial correlations over the regions, which are controlled by the decay parameter σ , affect allocations for narrower bottlenecks. Figure 3.3C demonstrates differences in allocation at narrow and

intermediate bottleneck widths. The σ change can cause a faster or slower approach to the convergence but does not influence where the allocation becomes a fixed proportion (the convergence point). For increasing σ , spatial correlations decrease, meaning that receptor responses are more independent from each other. Only neighbouring receptors may be responsive together, and more distant ones are not correlated. This leads to the convergence limit only being approached at wider bottlenecks for larger σ . Therefore, the correlation structure has a large influence on the initial extent of expansion of the region.

3.3.2 Heterogeneous statistics

Above we focused on how changing the receptor density ratio alone could affect allocations. Here, we demonstrate the effect for changing the activation of the receptors between the regions whilst holding the density ratio constant (both regions have an equal density of receptors, see inset in Fig. 3.3D). Differences in stimulus statistics lead to change in the activation levels of receptors and their response variances. Here a proportional allocation to density is exactly 50% between both regions. This is shown in Fig. 3.3D by the purple dashed line.

Similar to densities, heterogeneous receptor activation also demonstrate effects subject to the width of the bottleneck. However, there are some notable differences. First, representations of the higher activation region are expanded for the whole bottleneck (see example in Fig. 3.3D), with the extent of the expansion determined by the bottleneck width. For example, as with heterogeneous density, the higher activation region can be exclusively represented at narrow bottlenecks. This effect is more extreme for heterogeneous activation. As with differing density ratios, the allocations converge at a limit of $1/(1+a)$ where a is the activation ratio. This limit has a level of expansion or contraction that is more extreme in the intermediate bottlenecks (where convergence occurs) than the density case. For example, compare the difference between the yellow and purple dashed lines in both Fig. 3.3A and D. This convergence limit demonstrates a fixed allocation of both regions. Following this convergence, eigenvalues from the higher density region have either all or nearly all been assigned. After this, only eigenvalues from the lower activation region are left, meaning that the rest of the bottleneck allocation is assigned to the lower region.

As with heterogeneous densities, the non-linear effect for changing the statistics ratio is also due to a scaling of the eigenvalues, see Fig. 3.5B. In both cases, eigenvalues are scaled by the same amount. However, when the density is equal between the regions, the same number of eigenvalues are assigned for both regions. Here all the eigenvalues for the higher activation are assigned before those for the lower activation region, meaning that the allocations always approach the proportional density from below (curves are all below proportional allocation for the entire bottleneck in Fig. 3.3E), and there is no over-representation of the lower activation region.

Comparing the extent of spatial correlations, increasing σ leads to a more exclusive representation of the higher activation region, which can occur at even larger, intermediate bottlenecks. Here it is also possible that the allocations do not meet the convergence, with all eigenvalues allocated to

the higher density region before this is reached. Figure 3.3F demonstrates changes to σ value.

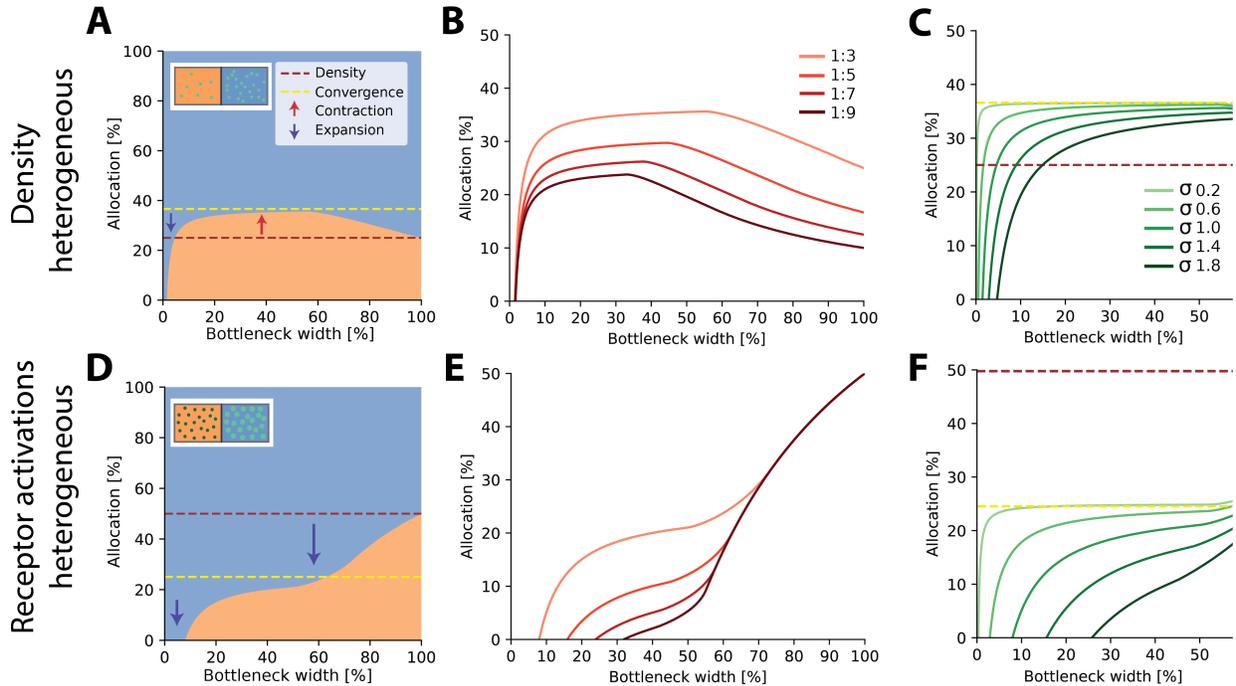


Figure 3.3: Optimal resource allocation for heterogeneous receptor densities or input statistics leads to complex trade-offs. **A.** Illustration of resource allocation for heterogeneous receptor density but homogeneous stimulus statistics over all bottleneck sizes. Orange region denotes the lower density and blue, higher density, with a ratio of 1:3. Dotted lines show proportional representation according to receptor numbers (purple) and convergence of the optimal allocation in the limit (yellow). Arrows indicate contraction (red, up) and expansion (black, down) of the higher density region representation. **B.** Bottleneck allocation boundaries for different density ratios (given as low:high). The area below each line corresponds to the low-density representation, while the area above corresponds to the high density one, as in panel A. **C.** Effect of changing the extent of the spatial correlations (parameterised by the decay value σ , see Methods for details and Fig. 3.2 for an illustration of the covariance function for different values of σ). Density ratio is set at 1:3 for all σ . Increasing σ leads to expansion of the higher density region for a larger initial portion of the bottleneck. In all other panels σ is set to 1. **D-F.** Same as in row above but for homogeneous density and heterogeneous receptor activation. **D.** Illustrative example with the blue region having higher receptor activation. Note that the representation of the higher activation region is expanded for all bottleneck widths. **E** Allocation boundaries for different activation ratios. The representation of the high-activation regions is expanded for all bottlenecks. As the activation ratio increases, the highly active region allocation is expanded for wider bottlenecks. **F.** Changing the extent of spatial correlations (σ) has larger effects when the activation is heterogeneous (set at 1:3 for all σ) compared to heterogeneous density (panel C). See appendix Fig. S.1 for an equivalent figure considering one-dimensional receptor arrangements.

3.3.3 Interplay between stimulus statistics and receptor density

When stimulus statistics (receptor activation ratios) and receptor density differences are considered separately, representations of input regions can expand or contract depending on the width of the bottleneck. Some similar effects are seen between the two factors, with narrow bottlenecks being a region where the spatial extent of correlations leads to the largest effects on the allocations. However, in sensory systems, both density and input statistics vary across regions and influence the allocations. For example, in touch, alongside density differences across the hand, we also contact and manipulate objects in varying ways, leading to different activation patterns for each digit.

Similar to considering the factors separately, there is an expansion of narrow bottlenecks; however, this can be more extreme, especially for large activation ratio differences. A key finding is that the activation ratio is more influential in determining the allocation than the density ratio, with the extreme sole-representation of a single region occurring into even intermediate bottlenecks for highly variable activation. The corresponding change in density does not lead to large differences between initial representations. For example, Fig. 3.4A, left panel, there is a wider spread of possible allocations when the activation is varied than the corresponding density (right panel). The convergence for intermediate and wide bottlenecks is affected by both factors, and is calculated as $1/(1+a\sqrt{d})$, where a is the activation and d is the density ratio.

Possible allocation schemes across different combinations of activation and density are shown in Fig. 3.4B. This demonstrates whether region one is either solely expanded, solely contracted, or both expanded and contracted over varying widths of the bottleneck. Considering a fixed density (1:2 for region 1:region 2, shown by horizontal dashed line), then the allocations of region one can be in any of the regimes of expansion, contraction or both, and this depends on the activation ratio. However, when only information about the activation ratio is known, then more is known about the possible regime (see the vertical dashed line, for example). When activation is greater over region two, then region one will either be contracted or both expanded and contracted. The converse is true when there is greater activation on region one, with either full expansion over the bottleneck or expansion and contraction. Therefore, when only information about density is known, we cannot know which of the three regimes we will be in. Therefore, this knowledge of the activation ratio is more useful in knowing the possible representations, although not the specifics or extent of the expansion and contraction.

Finally, the two factors are likely positively correlated; for example, those regions with higher densities will also experience a corresponding higher activation. This is seen in touch, where the fingertips are most densely innervated (Vallbo and Johansson, 1984), and also have a typically greater activation (Gonzalez et al., 2014). When the higher density region's activation is greater than \sqrt{d} , the representation of this region is always expanded for the entire bottleneck. The extent of this expansion is greater for narrower bottlenecks; for example, see Fig. 3.4A, orange lines.

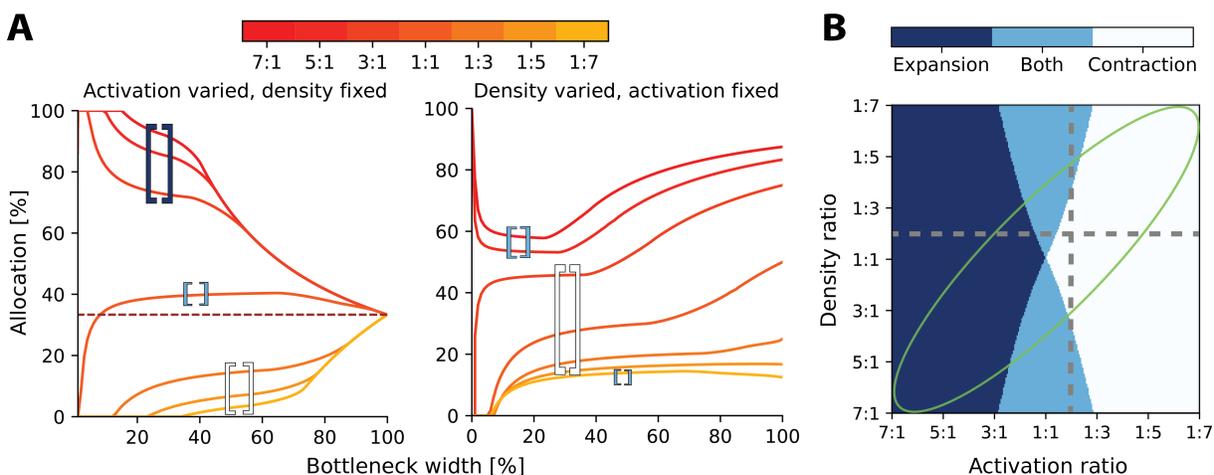


Figure 3.4: Interactions between heterogeneous statistics and density. **A.** Expansion and contraction for a region with constant density or activation (baseline), while both density and activation are varied for the other region. All ratios given are baseline:other region. Symbols: *Dark blue brackets* denotes expansion of the baseline region. The extent of brackets shows all combinations with expansion. *Light blue*, expansion and contraction of the baseline region. *White* indicates contraction of the baseline region. Left: Fixed density ratio, while activation ratio is varied. Right: Fixed activation ratio, while density ratio is varied. **B.** Possible expansion/contraction regimes for baseline region based on combinations of density and activation ratios. Grey dotted lines denote the possible allocation regimes for a region if only the activation ratio (vertical) or density (horizontal) is known. Examples shown in panel A are for these parameters. Green ellipse demonstrates ratio combinations where activation and density are correlated. See Fig. S.2 for a comparison of how receptor density and activation interact between one-dimensional and two-dimensional receptor arrangements.

Explaining allocation and eigenvalue ordering in the 1D case

The 1D solution can be used to easily investigate the effect of spatial correlations on eigenvalue decay for the two regions. Figure 3.5 demonstrates eigenvalue patterns for heterogeneous density (A) or activation alone (B), and then the effect of both combined (C). For the case of heterogeneous activation and density, the eigenvalues are scaled by density ratio, such that in the example given, the higher density or activation eigenvalues are doubled for a 1:2 ratio. This leads to the sole allocation of those eigenvalues in narrow bottlenecks. Once the rank of the combined set increases (bottleneck widens), the eigenvalues of the lower density region are included in the sorting. The sorting initially follows the same order for both factors, as long as other parameters are fixed, such as σ . Differences occur in the heterogeneous density case, as there are double the amount of eigenvalues in the higher density region, meaning that these still need to be allocated in the combined set. For combined density and activation, the effect is a double scaling for increasing both ratios to 1:2. This results in a larger initial expansion of the higher region for 1D.

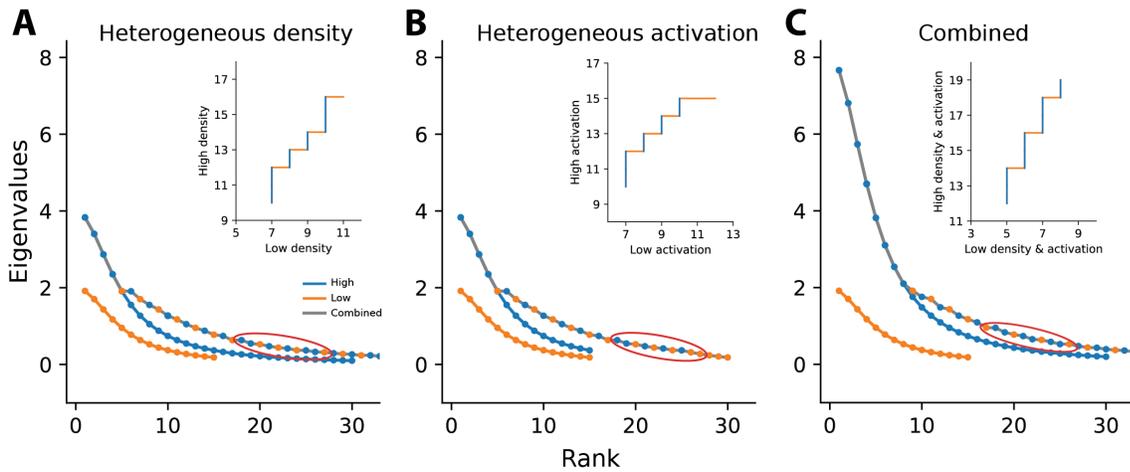


Figure 3.5: Illustration of eigenvalue sorting and resulting allocation. Main panels: Eigenvalues for heterogeneous density (A), heterogeneous activation (B), or combined (C). For each, the manipulated ratio is set as 1:2. For simplicity, the example considers 1D regions. Orange markers show eigenvalues from the baseline region, identical for all panels. Blue markers show eigenvalues from the second region with higher density, activation, or both. Markers connected by the grey line show the combined set of sorted eigenvalues from both regions. Insets: Cumulative allocation for the set of eigenvalues marked by the red ellipse in the main panel. Horizontal orange lines indicate eigenvalues from and therefore allocation to the orange region, whilst vertical blue lines indicate allocation to the blue region. **A.** Heterogeneous density. Increased receptor density in the blue region causes higher spatial correlations between neighbouring receptors and, therefore, larger eigenvalues, leading to a scaling of the eigenvalue curve. Increased density also leads to more receptors in the blue region, increasing the number of eigenvalues from this region and the total number of eigenvalues considered. **B.** Heterogeneous activation. Increased variance in the blue region leads to a scaling of the eigenvalue curve. However, the number of receptors in the blue region is not affected, explaining differences observed compared to A. **C.** Heterogeneous density and activation. Effects from A and B are combined, leading to a double-scaling of the eigenvalue curve, plus an increase in the total number of eigenvalues.

3.3.4 Resource limits determine the extent of changes under plastic manipulation

Throughout most of an organisms' life, both the available sensory processing resources and density of afferents remain relatively constant. However, the stimulus statistics can change. For example, when learning new skills or adapting to a new environment. Changes to stimulus statistics affect the incoming sensory information, which is likely to change the sensory representations. For example, many empirical studies have demonstrated changes to somatosensory cortex representations when manipulating inputs to the hand. Increasing stimulation to some areas of a digit increase its corresponding representation size in cortex (Jenkins et al., 1990). This is mediated through a variety of plasticity mechanisms which are further discussed in chapter 6.

Here, we investigate how representation allocation may change given a change in the input activation ratios and whether the resource limits of the bottleneck would lead to any differences. We first calculated the allocations given certain density and activation ratios for different bottleneck widths. Following this, we changed the receptor activation ratios (modelling a change in the inputs) and recalculated the allocations for each region. This is demonstrated in Fig. 3.6A. When the receptor activation is increased more for one region, the effect's extent depends on the bottleneck width—the largest changes are seen for smaller bottlenecks. Effects reduce with increasing bottleneck widths. We demonstrate this over a range of activation ratio changes across regions for different densities and three bottleneck sizes (narrow, intermediate, wider) in Fig. 3.6B.

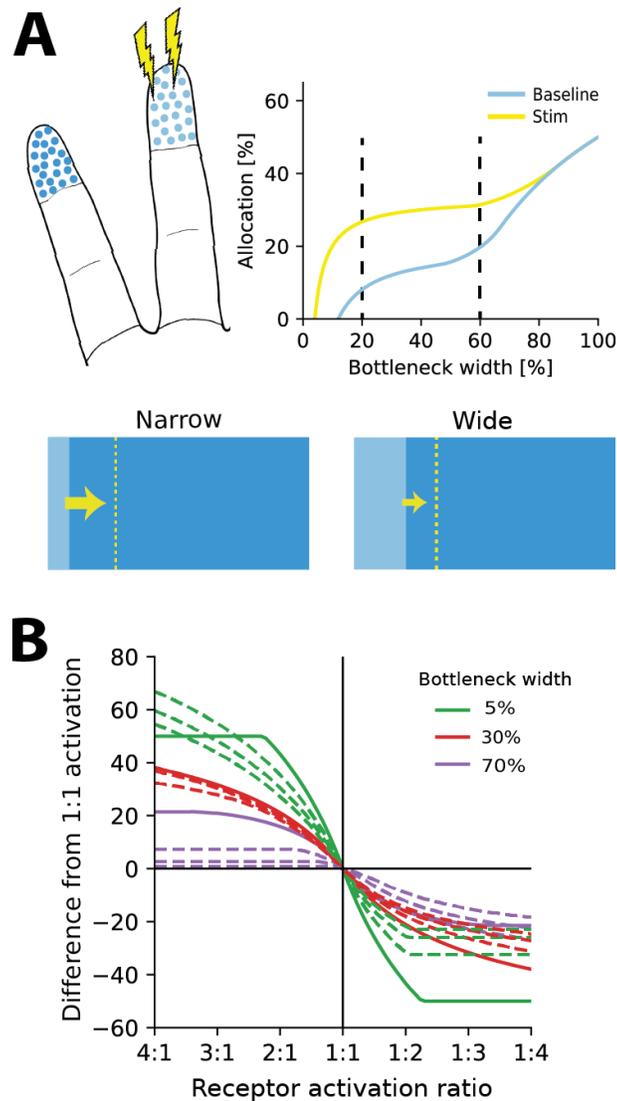


Figure 3.6: Re-allocation to account for changes in stimulus statistics. **A.** Left: Illustration of problem setup. Shown are two fingertips with fixed receptor density and known statistics (baseline condition). Increased stimulation is applied to the middle digit (yellow symbols, stimulation condition), leading to changes in optimal allocations. Right: Optimal allocations for baseline (blue) and stimulation (yellow) conditions across all bottleneck widths. Stimulation of the middle finger increases its representation, but the magnitude of the effect depends on the bottleneck width. Bottom: Changes in allocation of the middle digit for two bottleneck widths (indicated by vertical black dashed lines above). The yellow dotted line denotes an increase in allocation post-stimulation. This effect is proportionally larger for the narrow compared to the wide bottleneck. **B.** Change in allocation when receptor activation for an input region increases (left half) or decreases (right half), mimicking a change in stimulus statistics. Drastic changes in cortical allocation are observed for narrow bottlenecks (green lines), while wider bottlenecks cause more moderate change. Solid lines denote equal receptor density across both regions, while dotted lines show varying density ratios between 1:2 and 1:4.

3.4 Discussion

In this chapter, we investigated efficient coding methods of decorrelation in calculating allocations of sensory regions under limited resources, given varying receptor densities and activation ratios. We demonstrate that these factors lead to a non-linear scaling of the region representations, which do not follow an allocation proportional to the receptor densities, but rather depend on the width of the bottleneck. For example, when the bottleneck is narrow, one region is typically expanded and favoured in its representation. For intermediate to wide bottlenecks, there is a convergence towards a ratio of constant allocation between the two regions. The exact regime of expansion and contraction of a region over the bottleneck depends on both the activation and density ratios. For 2D surfaces, the activation ratio better indicates whether a region is expanded or contracted. Considering a change to the activation of a region, for example through changes in hand use, smaller bottlenecks will result in a larger change of allocation between regions, compared to a larger bottleneck.

3.4.1 Comparison with previous approaches

Most previous work investigating receptive field placement and structure has focused on infinite populations of neurons. Few have considered the effect of limiting the number of output neurons (resource bottlenecks) on the output representations or considering that the input receptors may also vary in their densities over the input space.

Using Fisher information, models have demonstrated non-uniform allocation of output neurons, with higher probability stimuli represented by a greater number of neurons. The use of Fisher Information is a proxy for mutual information, enabling the calculation of optimal output neuron density and tuning curve placement given the distribution of sensory stimuli (Ganguli and Simoncelli, 2010, 2014, 2016; Yerxa et al., 2020). However, these approaches maximise information between the output population and the stimulus distribution itself, such that allocating additional neurons to a given region of the input space will always increase the overall amount of information represented. In contrast, we only aim to identify which region the RFs will fall in, but not their precise structure of receptive fields. Furthermore, we assume a finite number of input receptors in each region, determined by their relative densities. Once information from one region has been fully captured in the bottleneck, our method does not allocate further resources to that region. Therefore, this also limits the total size of our neural population—the population size cannot exceed the total number of input receptors. In contrast, Fisher information does not consider the placement of receptors within a region, so these can potentially be infinite. Therefore, the size of the neural population can always be increased. Adding neurons to the population would lead to increasingly narrow tuning curves for a region with a higher likelihood of incoming stimuli. This implicitly assumes that there is always a high enough density of receptors for adding additional resources to be useful. In our model, we would expect that the amount of resources is inversely proportional to the width of the tuning curves, up to a limit of the spacing of the receptors. For

a region with a 1:1 mapping (where all receptor information is captured), the tuning curves would be expected to cover one receptor. Adding additional resources, for example, two output neurons representing a single input receptor, would lead to redundancy in the neural population.

Prior research on bottlenecks has mostly focused on the effect of different levels of compression on receptive field structure. Doi and Lewicki (2014) compared increasing eccentricities of the retina and predicted typical receptive fields of the retinal ganglion cells for differing convergence ratios. Work utilising deep neural network architectures have also found effects of network depth and layer sizes which affect compression (Lindsey et al., 2019).

In this work, we use eigendecomposition to reduce the redundancy as this decorrelates the inputs. Our cost-function therefore equates to minimising the mean-squared error of the reconstruction when a bottleneck is imposed. Cost functions explored in other works have considered factors such as input relevance to the system, for example, the Information bottleneck theory (Tishby et al., 2000; Tishby and Zaslavsky, 2015). This aims to find low-dimensional representations preserving specific information of the output variable, whilst compressing the input. Future research may investigate whether the effects of varying densities and activation on representations may change depending on the cost function or goals of the system.

Although it is not the focus of our investigation, it is possible systems are optimised based on behavioural relevance. For example, encoding of stimuli that is not behaviourally relevant may be a poor utilisation of coding capacity. Some evidence of this has been found in sensory systems, for example, grasshoppers tend to encode behaviourally relevant sound frequencies used for vocalisations, rather than all sounds in the environment (Machens et al., 2005). This could lead to different solutions to the resource allocation problem depending on a person’s individual goals. One branch of efficient coding theory has investigated this— predictive coding. These models consider how encoding behaviourally relevant information aids in predicting future inputs (Bialek et al., 2001). The system solves the optimising of compressing information about the past and maximising our predictability of the future- here we need to focus on the signals which are maximally informative of our environment (Palmer et al., 2015), and that are useful in decision making and action selection. Some properties of visual and auditory RFs have been captured using these methods (Singer et al., 2018).

3.4.2 Bottlenecks in tactile processing

In vision, the optic nerve forms a clear bottleneck between the retina and cortex, with many fewer fibres compared to the number of photoreceptors (Curcio and Allen, 1990). In touch there is no clear bottleneck location, however, this could be located along the pathway to cortex, such as at the level of the cuneate or gracile nucleus in the brainstem (for tactile information from the upper and lower body, respectively), or the thalamus. Somatosensory information in all of these areas is organised into topographic representations, with nearby body locations represented adjacent to each other (Xu and Wall, 1999; Florence et al., 1988; Suresh et al., 2017).

There exist few to no studies which have investigated the volumes or numbers of neurons in the somatosensory areas of these regions. Instead, volumes and neuron counts have only been assessed in wider regions of the brain. In total, the brain has approximately 86 billion neurons (Azevedo et al., 2009). Across the brain, regions differ in both their volumes and density of neurons (Collins et al., 2016), therefore brain volumes alone cannot be equated to neuron numbers. The brainstem, diencephalon and striatum combined equates to around 2-8% of the total brain volume but contains only 0.8% of the total neurons (Azevedo et al., 2009). This equates to around 690 million total neurons. Many structures within these regions make up large proportions of this, such as the basal ganglia, with approximately 415 million neurons (Karlsen and Pakkenberg, 2011). Within the residual count of neurons, a vast array of other structures need to be accounted for, including regions within the pons, midbrain, hypothalamus and thalamus. In the brainstem alone there exists tracts for both ascending and descending pathways of information from tactile mechanoreceptors, proprioceptive, motor and pain receptors to name a few, alongside those responsible for critical cardiac and respiratory functions (Purves et al., 2018). Therefore, it is possible that the amount of resources allocated to processing tactile inputs is more restricted at lower levels of the tactile pathway compared to cortex.

One study investigated changes in macaques between 12-21 years after dorsal rhizotomy (transection of dorsal roots of the spinal cord). Properties of the cuneate nucleus (CN), external cuneate nucleus (ECN), gracile nucleus (GC) and VPL of the thalamus were measured in both the deafferented (injury) and intact side (Woods et al., 2000). Thalamic VPL neuron numbers were counted in three monkeys, whilst upstream dorsal column nuclei regions were measured in seven. Stereological neuron counting (where tissue slices are stained and cells are counted) placed the intact side at approximately 277,000 neurons for the CN, ECN and GC combined. Comparatively fewer neurons are found in the VPL at approximately 168,000, which suggests a potential bottleneck between these two regions. Comparing the periphery and early brainstem areas to consider another early bottleneck requires knowledge of the absolute afferent fibre counts for the macaque, which is unknown. Cutaneous innervation densities of the hand are similar in macaques to humans, however other regions have not been measured. Furthermore, macaques lack Ruffini (type SAI) afferents found in humans (Paré et al., 2002). There are also several limitations of this study to consider. First, the numbers are calculated based on the intact side of the brain, and this assumes that the deafferentation did not have bi-lateral effects. Second, stereological counting assumes homogeneous cell densities over a region, in this case calculating total numbers based on smaller volumes of tissue which are then extrapolated³. It is known that neurons within these regions are not homogeneous, and instead, they consist of different cell types and densities (Loutit et al., 2021). For example, feedback connections exist from cortex regions to the CN via inhibitory interneurons (Aguilar et al., 2003), which may contribute to the total neuron counts. Finally, these numbers have not since been

³More modern techniques such as isotropic fractionators ('brain soup'- where tissue is homogenised and cell nuclei are recovered and counted) can count whole regions (Herculano-Houzel and Lent, 2005), but early somatosensory regions in primates have not yet been assessed with this method.

reproduced in other studies for confirmation or measured in any other non-human primate species.

A criticism of efficient coding theory is that neuron numbers increase substantially in cortex, suggesting a lack of redundancy. However, there may be differences in the coding capacity of cortex compared to upstream regions (Simoncelli, 2003). For example, cortical neurons may use lower firing rates than other regions, requiring larger neuron numbers. Furthermore, interconnections between brain regions in cortex are vast, and many neurons may be used in signalling information between brain regions. For example in somatosensory cortex, incoming information from upstream regions is mostly received by neurons in layer IV (Purves et al., 2018), whereas those in other cortical layers connect to thalamic, motor and higher cortical regions (Harding-Forrester and Feldman, 2018).

3.4.3 Implications for sensory processing

The densities of afferents in the hand have been extensively characterised (Vallbo and Johansson, 1984), however, due to the complexity of data collection on behavioural use, there is little empirical data available on typical distributions of hand contact and input statistics. It is therefore difficult to directly test the model in most species (however, see chapter 4 for a specific comparison using the star-nosed mole, where these factors have been quantified more extensively). Nevertheless, several qualitative predictions can be made that could be compared to available data or form the basis of future experiments.

Empirical evidence suggests that both receptor densities and stimulus statistics contribute to allocations in neural populations. For example, there are non-uniform distributions of both statistics and densities in primate somatosensory processing. These are also broadly correlated, such that the stimuli are more likely to occur in higher density areas, such as the fingertips (Vallbo and Johansson, 1984; Gonzalez et al., 2014). Recent evaluation of receptor densities has found that they can explain some of the magnification of cortical somatosensory regions, but do not fully account for the representations (Corniani and Saal, 2020). There is also impact of experience, which has been demonstrated in many studies of plasticity in NHPs (Xerri et al., 1996, 1999), see chapter 6. As both the mechanoreceptor densities (Verendeve et al., 2015) and hand use statistics vary across primates (Fragaszy and Crast, 2016), a cross-species study could investigate whether the model can predict the effect of these differences on allocations.

There is a qualitative agreement between allocations in sensory systems and predictions of the presented model. The most prominent sensory bottleneck, the optic nerve is a rather narrow bottleneck (approx. 12-27%)⁴. Furthermore, the retina contains varying densities of cone receptors, with the density depending on the eccentricity. In the fovea, there is a much higher density of cones than the periphery (Wells-Gray et al., 2016; Curcio et al., 1990). Considering these two factors, the model would qualitatively predict a large over-representation of the higher density

⁴Assuming 0.71 - 1.54 million retinal ganglion cells with 80% Midget cells (Curcio and Allen, 1990; Perry et al., 1984) and 4.6 million cones (Curcio et al., 1990)

foveal regions, agreeing with experimental observations (see Edmondson et al., 2019, for further qualitative evidence).

Further tests of the model could be done by comparisons between individuals, which may experience a different number of fibres in the optic nerve, in line with population variation. Those with larger numbers of fibres would be expected to have relatively fewer of these devoted to the fovea (this suggests a larger bottleneck), based on the assumption that these have equal densities and numbers of receptors in the retina. These could be further compared across species, as photoreceptor densities and optic nerve fibre counts vary between primate species (Finlay et al., 2008). While it is likely that densities (Finlay et al., 2008; Verendeev et al., 2015), statistics (Fragaszy and Crast, 2016) and bottlenecks (Finlay et al., 2008) may differ, it is possible to test whether the model can explain allocations over a range of parameters. Similar tests were carried out to test the effect of bottleneck sizes in Lindsey et al. (2019). Finally, the sizes of cortical regions can be controlled in experimental studies with animals (Huffman et al., 1999), which could be manipulated as a test of bottleneck sizes.

Plastic expansion and contraction due to changes in stimulus statistics

In touch, changing the nature of tactile inputs can also affect cortical representations (Coq and Xerri, 1998; Merzenich and Jenkins, 1993; Xerri et al., 1996). Typically, increasing statistics of contact over a region leads to expansion of that region in somatosensory cortex. Our method suggests that this precise level of expansion would be dependent on the bottleneck width, with larger effects for narrower bottleneck sizes. For some sensory systems, such as touch, the size of the bottleneck between the periphery and cortex is unknown. If densities, statistics and cortical sizes are known for a sensory system, the bottleneck could be approximated using this method.

Changing bottlenecks during development and ageing

The width of the bottleneck specifying the resource limit might generally be considered fixed. However, a change in resources could occur due to changing numbers of fibres in periods such as development and ageing. In early development, the number of optic nerve fibres rapidly decrease from their initial number (Sefton et al., 1985; Provis et al., 1985). During ageing, fibre counts further decrease (Dolman et al., 1980; Jonas et al., 1990; Sandell and Peters, 2001). Both of these could be interpreted as a reduction in bottleneck size, which may affect allocations (the specific change would be dependent on the initial bottleneck size). In vision, the optic nerve is a narrow bottleneck compared to the number of receptors in the retina (Wells-Gray et al., 2016; Curcio et al., 1990). With the higher density of receptors in the fovea than the periphery, it is likely the model predicts a large over-representation of the fovea. For a decreasing number of fibres due to ageing, the model would suggest an increase in this over-representation given the shrinking bottleneck size. Furthermore, our model suggests that it is also possible for no differences in representation to be seen if the change in bottleneck width remains in an intermediate zone where neurons are allocated

at a constant ratio. In this case, although the total capacity reduces for both regions, the relative change for both is constant. For example, considering receptive fields, the same proportion would be allocated to both regions. However, these receptive fields may expand for both relative to their previous sizes (assuming an inverse relationship between representation size and receptive field size, (Sur et al., 1984; Nelson et al., 1980)).

There is some evidence that receptor densities may also change; for example, ageing leads to a reduction in tactile receptors in older adults (García-Piqueras et al., 2019). If the number of receptors decreases but the bottleneck remains the same size, the bottleneck effectively becomes increased in capacity compared to the number of receptors. Whether allocations change is also affected by which region these receptors are lost from. Allocations may not substantially change if receptors decrease in higher density regions where these are not taken advantage of, for example, if they have fairly wide-ranging correlations.

Expansion and contraction along the sensory hierarchy

Sensory information is projected to cortex along the tactile pathway, through the brainstem and thalamus to cortex. Throughout these regions, varying levels of magnification of inputs can be seen. In vision, the magnification of the fovea increases between V1 (primary visual cortex) to higher areas such as V4 (Harvey and Dumoulin, 2011).

Our model utilises decorrelation to maximise information in efficiently coding inputs. Decorrelation has shown to be an estimator for optimal coding of representation in lower noise settings (Chalk et al., 2018). Typically lower noise is seen in lower-level sensory processing, and our results likely best explain results of early sensory processing and early cortical representations (primary regions). Previous work has demonstrated that when noise is higher, the optimal code may no longer be based on decorrelation (Hermundstad et al., 2014). Whether these changes would affect neural allocations is unclear. However, these higher noise regions are also likely to be further downstream in the pathway, such that their inputs would be affected by the earlier decorrelation of signals.

Perceptual consequences

Here we consider whether the allocation of output neurons can provide testable predictions to changes to perceptual processing and abilities. In the model, we do not consider the exact shape or locality of receptive fields within regions, only which region they are assigned to. Typically the more resources allocated, the smaller the receptive fields in order to take advantage of the resources. This may then offer higher perceptual spatial resolution. Studies have demonstrated that cortical magnification is correlated with perceptual acuities in both visual (Duncan and Boynton, 2003) and tactile processing (Duncan and Boynton, 2007). However, the receptor densities themselves provide absolute limits on the spatial acuity. If output neurons are assigned proportionally to regions, then the perceptual abilities would be similar across both regions regardless of the densities. However, our model demonstrates that most of the time, this would not be the case, and instead, the allocation

would vary. For example, higher density regions would be allocated more resources for narrower bottlenecks, offering a relatively higher spatial resolution than a proportional allocation. At a wider bottleneck, the spatial resolution would be better for lower density regions than expected. In vision, spatial resolutions decline faster than what would be expected, given an increase in eccentricity, which suggests a narrow bottleneck (Anderson et al., 1991).

Non-proportional resource allocation has perceptual consequences, where neurons have narrower tuning curves over high probability regions, and therefore finer grain discrimination is possible. This can be seen in both the retina (foveal vs peripheral receptor densities, (Wells-Gray et al., 2016)) and in touch (for example, fingertip vs finger base (Johansson and Vallbo, 1979)). Particularly in touch, the higher density of mechanoreceptors on the fingertips corresponds with typical hand usage—fingertips have higher contact statistics compared to other parts of the hand (Gonzalez et al., 2014). Perceptually, the fingertips also have the lowest spatial discrimination thresholds (Johnson and Phillips, 1981). Although we do not explicitly investigate the shapes of the tuning curves, it is likely that as more resources are assigned to a region, these curves would become narrower, thus leading to finer discrimination. Our model suggests that spatial acuity should be better on regions with higher activation probability even when receptor densities are equal. Improvements may also be seen if representations are in part driven by usage. Spatial discrimination can improve with training or even passive learning (Van Boven et al., 2000; Godde et al., 2000). If information is limited, it may be optimal to change the representations to reflect the change in input statistics. If this is the case, then training may offer improvements to some digits at the detriment of others. Whether this is indeed the case has, to our knowledge, not yet been empirically tested. Nevertheless, the absolute limits on spatial acuity seem to be determined by the density of receptors. The total number of receptors remains relatively constant across individuals; however, finger size can vary. This means that those with smaller digits can achieve finer spatial acuity as the receptors are more densely packed (Wong et al., 2013; Peters et al., 2009).

The correlations between densities of receptors on the hand and activation may imply evolutionary constraints. It is unclear whether the increased representation of the fingertips is due to usage or density and whether the placement of receptors is optimised for this higher activation (Plumbley, 1999). Other senses show adaptation to activation by the receptors; for example, in olfaction, receptors have been found to alter their locations to continuously adapt with changes in the statistics of environmental odours (Teşileanu et al., 2019).

3.4.4 Limitations and future work

Here we showed how a simple model (noiseless, linear second-order statistics) demonstrated complex trade-offs in allocations under resource constraints. Output neurons were not allocated proportionally to input regions but instead were influenced by the receptor densities and stimulus statistics. We, therefore, suggest that similar trade-offs may occur in more complex models; however, the precise allocations may vary depending on assumptions. However, the use of decorrelation and

principal component analysis is widely used in models and other algorithms, even those more complex. For example, independent component analysis (ICA) incorporates higher-order statistics but uses the whitened input signal. When this is under-complete, the bottleneck case would be obtained through dimensionality reduction (Hyvärinen and Oja, 2000), such that the inputs would already align with the effects demonstrated in this work.

Sensory signal transmission is affected by noise; for example, in the visual system, incoming signals are subject to blurring from the retinal optics, photoreceptor and neuron noise. One limitation of our model is that we do not consider these effects. It is likely noise will have some effect on the allocations; however, we have demonstrated that even in a simple, second-order model, the bottleneck width has large effects on the representations. It is unlikely that the inclusion of noise would reverse this effect entirely. Some models that do incorporate sensory noise and maximise reconstruction accuracy also use an undercomplete set of principal components to reduce the dimensionality of the sensory signal (Doi and Lewicki, 2014). Between Doi and Lewicki (2014) and Hyvärinen and Oja (2000), the resulting receptive field structure differs, but their allocations (which determine where they will fall on the sensory sheet) is dictated by the same principles described earlier.

Derived allocations in the model assumed a negative exponential covariance function. This enabled us to calculate allocations and convergences (limits) analytically for larger populations than for numerical calculations. Negative exponential covariance functions consider that the receptor correlations decrease with increasing distance and are likely a good match for touch sensing. Other covariance functions occur in some sensory systems, such as vision, which is well characterised, with eigenvalue decay following the power-law $1/f^2$ (Field, 1987). We have previously demonstrated this numerically with covariances calculated from natural images, considering changing density between two regions (Edmondson et al., 2019). Precise allocations may differ with the choice of covariance; however, as long as this decays with increasing distance, our results should qualitatively hold.

In the tactile system, the statistics may change not only in terms of the numbers of contacts for each region, but also the size or area of the contacts (Catania and Kaas, 1995; Gonzalez et al., 2014). It is likely that both density and statistics are positively correlated, while they are both negatively correlated with the relative sizes of different stimuli. For example, when densities increase (e.g. on the fingertips), smaller contacts are also made with stimuli in the form of precision grips (Napier, 1956). In the model, the extent of stimuli size contacts is related to the parameter σ . In our main results presented here, we considered the σ constant over both regions to focus on the effect of manipulating the densities and statistics. However, all three of these parameters will lead to more complex allocation regimes, and as such σ differences will need to be considered in the calculation of the limit. Nevertheless, in the following chapter 4, we demonstrate the incorporation of varying spatial correlations, alongside densities and activation differences. In this case, the allocations are still non-linear, dependent on the width of the bottleneck.

Chapter 4

Efficient coding in the Star-Nosed Mole

4.1 Introduction

In this chapter, we will apply the model presented in chapter 3 to the tactile sensory system of a model organism, the star-nosed mole. We compare the predicted model allocations for each region of the animal’s star-shaped nose to empirically measured somatosensory cortex regions sizes. We investigate between models considering only accurate density or statistics ratios, and then both combined. Furthermore, we predict the potential size of the mole’s sensory bottleneck based on the best fits over all bottleneck widths.

4.1.1 The star-nosed mole

The star-nosed mole (*Condylura cristata*) is a small mammal found in wet lowland areas of North America, known for its unusual, highly specialised star-shaped snout (see Fig. 4.1A), (Ayers, 1884). The mole has been well-studied in the context of understanding how environmental adaptations can lead to specialised systems (Catania, 2020), therefore both the anatomy and behaviour of the mole have been extensively characterised.

The star-nosed mole is a predominately underground dweller with large claws adapted to dig through soil and build extensive tunnel systems. These living conditions result in the mole being functionally blind. Furthermore, the mole is an insectivore, which requires a substantial amount of food to survive. Their diet typically consists of earthworms, small insects and larvae (Catania and Kaas, 1997). To quickly detect and consume prey, the mole relies on active tactile sensing (Catania, 2020), facilitated through a specially adapted snout constructed of 22 mechanosensory appendages. These are arranged in a star-like pattern, fanning out from the centre of the mole’s nose (see Fig. 4.1B). The star is symmetrical down the midline resulting in 11 pairs of rays. The 11th pair,

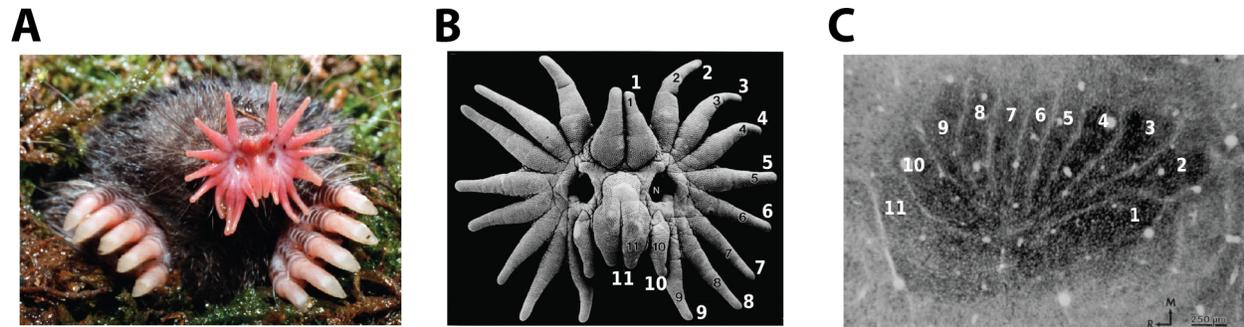


Figure 4.1: A. Star-nosed mole. The star-nosed mole lives underground in extensive tunnel networks built using its long claws. As a result, it relies on its sense of touch over visual information, using its highly specialised tactile ‘star’ to detect prey. From Catania et al. (2011), image copyright Kenneth Catania. **B. Close-up of the star.** Two mirrored sets of 11 tactile rays. Ray 11 is closest to the mouth and is known as the ‘tactile fovea’ of the star. Image from Catania and Kaas (1995). **C. Somatosensory cortical representations of the star.** Rays are topographically represented, with each separated by a septa boundary. Ray 11 has the largest representation. Image from Catania and Kaas (1995).

located on the ventral midline close to the mouth, is involved in nearly all prey interactions. It has therefore been likened to a tactile version of the visual ‘fovea’ (Catania and Kaas, 1997; Catania and Remple, 2004). Similar to other sensory systems, such as vision (Harvey and Dumoulin, 2011), these two central appendages have a greater cortical magnification of their representation compared to other rays in the mole’s primary somatosensory cortex (see Fig. 4.1C).

4.1.2 Anatomy of the star

The whole nose is approximately 1cm wide, with each ray varying in size— rays two and three are the largest and ray ten the smallest. The arrangement of the star shape means a larger surface area of the environment can be explored when searching for prey (Catania, 1996). Each ray is tiled with sensory receptors known as Eimers organs, and around 30,000 covers the entire star (Catania, 1995b). Eimers are Merkel type receptors and are highly sensitive to fine tactile discrimination, enabling accurate identification of prey (Catania et al., 1993; Catania, 1995a, 1996). Eimers organ density within the rays differs by approximately 1:3 between close to the nostrils and the very tip, respectively (Sawyer and Catania, 2016). The Eimers organs are innervated by approximately 106,000-117,000 myelinated fibres (Catania, 1995b); this means the star is more densely innervated than the whole human hand (at approximately 17,000 fibres, Catania and Kaas (1996); Johansson and Vallbo (1979)). Multiples of these fibres innervate each Eimer, approximately four each on rays one to nine, and 5.6 - 7.1 per organ for rays 10 and 11. These fibres are all connected into a single nerve branch for each ray (Catania, 1995b). The total density of sensory fibres between rays differs, with those closest to the mouth exhibiting higher densities. This is also evident in other sensory systems, which feature small areas of high resolution, surrounded by larger regions of lower resolutions.

4.1.3 Behavioural use of the star

The tactile star is primarily used in prey foraging, which follows a stereotypical movement pattern. First, prey is typically detected by the outer ray pairs (1-10), which extend further into space. The prey is then funnelled towards ray 11 — the tactile fovea— where it is examined and possibly eaten (Catania and Kaas, 1997; Sawyer and Catania, 2016). In an experiment examining the ray contact made to prey items during hunting, it was found that in all trials, the moles made mechanosensory saccades. During these saccades, the prey is positioned onto the 11th rays for one or more discrete touches before eating.

When interacting with each stimulus, every ray touch is a discrete behaviour; the rays contact the prey item and are then lifted off again. There are no side-to-side scanning movements or continuous manipulation (Catania and Kaas, 1997). Sensors from a single ray are stimulated when contacting objects, with little simultaneous ray contact. Furthermore, the surface of each ray is mostly separated from all others; and they are only joined to the mole's snout at the base of the nostrils. Within the rays, afferents are activated alongside neighbours, forming a smooth gradient of correlation without discontinuities (Catania and Kaas, 1995). The activity of sensory fibres is therefore correlated within the rays, but less so between them (Catania and Kaas, 1996).

4.1.4 Somatosensory representations of the star

The somatosensory cortex of the star-nosed mole contains a map of the body surface, from mouth to trunk, in a rostrocaudal ordering (akin to a homunculus). Evidence of multiple sets of ray representations may be akin to primary somatosensory, S1, and secondary somatosensory areas. In addition, a third, less prominent representation has been noted, which could be higher somatosensory representations (Catania and Kaas, 1995). This suggests that, as with other mammals, the star-nosed mole may process tactile sensory stimuli in a broadly hierarchical manner.

Similar to the specialised barrel cortex representations of the whiskers in rodents (Petersen, 2007), the mole has a representation of each of the rays arranged in a pinwheel formation (see Fig. 4.1C). The nose occupies approximately 50% of the cortex area which is devoted to touch (Catania and Kaas, 1997). A septa boundary separates each ray, clearly demarcating their representation. These boundaries lead to a lack of correlation between activity from different rays (Catania and Kaas, 1995). Under the assumption that correlated regions of the peripheral sheet typically map together in cortex, these boundaries may suggest a lack of correlation between ray afferent input.

The sizes of the cortical representations are not proportional to the physical size of the rays; instead, there is cortical magnification. Ray 11 has the greatest over-representation, taking up approximately 25% of the somatosensory ray representation (Catania and Kaas, 1995). The mean area of cortex for ray 11, 'the somatosensory fovea', is around four times larger than each of the rays 2-7. Correspondingly, the receptive fields on these rays are around 40% larger than those on ray 11, at around 0.82mm^2 versus 0.59mm^2 (Sachdev and Catania, 2002). Overall average sizes of

receptive fields across rays were less than 1mm^2 .

Studies of cortical magnification versus the number of Eimer's organs or peripheral innervation density have demonstrated that magnification of rays 10 and 11 is beyond what would be expected given the number of Eimer's organs (which is roughly proportional to the surface area of the ray). Furthermore, innervation density also does not correlate with cortical representation sizes. Instead, the magnification can be explained better by the pattern of activation seen across the nose during feeding, with a strong positive correlation between the number of contacts on each ray and cortical sizes (Catania and Kaas, 1997).

4.1.5 Aim of this chapter

Previous work by Catania and Kaas (1997) demonstrated that the cortical representation sizes of the rays do not correspond to the innervation densities but are more closely correlated to the ray usage. Here we test whether the model described in chapter 3 can explain the empirical allocations of each ray's representation in somatosensory cortex, when considering either the innervation densities, usage statistics of the rays or both factors combined. Between ours and the calculations by Catania and Kaas (1997) there are some differences. First, our model is able to consider allocation when both factors are taken into account. Second, previously calculated correlations of usage and ray size were not directly proportional; therefore, even a single factor model may make different predictions to Catania and Kaas (1997).

In the model we utilise empirical estimates of innervation densities and ray contacts which have been extensively documented, as described above. We compare whether the suggestion that the rays contribute more to the cortical allocations than densities holds under the efficient coding model. Finally, we also estimate the possible size of the bottleneck given the resulting fits.

4.2 Methods

We utilise the 2D form of the model in chapter 3, which enables the incorporation of both the densities of receptors and the statistics of the input stimuli. For each ray, we calculate the eigenvalues analytically using equation S.24. These are the eigenvalues of an exponentially decaying input covariance matrix, which considers each ray's density of receptors and activation. We follow the procedure outlined in the previous chapter 3, where the resulting eigenvalues for all regions are sorted in order of their magnitude (see section 3.2.3). Each ray's percentage allocations were calculated at all possible bottleneck widths (determined by the number of retained eigenvalues from the sorted set). Three models were compared, first the 'density only' model, where accurate receptor densities are included; however, receptor activation is fixed constant across all rays. Second, the 'usage only' model includes the receptor activation variances but fixed, equal densities across all rays. The final 'full model' combines both accurate densities and receptor activation variances. Model allocations for each ray were then compared to the cortical allocation empirical data. In

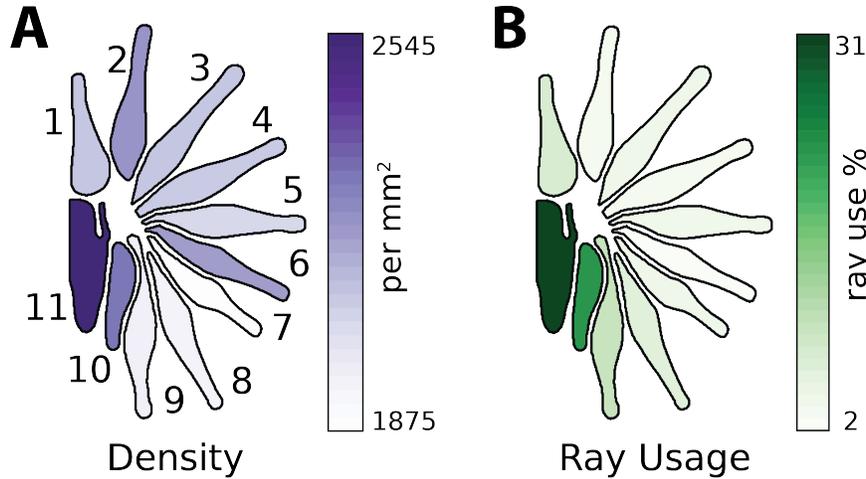


Figure 4.2: A. Innervation densities. Innervation densities for each ray, darker colours indicate ray is more densely innervated. **B. Typical usage of each ray.** Darker shades of green indicate higher usage of the ray. At the centre of the nose, ray 11 is most frequently involved in prey interactions.

the following sections, we outline the calculation of the required parameters of the model based on empirical data from star-nosed mole studies. These include the innervation densities, average contacts with stimuli and typical extents of ray stimulation.

Size and innervation density of the rays

The sizes of each ray were taken from Sawyer and Catania (2016), and are detailed in Table 4.1. Using these, we calculated the innervation density for each ray from the typical number of fibres in each (Catania and Kaas, 1997). A comparison of the average innervation density for each ray is demonstrated in Fig. 4.2A. For the ‘density only’ and ‘full model’, the innervation densities for each ray are set as in table 4.1. For the ‘usage only’ model the densities for all rays were set to the mean value of 2124.

Ray	1	2	3	4	5	6	7	8	9	10	11
Sizes	2.61	2.93	2.91	2.74	2.34	2.09	2	2.55	2.2	1.52	2.4
Sizes %	9.93	11.14	11.07	10.42	8.9	7.95	7.61	9.7	8.37	5.7	9.13
Innervation	5470	6510	6110	5720	4780	4600	3750	4940	4310	3490	6110
Innervation %	9.80	11.67	10.95	10.25	8.57	8.25	6.72	8.85	7.73	6.26	10.95
Density	2095	2221	2099	2087	2042	2200	1875	1937	1959	2296	2545

Table 4.1: Ray sizes and afferent innervation Size of each ray in mm² and corresponding size as percentage of entire star, from Sawyer and Catania (2016). Number of myelinated fibres and corresponding percentages of each ray from Catania and Kaas (1997). Density is the number of fibres per mm².

Variations of ray usage

Approximations of receptor activation on each ray were calculated from recorded interactions of prey foraging behaviours (Catania and Kaas, 1997). The number of prey item touches during each interaction were recorded out of a total possible number of touches. This total was determined by the recording camera’s frame rate and duration of the stimuli interaction. The distribution of touches for each ray is shown in table 4.2, and Fig. 4.2B.

The ray contacts for interactions were converted to ray activation probabilities under a Bernoulli distribution, denoting the chance of the ray being active or not. This was a suitable choice model as each ray’s contact with prey objects is typically discrete, without sideways movements across the prey stimuli (Catania and Kaas, 1997). As the data contained only prey behaviours, we assumed that there would be more off-touches, accounting for periods when the mole is not hunting and engaging in other behaviours. This is shown in table 4.2. Considering a Bernoulli distribution, the variance is $Var[X] = p(1 - p)$. When the number of ‘on’ stimuli are sufficiently small, this becomes a linear relationship around low variances. In both the ‘usage only’ and ‘full model’, the variances for each ray were set as in table 4.2. For the ‘density only’ model the variance was set to the mean of these values, 0.03, for all rays.

Ray	1	2	3	4	5	6	7	8	9	10	11
Usage	15.15	7.29	7.72	7.46	7.97	4.66	8.98	12.00	19.17	44.61	60.61
Usage %	7.74	3.73	3.95	3.81	4.07	2.38	4.59	6.13	9.80	22.80	30.98
Variance	0.03	0.01	0.02	0.01	0.02	0.01	0.02	0.02	0.04	0.08	0.11

Table 4.2: Ray usage parameters Usage denotes the empirical data of the number of contacts for each ray made during prey hunting and consumption interactions, from Catania and Kaas (1997). From this, Bernoulli distributed variances of ray use were calculated.

Fitting covariance functions

The decay rate of the covariance exponential σ was determined for each ray based on the average contact with each type of stimuli, and their sizes (Catania and Kaas, 1997). Larger stimuli covering more of the ray will lead to increased covariances of receptors and thus a flatter exponential decay. One σ value is estimated for each ray, based on the frequency of contacts with different sizes of prey stimuli. We characterised the σ based on the model data in Catania and Kaas (1997), which details the average number of each sized stimuli contacted by each ray. Typically, all rays are contacted by larger stimuli, but smaller stimuli are more frequently contacted by the rays closer to the mouth- rays 10 and 11. Examples of the distribution of stimuli sizes contacted by each ray can be found in supplementary Fig. S.5. To calculate the typical decay of the stimuli, we created a simulation model of each ray, where model circular stimuli of varying sizes were contacted with the rays, and typical extents of activation across the surface were calculated. Full details of the model are provided in the supplementary information B.1. An appropriately decaying exponential

function was fit to the estimated covariances of receptor activation, given their distances, to obtain a σ parameter for each ray. These are displayed in table 4.3.

Ray	1	2	3	4	5	6	7	8	9	10	11
σ	0.99	1.00	1.02	1.01	1.00	1.01	0.99	1.04	1.10	1.16	1.27

Table 4.3: Model sigmas Covariance decay sigmas for each ray, calculated using the prey contact model with varying stimuli size.

Cortical sizes

The cortical sizes for each ray are taken from Catania and Kaas (1997), and are shown in table 4.4. Ray 11 has the largest area, occupying approximately 25% of the primary somatosensory cortex devoted to ray representations.

Ray	1	2	3	4	5	6	7	8	9	10	11
Cortex size	30.00	27.10	20.10	20.10	19.10	18.10	18.00	23.10	29.00	33.10	79.10
Cortex %	9.47	8.55	6.34	6.34	6.03	5.71	5.68	7.29	9.15	10.45	24.97

Table 4.4: Cortical size of each ray representation from Catania and Kaas (1997).

Optimal allocations

The full range of allocations under all possible bottleneck widths was calculated for each model. For each bottleneck, the root-mean-squared error (RMSE) was calculated between the model predicted sizes and the empirical cortical sizes of each ray. Allocations at the bottleneck with the lowest RMSE for each model were selected to compare between them. Several other error measures were tested including the mean absolute error, which gave similar results for the bottleneck sizes as the RMSE, see appendix B.3.

4.3 Results

Allocations of each model for all bottleneck widths are shown in Fig. 4.3A. First, for the density only model, we find little effect of the bottleneck width (Fig. 4.3A, top left). Here, the allocations are relatively consistent for all rays across bottleneck sizes, and are close in their proportions to an allocation simply according to each ray’s innervation density. This could be due to the similarities of the σ parameters for most rays and correlations across the sheet being broader. In contrast, the σ parameter for ray 11 was larger, as it typically interacts with smaller stimuli. This may have influenced the convergence to be at a slightly higher proportion than the corresponding innervation proportion (see Fig. 4.3A, top left. Ray 11, shown in yellow, has the highest proportion out of all rays). The corresponding RMSE between the predicted model allocations and empirical cortical

allocations at each bottleneck is shown in Fig. 4.3B. For the density model, the error is fairly consistent across all bottleneck widths (mean 4.64, purple line in the figure) in correspondence with the similar allocations across the bottleneck in Fig. 4.3A.

In comparison, the models with either usage alone (Fig. 4.3A, top right) or both factors combined (Fig. 4.3A, bottom left) have more variable allocations for each ray. Here, ray 11 (shown in yellow) has a large over-representation at smaller bottlenecks beyond the cortical proportional sizes. In both models, this is followed by ray 10. Although this ray is small in size, it also has a higher typical usage than other rays, due to its proximity to the mouth. For the ‘full’ and ‘usage’ models, representations are similar at both narrow and intermediate bottlenecks. The RMSE values between predicted and cortical allocations for these models are shown in Fig. 4.3B, green and pink curves, respectively. The minimum RMSEs are 2.21 and 1.97, with bottleneck sizes of 38% and 46%. The corresponding best-fitting bottleneck widths are marked on Fig. 4.3A by the dashed line.

Comparing between all models, the ‘usage’ and ‘full’ models provide the best comparative fit to the cortical data, with a marginally better fit for the ‘full’ model, see Fig. 4.4A. The correlation for each ray’s representation size between the empirical allocations and the ‘full’ model is shown in Fig. 4.4B; corresponding plots for all models can be found in supplementary Fig. S.7.

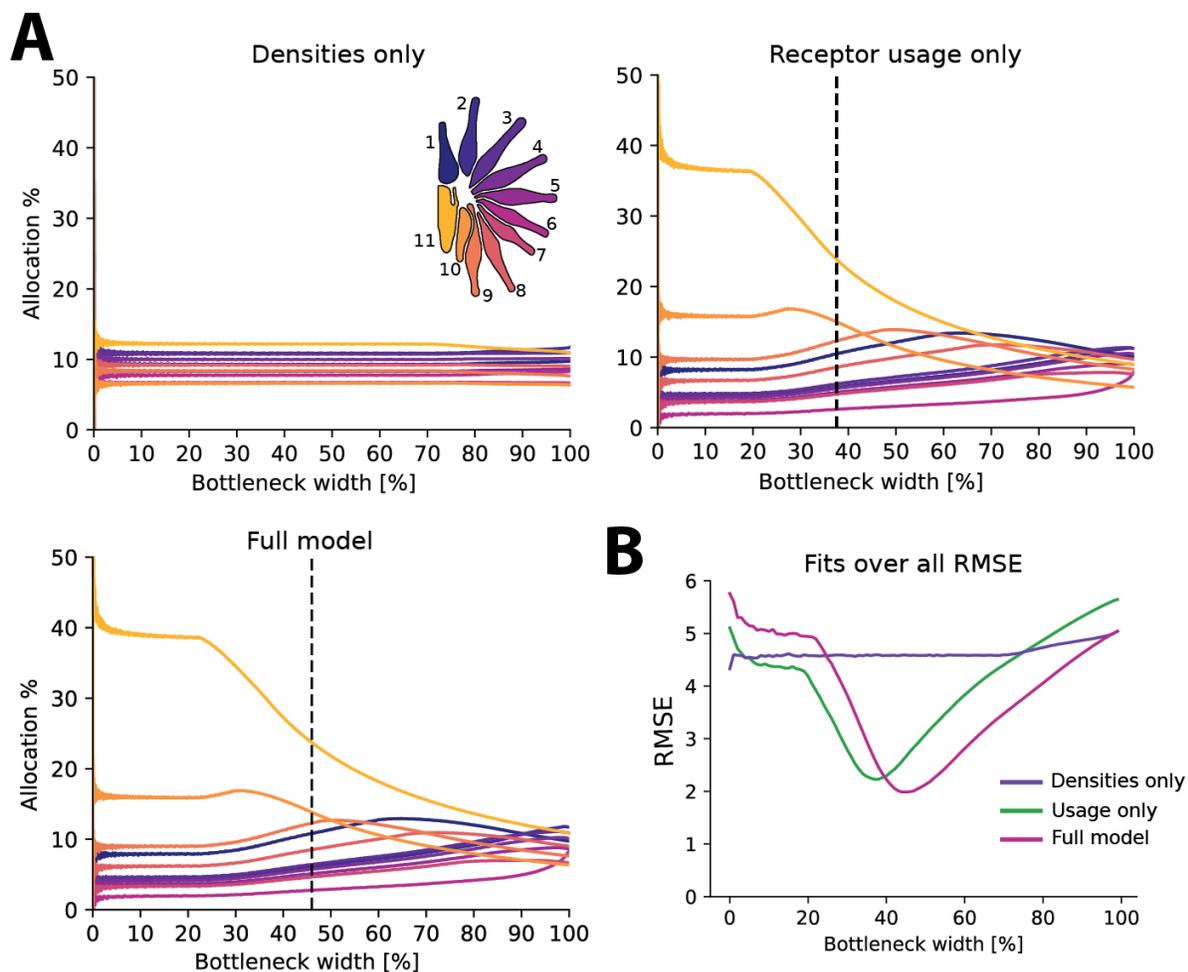


Figure 4.3: **A.** Percentage allocations for each ray over increasing bottleneck widths. At 100% bottleneck, allocations are 1:1 between the number of receptors and cortical units. Dashed lines indicated the lowest RMSE for receptor usage only and full models. Dashed line is omitted for densities only as allocations are fairly stable across all bottleneck widths. Curves are initially similar for usage only and full models but vary for intermediate and wide bottleneck widths. **B.** Total RMSE for each allocation across all bottlenecks. RMSE for the ‘densities only’ model is consistent across all widths. RMSE is greater for ‘usage’ and ‘full’ model as smaller bottlenecks but decreases for intermediate sizes. These also demonstrate a poor fit at wide bottlenecks.

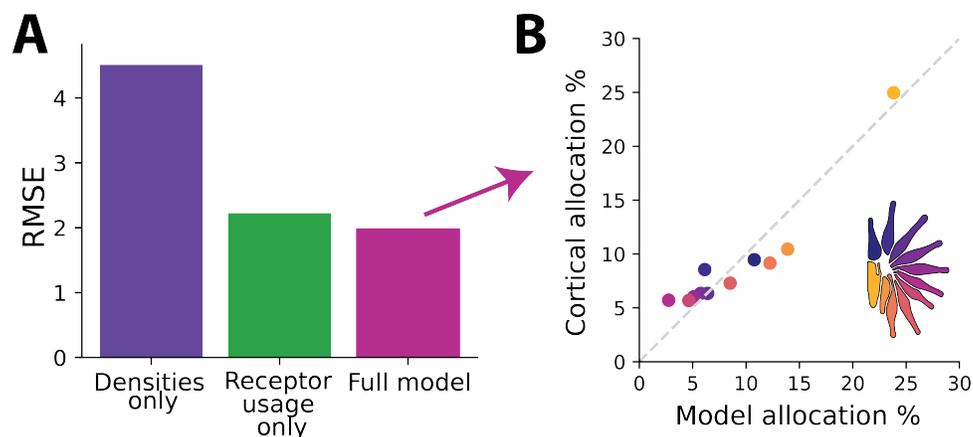


Figure 4.4: **A.** Comparison of RMSE values at the best fitting bottleneck for each model. Usage only model and ‘full’ model provide the best fit, with the ‘full’ model including density marginally lower. **B.** Corresponding fit for ‘full’ model compared to cortical allocation empirical values.

4.4 Discussion

Here we investigated whether a decorrelation model based on efficient coding could explain the pattern of cortical allocations for the star-nosed mole. In the somatosensory cortex, the size of the cortical rays is not proportional to their corresponding physical size. Instead, certain rays, such as those closer to the mouth, occupy a much larger region of the ray representation. For example, Ray 11 has the largest representation at approximately 25%, even though other rays are larger in size. Previous work by Catania and Kaas (1997) demonstrated that usage of tactile rays correlated more closely with cortical allocations than the innervation densities across the rays. Here we tested three models, two considered heterogeneous innervation densities and ray statistics whilst holding the other variable constant, and the final ‘full’ model included accurate values for both factors. Using the decorrelation model presented in chapter 3, we obtained a similar pattern of results, such that the models which included receptor usage statistics obtained a better fit to cortical activation patterns. In comparison to Catania and Kaas (1997), we could also estimate a model with both factors combined. In this case we found a marginally improved fit.

For the two best-fitting models (usage and combined), we estimated the bottleneck between 38% and 46%. There is no obvious bottleneck location in the star-nosed mole, as with the somatosensory pathway in humans. Empirically measured receptive fields in mole’s S1 cortex are small, particularly on the 11th ray (Sachdev and Catania, 2002). This suggests that the bottleneck is likely not extremely small in the ‘narrow’ bottleneck range. In our model, we do not measure precise properties of receptive fields, such as their size, only which region they would be placed. However, previous modelling work has shown (or typically assumes) that regions of an input space with more assigned resources have narrower neuron tuning curves and corresponding increases in discrimination ability (Ganguli and Simoncelli, 2016; Doi and Lewicki, 2014). Although we do not

calculate receptive field sizes, they can be calculated by adding additional constraints in our model (see section 3.2.2), which could be considered in future work.

Information compression through a bottleneck could occur at multiple locations along the somatosensory pathway from the peripheral rays. Aside from cortex, ray representations have only been measured in the PrV (principal sensory nucleus) of the brainstem (Catania et al., 2011). The PrV first receives afferents from the trigeminal nerve, which carries information from the 106,000 to 117,000 total fibres of the rays (Catania and Kaas, 1997). These are then projected to the thalamus before reaching somatosensory cortex. The PrV of the star-nosed mole is larger than those reported for mice and rats, reflecting the prominent tactile capabilities of the mole. Within the PrV, rays are represented topographically, with 11 distinct regions (Catania et al., 2011). Given the ordered representations of both the PrV and somatosensory cortex, it is likely that the thalamus also has a topographic representation, similar to that of primate body regions (Qi et al., 2011; Padberg et al., 2009). In both the PrV and S1, the representation of the rays is given in terms of their volume, therefore the neuron counts or neuronal densities for each ray have not been characterised.

Within the PrV, ray 11 is more magnified than other rays, but not at the extent of the cortex. It has approximately 11% of the total afferent fibres and 14% of PrV volume versus 25% of S1. Considering the difference between PrV and cortex volumes per *afferent fibre*, S1 has a large disparity between different rays, whilst in PrV, volume per fibre is more equal, with much less prominent over-representation of rays 1, 10 and 11 (Catania et al., 2011). Afferent fibre numbers and trigeminal volume were found to have a correlation coefficient of $r=0.84$. This may suggest a wider bottleneck at the level of the PrV, allowing more information from all rays, and therefore more closely resembling the overall fibre counts. Further over-representation, particularly of ray 11, may occur from a constriction at downstream regions, such as the thalamus. However, somatosensory ray representations within these regions have unfortunately not been characterised to test this idea.

Modelling assumptions, limitations and future work

In the model, we use the number of afferent fibres rather than the number of Eimer's organs for our measure of innervation density. Within each ray, Eimer's organs tile the surface and these structures detect the deformation on the surface and detect the stimuli. The number of Eimer's organs is roughly proportional to the surface area of the rays, given that Eimer's are broadly similar in their size. In total, there are approximately 25,000 Eimer's on the rays, which are innervated by 100,000 afferent fibres (Catania and Kaas, 1997). The number of innervating fibres and Eimer's organs overall is highly correlated between animals (Pearson's $r=.988$) (Catania and Kaas, 1997). Although the Eimer's structures initially sense the stimuli, we model the number of afferent fibres per mm^2 as the density rather than the number of Eimer's. This is done for several reasons. First, the discrepancy between the number of fibres and Eimer's organs means there are multiple fibres (approximately five) innervating each Eimer. Within each Eimer, there are multiple components forming the structure, including a Merkel process at the base, lamellated corpuscles and free nerve

endings that terminate at the tip (Catania, 1995b, 1996). These are arranged in a hub-and-spoke layout detecting stimulation at different locations of the Eimer tip. In two of the rays, 10 and 11, the innervation densities are not in line with rays 1-9, and instead, there are more fibres for each Eimer (Catania and Kaas, 1997). Although the exact differences within the Eimer concerning how these additional fibres are connected to the free nerve endings have not been investigated, there is an overall lower convergence of information detected by the Eimers for these rays. This could only be captured in the model by using the number of innervated fibres, rather than the number of Eimers organs (which is a closer proxy of ray size). The higher number of fibres has also been suggested as a possible basis of the higher resolution capabilities for object detection on these rays (Sawyer and Catania, 2016). Second, the afferent fibres transport the information detected by the Eimer's organs from the rays to cortical regions. Therefore it is the number of these that place the limit on the amount of information travelling to cortex (akin to the number of fibres in the optic nerve). Therefore it is this quantity that would be compressed along the pathway to cortex in terms of a resource bottleneck.

We considered the afferent fibre distributions within each ray as homogeneous. However, the distribution of Eimers organs innervated by these afferents differs by a ratio of around 1:3 between the base and ray tip (Sawyer and Catania, 2016). Assuming the number of fibres for each Eimer is roughly constant, leads to a higher fibre innervation at the tips. We were unable to account for this in the model as there is no empirical data for cortical sizes of different subdivisions of the rays. Furthermore, there is no usage data including possible differences in the manipulation of prey items between the tip and base. However, given that the different cortical sizes are explained mainly by ray activation rather than the densities, it is unlikely that considering this would significantly change the results. Furthermore, many of the stimuli typically contacting the rays are relatively large and would cover regions of the ray beyond just the tip. The fovea-like properties of ray 11 are also not related to the higher densities of myelinated fibres (Catania and Kaas, 1997; Sawyer and Catania, 2016).

For the prey model and calculation of receptor covariances given the stimuli, we utilised an idealised, flattened 2D representation of the rays and all stimuli were circular. In prey interactions, stimuli can be many different shapes, and the ray shape deforms under contact with stimuli. As the rays are 3D, it is more likely that only edges or parts of the stimuli are contacted rather than the whole surface. This may further explain why the rays are so densely packed with receptors, as such a packing would not be resource optimal if stimuli are mostly larger than the rays. In our simulation we decreased the size of stimuli to account for this, however future work could consider estimating the extents, sizes and contacts with stimuli more precisely by utilising a more complex model of stimuli interactions.

Our model parameters were estimated from published empirical data of ray densities and usage. The best fit of each model and bottleneck sizes were calculated using empirical cortical sizes of each ray. Sources of error in these measurements and variation between tested animals (for example

individual differences), could lead to variations in the results. Here we discuss the measurements of each of these parameters and potential sources of error. The sizes of the rays were calculated by Sawyer and Catania (2016), with measurements taken from one side of each star for 11 moles (stars tend to be symmetric with no laterality effects). The surface area of the star varied between moles, with a coefficient of variation (CV) at 15.6%. However, within each mole, the relative sizes of the ray were consistent. We estimated the afferent densities based on the number of fibres innervating each ray. The number of Eimer's organs and afferent innervation of rays were taken from four moles. The total number of Eimer's organs (a broad proxy for size) varied between approximately 23,500 to 27,000 in the moles; nerve fibre counts were between 106,000 and 117,000. As the number of afferent fibres per Eimer is relatively constant, moles with larger rays (and therefore more Eimer's) have a larger number of innervating afferents. The number of Eimer's organs is highly correlated to the number of afferent fibres across animals (Pearson's $r=.988$) (Catania and Kaas, 1997). Ray usage variances were estimated using data from ten behavioural interactions with prey, as reported in Catania and Kaas (1997). In these trials, small pieces of earthworm were used, however, it is unclear if these stimuli sizes are similar to what is encountered in the wild. Variations in prey size could affect which rays are utilised when funnelling prey towards the mouth, and therefore, the overall distribution of ray contacts. Later experimental tests with over 500 trials using a larger range of stimuli sizes demonstrated a similar distribution of ray contacts (Catania and Remple, 2004), however, as before, the exact distribution of each prey size encountered in the wild is unclear. Cortical measures were taken from four moles (Catania and Kaas, 1997). The estimated area of ray 11 had a narrow range from 22 to 26% of ray somatosensory cortex between moles, with a mean of 25%. Although the cortical sizes of each ray, innervation densities and number of Eimer's organs (ray size) were measured in the same moles, individual data points for each parameter were not presented in Catania and Kaas (1997). Therefore, individual mole cases could not be reproduced in the model.

Chapter 5

Self-organising models of somatosensory representations

5.1 Introduction

This chapter will investigate a self-organising model of sensory cortical representations that has been popular for understanding the effects of plasticity. In previous chapters, we created and tested a more abstract model of the effects of receptor density and input statistics. Here, we examine these factors in the Kohonen SOM model, where regions compete for representation, and the model learns the mapping of the input space over time from a random initial state. Furthermore, the model learns a set of feedforward connections between the inputs and the map, enabling receptive fields to be examined.

5.1.1 Somatosensory models

For more than 50 years, the topography and feature mappings of early sensory cortical areas have been investigated with computational modelling. As a result, a plethora of models and theories of map formation have been proposed, including self-organising maps (SOM) (Kohonen, 1982; Willshaw and von der Malsburg, 1976), dynamic neural fields (Takeuchi and Amari, 1979; Amari, 1980; Zhang, 1991) and neural group theory (Edelman, 1981), to name a few. For the somatosensory system, the predominant interest of these models is cortical plasticity, particularly reproducing classic plasticity manipulations first demonstrated in the '70s and '80s by Merzenich and colleagues (Merzenich and Brugge, 1973; Merzenich et al., 1984). However, few have explicitly studied the effects of the structure of the inputs and sampling on the initial map (i.e. pre-plasticity manipulation). This chapter investigates the effects of manipulating the receptor densities and statistics of input in a model previously used to produce topographic representations—Kohonen SOM. We investigate whether this simple model can reproduce key characteristics of somatosensory cortex, such as magnification of either high density or higher activated regions and inverse-receptive field

magnification patterns. As a result, we demonstrate potential problems with this model in reproducing cortical maps and discuss improvements to future somatosensory models.

Sensory cortical map formation is influenced by both the environmental inputs and map-level mechanisms, including biological constraints on network architecture and learning mechanisms. Although initial topography is likely set up by genetic processes, which are important for broadly correct map development (Willshaw and von der Malsburg, 1976) (see chapter 2, 2.4.1), self-organising processes could refine these representations, giving rise to the highly organised maps seen across cortex. Self-organisation is unsupervised and usually implemented by competitive interactions between neurons, with learning governed by Hebbian type mechanisms (Kohonen, 1982). Models based on these principles have had some success in simulating representations of primary visual (Martin and Obermayer, 2009), auditory (Ritter and Schulten, 1986) and somatosensory (Ritter et al., 1992) cortices.

Willshaw and von der Malsburg (1976) proposed that topographically ordered connections could be formed through competition in a step-by-step manner over time. Initially, broad connections between the input space and cortical maps are pruned back via modification of synapses, such that they project predominately to localised regions. Next, competition can be implemented through a set of local excitatory and longer range inhibitory processes. This strengthens responses of nearby map regions whilst weakening those of more distant regions. Finally, correlations between features in the input space lead to those with similar responses mapping closer together, forming topographic map representations.

Kohonen later proposed a model of self-organisation which abstracted elements of the Willshaw and von der Malsburg (1976) model, such as the explicit lateral connections. Kohonen SOM is based on the competitive interactions between neurons and incorporates Hebbian synaptic learning mechanisms (Kohonen, 1982). An example of the network structure is shown in Fig. 5.1. A connection between the model and real cortical maps was identified, as the SOM enables similarly structured inputs to be mapped in nearby locations in the map space. Subsequently the model has been used to investigate somatosensory (Ritter and Schulten, 1986; Ritter, 1990; Obermayer et al., 1990b; Ritter et al., 1992), auditory (Ritter et al., 1992) and visual cortex responses (Swindale and Bauer, 1998; Shouno and Kurata, 2001). The model learning process is unsupervised, such that there is no target activation and backpropagation of errors to calculate adjustments to the network weights. Instead, the topographic organisation emerges based on the structure of the input data presented (Kohonen, 1990, 1982), such that similar inputs are represented closer together (for example, in neighbouring locations). Weight connections become tuned to specific patterns in the input via competitive mechanisms. When an input is presented to the map, the most active unit and its neighbours have their connection weights strengthened, thereby increasing their affinity for these inputs. Competition between units means the model acts in a ‘winner-takes-all’ fashion, with the updated neighbourhood’s location being decided by the most active unit. The neighbourhood effect aims to abstract the lateral connections between cortical neurons, whereby close range neighbours

excite others with similar responses, and most distant regions are inhibited. Weight connections are updated using Hebbian mechanisms, based on synaptic plasticity (Hebb, 1949). Over time, both learning rates and neighbourhood update sizes typically decay to form a converged map. The model has no general closed-form analytical solution for this steady-state, converged mapping (however, see below for solutions in 1D under specific assumptions). Instead, the map is learned through an iterative process of input presentation and weight updates (Kohonen, 1990).

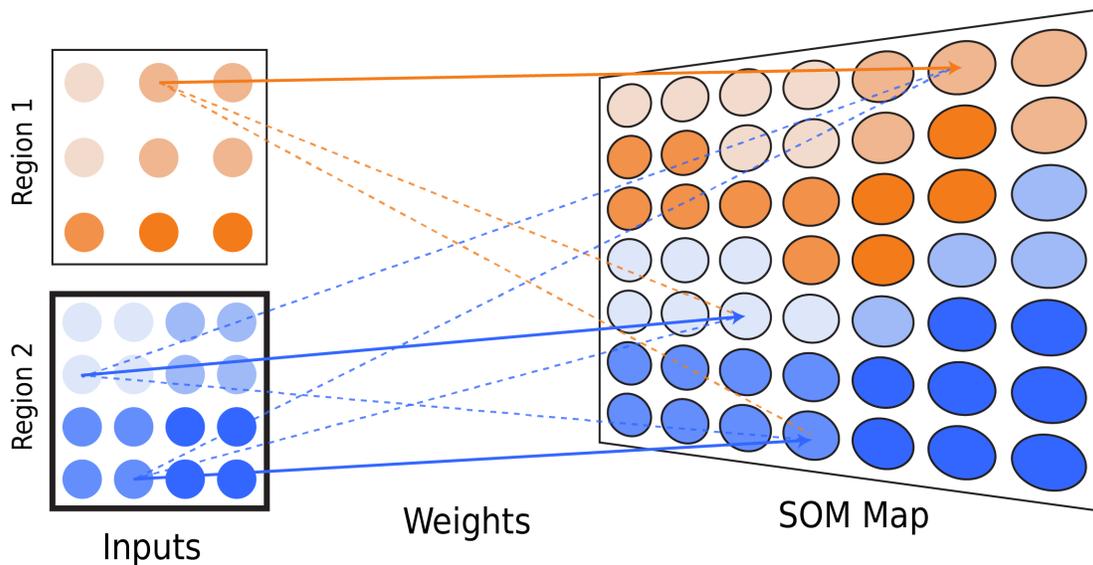


Figure 5.1: Kohonen SOM structure. Kohonen SOM has a feedforward style network structure. Map weights are learned between a set of inputs and each map unit. In our setup, the inputs are defined as two separate 2D regions which can vary in both their receptor densities and their stimulation probability. Region 1 demonstrates a lower density region, with fewer receptors within the space compared to region 2. The thicker outline on region 2 denotes a greater stimulation of this region, akin to the higher stimulation probabilities of certain regions of the body. When the model is initialised, inputs and map units have random weight connection strengths. During training, the weight connections are updated, with map units developing stronger connections to some regions rather than others. Here, solid lines denote strong weight connections and dashed lines denote weaker connections. The final trained weights can be visualised as a map representing different input regions. The map typically develops within region topography (see relative positioning of units representing each region with respect to their spatial positioning in the map). Furthermore, some input regions can have a larger number of input units with strong weight connections to map units than other regions. This leads to an over-representation of the region (see SOM map, which has more units representing the blue ‘region 2’ rather than orange ‘region 1’).

SOM as a model of somatopy

Qualitative comparisons have been made between the Kohonen SOM mapping and somatosensory maps of humans and NHPs. Most studies focus on local organisation of single body parts, such as the hand region; however, some have studied the global ordering of the body representation (Stafford and Wilson, 2007).

Early models utilised rudimentary hand surfaces, containing several digits and a palm (Ritter and Schulten, 1986; Ritter, 1990), with afferent receptors distributed over the surface. Initially, random

weighted connections are made between the receptors and a set of 2D map units, with inputs presented as activation of the model afferents across the hand surfaces. After some training, the models form stable representations of the input space, driven by decreasing the spatial extent of the weight updates and learning rates. In both studies, topographic representations emerged of the input space, which can be qualitatively likened to the natural topography seen in somatosensory cortex.

The model was subsequently scaled up by Obermayer et al. (1990a), using a more realistic hand shape with more receptors and a larger output map. As with Ritter and Schulden (1986); Ritter (1990); Ritter et al. (1992), topographic formations of the hand shape emerged, with neighbouring digits adjacent in the maps. Obermayer et al. (1990a) expanded on previous work by studying the receptive fields (RFs) of the map units and qualitatively compared them to empirical data. They found that RFs were mostly localised, covering regions on individual digits, with their sizes approximately equal to those of the tactile stimuli. Localisation occurs as most weight connections to other digits' receptors decay to zero in the learning process. Nevertheless, as with somatosensory cortex, some multi-digit RFs were noted. Other deficits were also seen in the model maps, such as incorrect digit ordering or separated, double representations of digits. Multi-digit representations have been noted in empirical data (Merzenich et al., 1984); however, digit ordering is typically more stable in experimentally observed maps than simulated ones (Martuzzi et al., 2014; Schweizer et al., 2008; Kolasinski et al., 2016; Merzenich et al., 1987).

These studies did not directly test the effects of manipulating the density of afferents and probability sampling over the input space. Obermayer et al. (1990b) randomly placed receptors over a hand-shaped input space, and Gaussian stimuli were placed at random locations (which would lead to approximately uniform probability sampling). In Ritter and Schulden (1986), receptors were also randomly placed, but in an attempt to model the change in densities seen across the hand, sampling of the input space was weighted towards the digit tips. However, as demonstrated in chapter 3, exchanging the density for statistics is not equivalent, therefore it is questionable that this effect is as intended. Furthermore, in Ritter and Schulden (1986), the resulting differences between the sampling of different areas were not directly assessed. Instead, these studies focused on the qualitative production of topographic somatosensory-like maps. Beyond this, the results of plasticity experiments in NHPs were implemented and reproduced. Merzenich and colleagues (Merzenich and Brugge, 1973; Merzenich et al., 1984) had demonstrated that a lack of input from amputation or nerve resection leads to reorganisation, with neighbouring region representations invading those of the removed or silenced digit. This can be thought of as a change to the input sampling distribution. Amputation is implemented in the Kohonen SOM after the map has converged to a steady-state. The map is further trained, but the inputs of the amputated region are removed (Ritter and Schulden, 1986; Obermayer et al., 1990b,a), leading to reorganisation whereby map representations of the neighbouring digit invade the representation of the removed digit. In qualitative agreement with the empirical studies of Merzenich et al. (1984), local magnification

of the neighbouring digits was seen. However, these plasticity studies consider changes after the initial map has been learned, which occur over more localised areas of the map and when learning rates are set lower. In this chapter, we are focused on the effect of receptor placement and input sampling distributions on the *initial* map development. Plasticity mechanisms of somatosensory cortex will be investigated further in the following chapter, Ch.6.

Cortical magnification factors depend on model assumptions

In the SOM model, it has previously been suggested that an optimal mapping under an information-theoretic view should reproduce the distribution of input probabilities, assigning more units to those regions which are more frequently activated (see also the work of Plumbley (1999)) Given a probability density, P , over the input space, the relationship between the mapping and the input density is governed by some exponent p , such that $M(w) \propto P(w)^p$. When the magnification perfectly reflects the input sampling, the exponent should be $p = 1$ (Kohonen, 1984). In cases where there is a uniform mapping and no adaption to the input statistics, $p = 0$. Note that this only considers the frequency of inputs in a region but does not consider the spatial density of sampling (for example, different afferent densities), which could further affect the magnification. The magnification factor of a region is calculated once the map has settled and reached its stable ‘stationary state’. Here, the learning rate is typically low, and weight vector changes become negligible, such that they do not alter the connection strengths enough for a map unit to flip its preference to another region. The exact magnification factor is difficult to calculate as there is currently no analytical solution for the stationary state of 2D maps (Cottrell et al., 2018). Even for 1D mappings, there is not always convergence to a stable state with the same ordering (Rojas, 1996). Some analytical convergence solutions have been found for 1D maps; however, these have strict assumptions and constraints on model parameters, such as uniform input distributions, a lack of neighbourhood decay over time, or restricted updates only applied to the two immediate neighbours (Cottrell et al., 2018; Bouton and Pagès, 1994). Ritter and Schulten (1986) found the magnification factor at convergence for the 1D continuum case (assuming infinite neighbours) to be $p = 2/3$. Here the map magnifies but does not sufficiently magnify high probability regions to be theoretically optimal. This is not in line with the original assumption that the magnification linearly follows the input density P (Kohonen, 1984). However, the magnification factor somewhat depends on the choice of parameters, such as the neighbourhood size and kernel type. It is unclear what the magnification factor follows a similar pattern for 2D maps, with under-representation of higher probability regions. It is possible to approximate magnification based on numerical simulations; however, there may be accuracy difficulties due to the stochasticity of initial model conditions (such as the model weights).

Beyond Kohonen SOM, several other unsupervised learning based models have been developed of somatosensory cortex, with varying degrees of complexity. Additions such as cortical columns and inhibitory-excitatory connections (Grajski, 2016), hierarchical processing with layer representation of the thalamus (Armentrout et al., 1994; Sutton et al., 1994; Weinrich et al., 1995) and subcortical

structures have been implemented (Grajski and Merzenich, 1990). Similar to Kohonen SOM, reproducing effects of plasticity studies, including amputation, map lesioning, digit syndactyly and multi-digit stimulation, is a popular line of work (Sutton et al., 1994; Mazza and Roque-da Silva, 1999; Mazza et al., 2004; Grajski and Merzenich, 1990; Armentrout et al., 1994; Bhand et al., 2011; Grajski, 2016; Miikkulainen, 1991). Common to many of these models and the Kohonen SOM is competition, whether implemented as an approximation with a winning unit and neighbourhood (as in Kohonen) or with explicit lateral connection and interactions dynamically changing between units.

Of these models, density differences were only considered in Mazza et al. (2004). However, the distributions of receptors over the hand surface were not explicitly modelled, and instead, the input region sizes (representing each digit) were increased for the higher density digits, coupled with uniform sampling. Furthermore, this is a complex, conductance-based model, meaning any effects of the density implementation are difficult to isolate and compare to other models. The inverse magnification rule has been demonstrated in models with Hebbian learning, multi-layer architecture and simple circuits of both excitatory and inhibitory cells (Grajski and Merzenich, 1990). Post-map development, over-stimulation of a region leads to magnification and corresponding smaller receptive fields over the region.

5.1.2 Aims of this chapter

As previous models have focused predominately on cortical plasticity, here we instead investigate the effects of manipulating the density of afferents and sampling distributions. Although more complex models have been developed, we first systematically investigate using the simplest model of topographic map development based on correlated inputs and Hebbian learning rules– Kohonen SOM. One previous study modelled the densities via increasing the sampling probabilities at typically higher densities (Ritter and Schulden, 1986). However, in the previous chapter, Ch.3, we demonstrated that manipulating the density and statistics of the input sampling can lead to different optimal representations. Here we test simple manipulations, changing the density ratio of the sampling between two separate regions. We discuss the assumption required of the structure of the inputs for the proxy sampling of the density to be valid. Finally, we test changing the sampling statistics between the regions and the resulting size of the receptive fields. Given increases in sampling probability, we investigate whether the inverse magnification rules hold.

5.2 Methods

5.2.1 Input sets

In the presented simulations, we change both the density (first and second investigation) and finally the statistics of the inputs sets. Here we describe the basic setup of each of these input sets and the stimuli utilised.

Receptor densities

For each Kohonen map trained, inputs can be placed on either of two separate input regions (see Fig. 5.1). In each of these, model receptors are uniformly distributed depending on their density ratios. The ‘baseline’ density region, $D = 1$, had 100 receptors total within the region of size 10x10. The second region either had the same density (1:1 ratio) or a higher density of points within the same sized space, at ratios of 1:2, 1:4 and 1:6. Figure. 5.2 (left to right) demonstrates examples of afferent placements for the baseline density and two examples of higher densities (1:2 and 1:6).

Region stimulation statistics

The stimulation statistics of regions can also vary, akin to the varying probabilities of tactile stimulation across the skin. This is implemented in the model by manipulating the number of active input patterns presented from each region during training. When a region has a higher probability of stimulation, the total input set will contain more input patterns where stimuli were placed on that region. For example, in the case of a 1:2 ratio, approximately 33% of the inputs will be from stimulation of region 1, and 66% from region 2. As with the innervation density, we test ratios of 1:1, 1:2, 1:4 and 1:6.

Stimuli sizes

As with previous work (Obermayer et al., 1990a), stimuli were modelled as 2D Gaussian bumps, with mean μ and standard deviation c , given by the following equation:

$$r_i = e^{\left(-\frac{(i_x - s_x)^2}{2c^2} - \frac{(i_y - s_y)^2}{2c^2}\right)} \quad (5.1)$$

where r_i is the response of the receptor i at position (i_x, i_y) to stimulus at position (s_x, s_y) . Circular Gaussians were used, such that they extend with the same standard deviation (here, c) in both axes. To account for any border effects of the activation over the map, the position centres of the stimuli (s_x, s_y) could be placed randomly within a boundary extending to the region’s size, with padding equal to the width of the Gaussian. This setup is a fair approximation to touch as objects can make contact with only the borders of regions; for example, in some precision grips, only the edges of the fingers may contact an object. This ensures that receptor activation is fairly even between the centre and the region edges. The complete set of responses I is a $m \times t$ array of receptor activation strengths to each stimulus, where m is the total number of receptors and t is the number of training patterns to be presented, corresponding with the total timesteps (training iterations of the model). Regions are considered completely independent, such that there is no overlap in the activation.

5.2.2 Kohonen SOM mapping algorithm and parameters

The Kohonen SOM takes input vectors of n -dimensions and maps these onto a lower, typically one or two-dimensional surface. For somatosensory modelling, 2D maps are most common to replicate the surface representation of the cortex. The feedforward network is fully connected, with all input

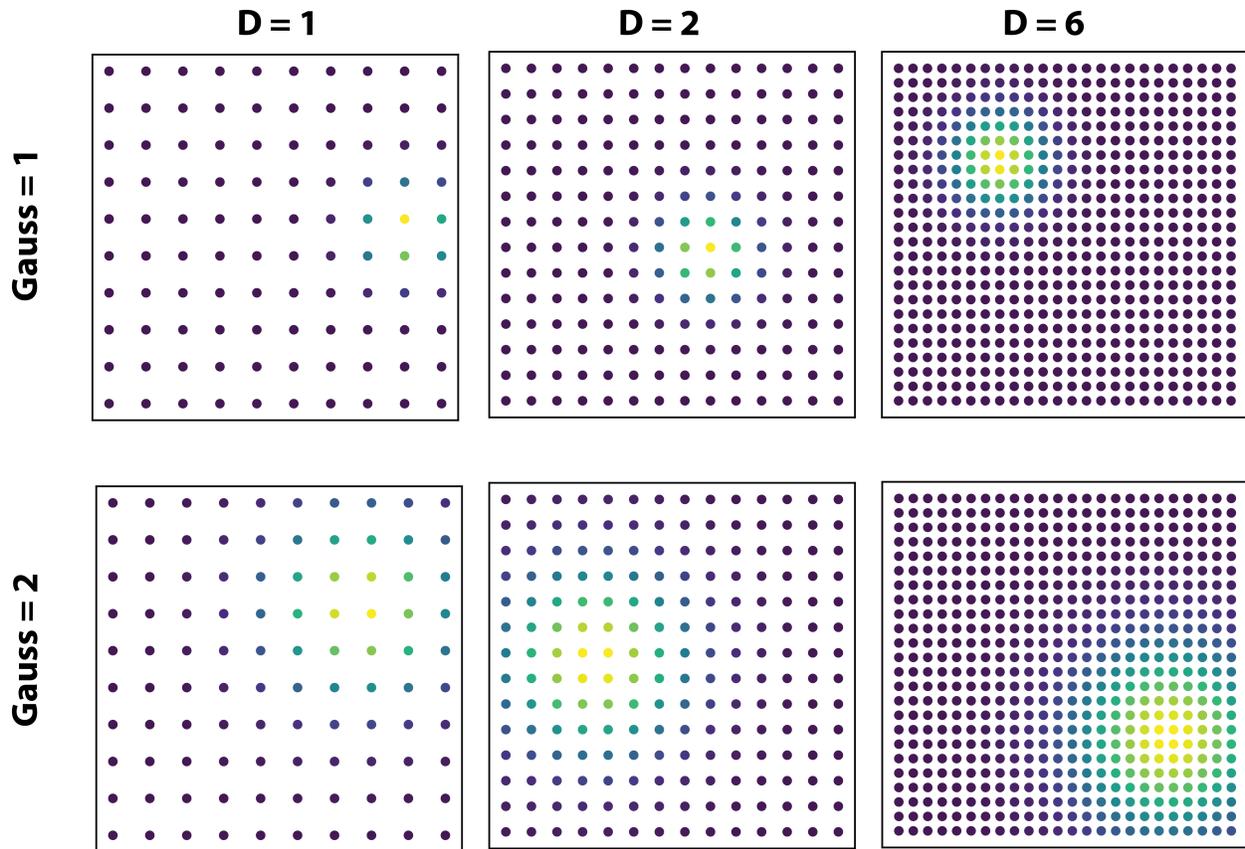


Figure 5.2: Example of stimuli coverage and resulting activation of receptors. Shown is the response of receptors to regions with an increasing density of receptors (columns) to two different sizes of Gaussian stimuli (rows). Colour of each receptor indicates the response level, with yellow indicating a higher response. Darker blue regions demonstrate little or no activation of a receptor. The placed stimuli centre lies close to the most activated unit. Between increasing density for each stimuli size, the extent of the stimulus activation over the region is the same; however, for higher density regions, more receptors are stimulated as more are packed into the region. The lowest density region ($D=1$) is set to the ‘baseline’ for each test, and either the same region or higher density regions are compared.

units connected to every map unit. These weights are the target for changes during the learning process. The model architecture can be thought of like a two-layer neural network, with inputs being learned by one output layer. Unlike somatosensory cortex, the model does not have explicit lateral connections, and instead, these are broadly approximated using the local neighbourhood bubble during the weight update. As the mapping is learned, initial random weights are updated to strengthen connections between map units and input sheet receptors. Strong weight connections are reinforced as learning progresses, and weak connections decay close to zero. Through the neighbourhood function, adjacent units learn to become responsive to similar inputs. As the neighbourhood size decays, broad similarities in responsiveness to more distant map units become weaker, but responsiveness to similar input spaces of adjacent units becomes stronger. This enables a topographic mapping to occur based on the correlations of the input receptors. The entire process of map training is described below.

We set the map to be a 2D sheet of units in our implementation. This can be likened to somatosensory cortex, in that each unit can be thought of as a small patch of cortex representing a population of neurons that are broadly tuned to the same inputs. The model weights between the receptors and map units are represented as a $u \times m$ matrix W , where u is the number of map units, and m is the number of receptors. Here u is the flattened representation of the 2D map to allow easy computation but can be reshaped into the 2D form. All weight connections are initially set to small random values, with a new set of random values selected to initialise each map.

The map is trained in an iterative fashion, where each iteration involves several steps. Firstly, the activation of each map unit, M , is calculated in response to a given input. From each input set (combinations of density and Gaussian size), inputs were presented randomly from the full set. Here we use the dot product between the weights and input from each receptor:

$$M = WI_i \quad (5.2)$$

where I_i is the i th input vector of size m , containing the activation for each receptor to a stimulus. The resulting M is a vector of size u . This vector can be reshaped into a 2D map, allowing visualisation of the spatial extent of the activation between map units.

Second, the set of weight connections are updated. The update is governed by two parameters, the learning rate η , and the size of the neighbourhood, θ .

From M , the location of the best matching unit (BMU) b is found, which is defined as the most active unit in M (highest valued). This enables competition between the inputs through a winner-takes-all mechanism. The BMU at time t is:

$$b(t) = \operatorname{argmax}(M) \quad (5.3)$$

Once the BMU $b(t)$ is found, a neighbourhood is created around this unit. This defines the area of map units that will have their weights from the receptors strengthened in the update. Here we use a ‘bubble update’, where the same update strength is applied to all units in the neighbourhood boundary. Other neighbourhood choices exist, such as the Gaussian kernel; however, we do not investigate the choice of the neighbourhood type in this study. The training schedule governs the size of this bubble and typically decreases over time. All map units included in the neighbourhood are those where their 2D Euclidean distance from the BMU is less than the current neighbourhood radius, $\theta(t)$. This neighbourhood of BMU $b(t)$ is therefore defined as:

$$N_{b(t)} = \sqrt{(l_1 - b(t)_1)^2 + (l_2 - b(t)_2)^2} \leq \theta(t) \quad (5.4)$$

where l_1 and l_2 are the vectors of locations for each map unit in each axis, $b(t)_1$ and $b(t)_2$ are the

axis locations of the current BMU. $\theta(t)$ is the current radius of the neighbourhood.

The update depends on a map unit's inclusion within the neighbourhood. Weights from units within the neighbourhood are updated using the current learning rate $\alpha(t)$:

$$\eta_j(t) = \begin{cases} \alpha(t), j \in N_{b(t)} \\ 0, j \notin N_{b(t)} \end{cases} \quad (5.5)$$

where j is the index of the map unit.

The rate of the learning, or how large the changes are to the weights is determined by the learning rate. The SOM update is:

$$W_j(t+1) = W_j(t) + \eta_j(t) \cdot (I(t) - W_j(t)) \quad (5.6)$$

where W_j is the weight at unit j , $\eta_j(t)$ denotes learning rate of j (as assigned above) and I_t is the input vector, both at time t . Training continues in this fashion, with the learning rate and neighbourhood size decreased for each consecutive presentation. As these values reduce, the map settles into a steady-state where weight updates become small, and the overall map representations are stable.

5.2.3 Selected SOM parameters

The Kohonen SOM algorithm has two adjustable parameters that decay exponentially with training time- the learning rate α and the neighbourhood size, θ . The selected parameter values are detailed in table 5.1. Neighbourhood sizes begin with a wide coverage of the map to set the initial placements of the two regions and decay to just a local update of the unit and its immediate neighbours.

Parameter	Value
Map size	14*14= 196
Neighbourhood size start	98 units
Neighbourhood size end	2 units
Neighbourhood decay	0.0005
Lr start	0.5
Lr end	0.001
Lr decay	0.001

Table 5.1: Kohonen map parameters

In choosing the number of map units, we tested several map sizes. Different to the study of chapter 3, we found little effect on the input allocations of reducing the number of map units, such that there are less than the number of input receptors (a bottleneck). This suggests that the cost function being solved by the Kohonen SOM is different to that of the previous study. Similarly,

there was also no effect of increasing the number of map units. This is shown in Fig. 5.3 for an example statistics ratio of 1:2. We therefore chose to approximate the total number of receptors of two ‘baseline’ density regions. The same map size was used in all simulations, set at 14 x 14 (196 total units).

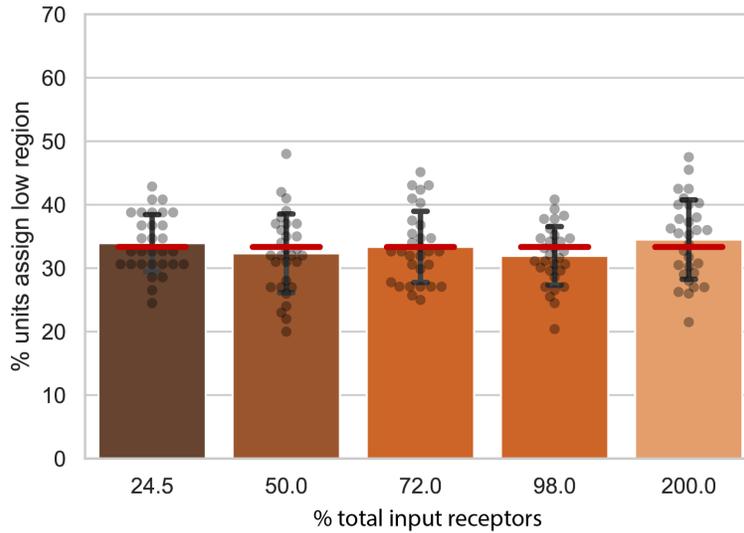


Figure 5.3: Map units assigned under different map sizes. Here we demonstrate that the assigned number of units to the baseline region is largely constant regardless of a change in the size of the map. Here each bar refers to a different map size, values are given as percentage of receptor units, where 100% means an equal number of map units to input receptors in both regions, therefore, less than 100% indicates a bottleneck. Example is for a statistics ratio of 1:2, with 30 maps generated for each.

5.2.4 Simulations

In the first investigation relating to the density of the receptors, we hold the statistics of input constant at a 1:1 ratio, such that stimuli are placed at a uniform probability between the two regions. A total of 40,000 inputs were created for each network run. Therefore responses to 20,000 stimuli from each region were calculated.

In the final investigation of the chapter, we instead manipulate the statistics of the input. The densities between regions are kept constant, and we use the ‘baseline’ density setup described above, $D = 1$, for both regions. The probability of a stimulus being placed on the ‘baseline’ lower-probability region versus the higher probability region is either 1:2, 1:4 or 1:6. The total samples were 40,000 as before, split between the regions according to the probability ratio. We also simulated an evenly sampled condition, where both regions were equally active (sampling ratio of 1:1, 20,000 inputs from each region). In this case, we expected the Kohonen map to allocate an equal number of output units to each region.

In all simulations, the effects of manipulating the Gaussian stimuli’ size were also investigated.

Here we set the parameter c in Eq. 5.1 to a width of either 1 or 2 (Fig. 5.2 top row versus bottom row). A minimum size of one was chosen as this still produces some co-activation of neighbouring receptors in the lowest density ‘baseline’ region (see top left of Fig. 5.2). Of course, between the densities, the spatial extent of activated receptors is the same, as this is calculated relative to the size of the region only. However, neighbouring receptors are more correlated in their responses at higher densities due to the tighter packing within the space. This can be seen by comparing different densities in Fig. 5.2.

5.3 Results

5.3.1 Effects of varying receptor densities in the Kohonen SOM

First, we investigated how changes in the density ratio between the two regions, when sampling statistics are held constant, affect the mapping representations in the Kohonen SOM algorithm. Once the map is trained to a steady-state, two broad clusters emerge (in most maps), each representing one of the two regions. The placement and shape of the clusters vary due to the random initial weight set-up and input presentation. An example of the mapping from weight connections of each receptor into the map is demonstrated in Fig. 5.4. Most maps form two distinct clusters of the regions; however, sometimes discontinuities emerge, such as the map shown in Fig. 5.4 right column. Here, a further small cluster representing the higher density region has emerged. Along the borders of region representations, map units can have some mixed selectivity for both regions.

We quantified and compared the size of each of the regions using a winner-takes-all (WTA) calculation of each map unit. For each map unit, the weights from each region’s receptors are separately summed, and the region with the largest total weight is selected as the winner. Fig. 5.4, bottom row, demonstrates the resulting winner-takes-all mapping from the two sets of region weights.

We found that for the 1:1 density, as expected, the mapping is approximately 50% devoted to each region (see Fig. 5.5A, darkest orange bar). There is some variation between individual maps due to the randomness of the initial weights and order of input presentations. This is demonstrated by the distribution of grey points. In the 1:1 case, even mapping occurs regardless of the size of the stimuli. For increasing densities of receptors, the division of map units between the regions remained relatively stable, with a slight decrease in the number of units assigned to the lower density region. However, even for a higher density difference of 1:6, this was only reduced by around 5-6% of the total map units. The standard deviation (due to individual differences in mapping) was approximately even across different densities, at between 5.8 and 7.47%.

The small decrease is due to the amount of total activation varying between the two regions, given an increased number of receptor points for the higher density. This leads to more weight connections between this region and map. Therefore, this increases the total weight connection sum from the higher density region. By equalising the mean of each of these datasets (normalising by the mean), this effect is eliminated (see Fig. 5.5B). The increased density in the second region likely offers some

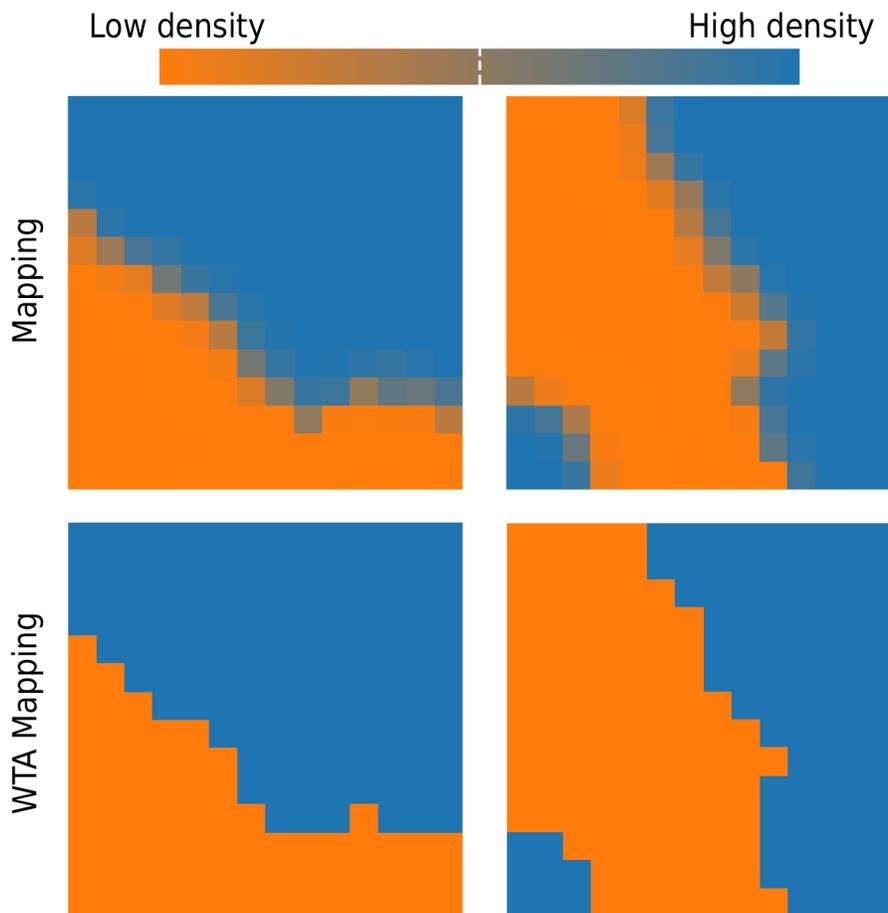


Figure 5.4: Example weight and winner-takes-all maps. Example maps and WTA maps. Each column shows one map example for density ratio 1:6. In the left column, the mapping is split into two clear clusters representing each region. However, discontinuities can occasionally emerge, for example, in the right column, where the high-density region is split into two clusters. **Top row.** Mapping based on the weights, colourbar indicates whether a unit is more strongly connected to either the low density (orange) region or high density (blue) region. Along the borders between clusters, units can have mixed selectivity. **Bottom row.** Resulting winner-takes-all mapping based on the weight strengths. Border units are claimed by the region they are weighted more strongly towards.

slight advantage at borders, such that the increased total weight makes the higher density region more often the ‘winner’ for these units.

5.3.2 Problems of implementing density via changing the stimulus statistics

In Ritter and Schulten (1986), the density increase at the digit tips was modelled by increasing the sampling of this region. Here we demonstrate that this assumption would be valid only in the case of sufficiently small stimuli. However, when using stimuli of these sizes, the assumption leads to disordered mappings and a lack of representation for the modelled ‘lower density’ region.

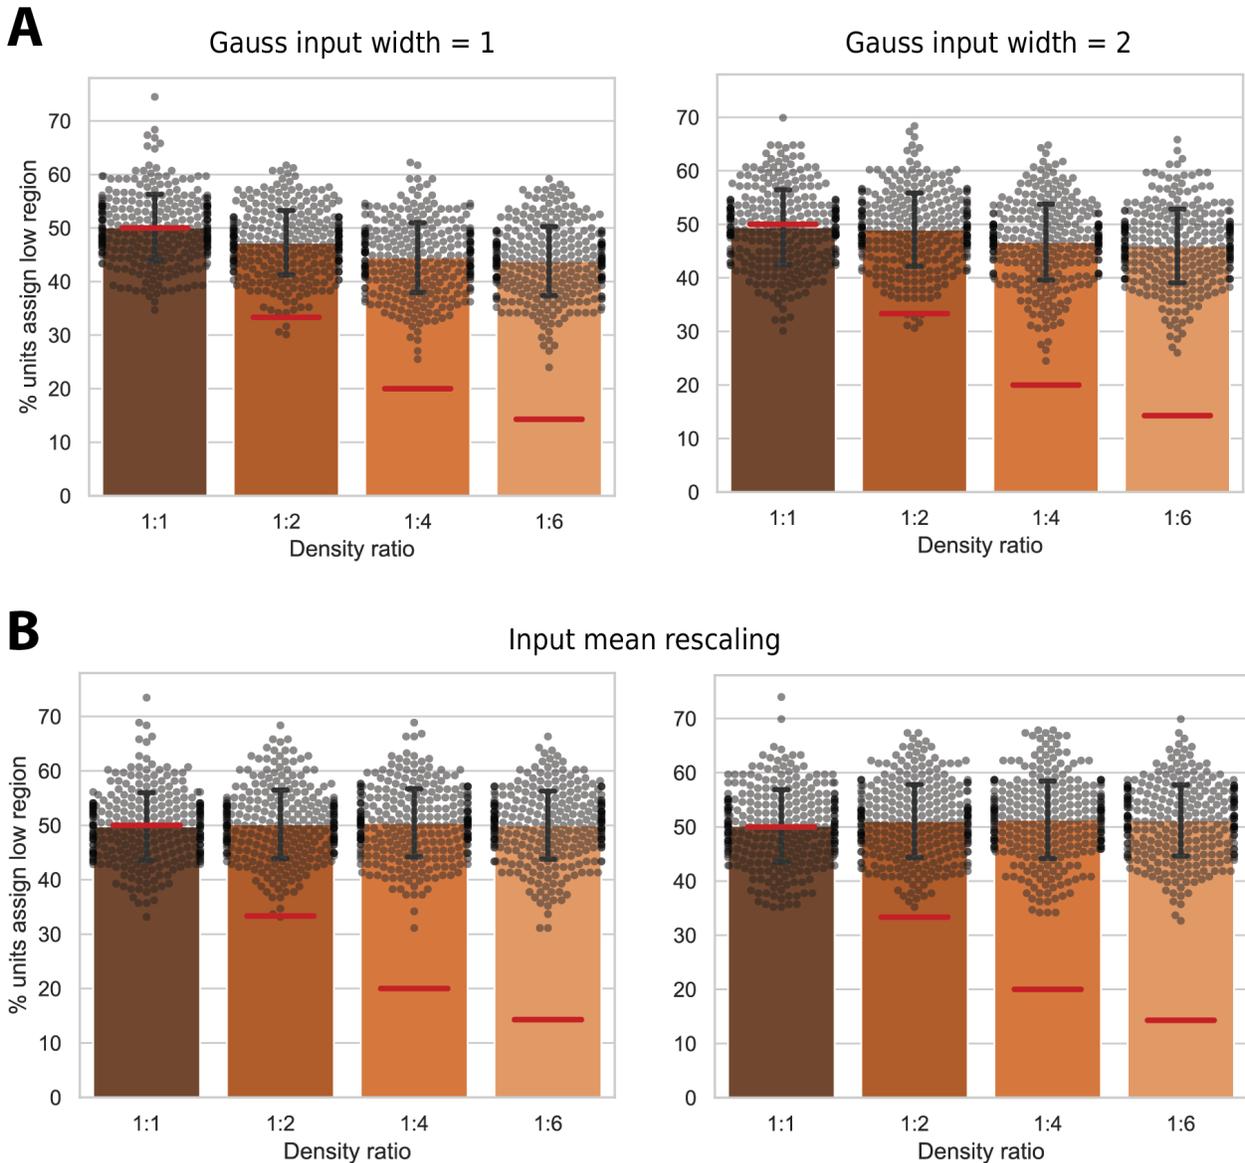


Figure 5.5: Allocations of map units (WTA) for increasing densities. **A.** Percentage of map units allocated to the lower density region for increasing density ratios. Results for the two Gaussian stimuli sizes are shown (left, smaller size, and right, larger stimuli size). Error bars indicate standard deviation. Red line on each bar denotes idealised allocations of the lower density region according to the density ratio. **B.** As above but with all inputs mean normalised.

If the stimuli sizes are of a small enough width, it is possible for single receptors only to be active. At a low enough size, the full width of the stimulus extent falls between the receptors, such that it is undetectable and does not strongly activate either (see Fig. 5.6A). Modelling a change in the density through a change in the statistics therefore makes the assumption that stimuli are small enough such that their random placement is only detected by receptors for a certain frequency of the total stimulation. The resulting input set would have an equal number of inputs for both regions; however, there would be approximately zero activation for any of the receptors in a certain

percentage of those for the lower density region. For example, for a lower density region, having approximately 75% of the input patterns zero (from the low spatial acuity and non-detection of placed stimuli) would lead to a 1:4 statistics ratio. This is further demonstrated in Fig. 5.6B.

Furthermore, of those active receptors, typically, only one is active in the lower density region at any time and few in the higher density region. In the SOM algorithm, this is problematic for two reasons. Firstly, the WTA algorithm selects and updates the maximally active unit. An input composed of zeros would lead to no update taking place. Second, Kohonen SOM relies on correlations in the inputs to create a consistent topographic mapping. When only single receptors are activated, this leads to discontinuities in the mapping and a large over-representation of the higher statistics region. In higher ratios, the lower density region can be abolished completely. See Fig. 5.6C.

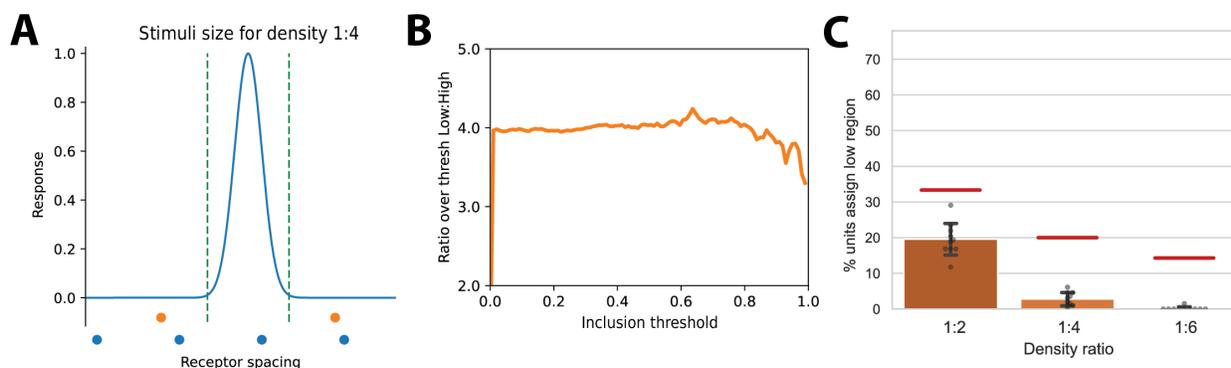


Figure 5.6: Example of density as statistics assumption. **A.** Calculation of input set and extent of Gaussian to recreate the effect of the statistics. Here, the density ratio is set to 1:4, with higher density points shown in blue (see bottom). Green dashed lines denote where the Gaussian approaches zero. With a small enough sized stimulus, the lower density points will be stimulated around 25% of the time. **B.** Ratio of total inputs for each region with at least one receptor over a threshold (see x-axis). Input set is constructed with Gaussian stimuli of the size shown in panel A. Ratio between the number of qualifying inputs in each region is approximately 1:4 for most thresholds. Ratio reduced towards a high threshold due to random placement of the stimulus, meaning the closest receptor may be further away and therefore activated less. **C.** Running the SOM algorithm for input sets of increasing density. Each stimuli size is adjusted such that the ratio of the activated inputs is equal to the required ratio (as in panel B). A total of 10 simulations were run. The small stimuli result in a lack of correlation between receptor activation. This leads to a disordered mapping, where low-density units are not clusters together. Furthermore, the statistics ratio results in the modelled higher density region having more map units than expected. The lower density region's representation is reduced (see blue and orange bars), and can be completely abolished for high enough densities (see modelled density 1:6).

5.3.3 Manipulating the statistics ratio and receptive field sizes

Next, we investigated the effect of changing the stimulus statistics. First, we show that the proportion of units allocated to each region closely follows the probability distribution of the inputs for each region. Second, we investigate the receptive field size of the map units, given the different contact ratios and two different stimulus sizes (as in previous sections). Here we demonstrate that the larger stimuli lead to larger receptive fields; however, there does not appear to be any relation-

ship between increasing the magnification of a region (over-representation of the higher statistics region) and the RF size. This is unlike real somatosensory maps, which demonstrate an inverse mapping rule, with more magnified regions having smaller receptive fields.

Map representation sizes follow statistics ratio

We first investigated the sizes of the two regions under a change in region sampling statistics. Here the activation ratio was set at one for the baseline region and increased for the second region to either two, four or six times. We tested two Gaussian stimuli sizes for each ratio, as in the previous investigations. All other map parameters were set as before, see table 5.1.

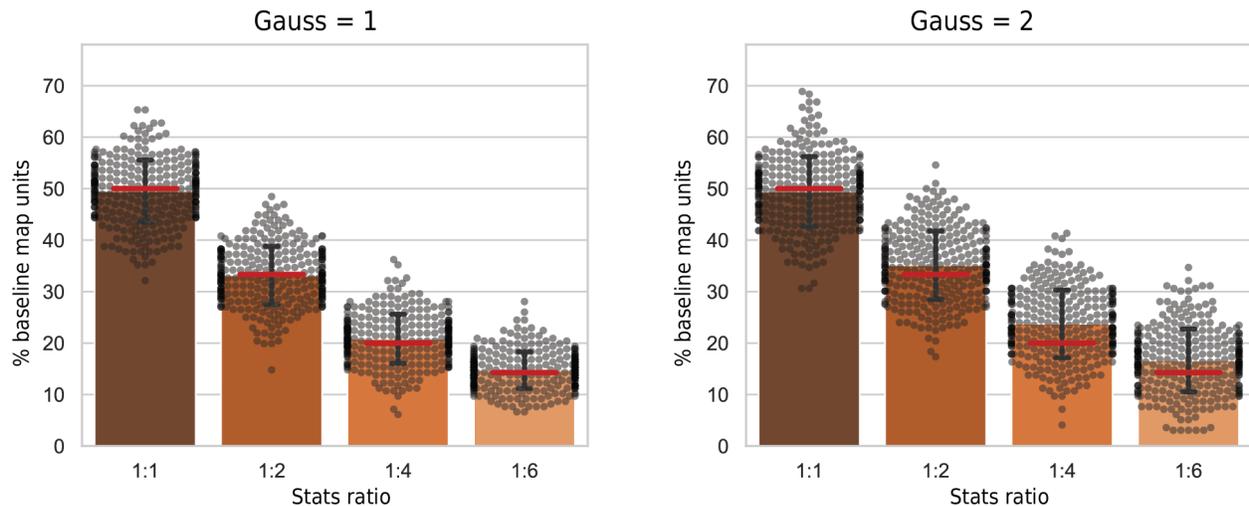


Figure 5.7: Changing sampling probabilities over regions. As in 5.5, plotted is the number of units assigned to the baseline (lower statistics) region. Bars denote increasing statistics ratios for stimuli with Gaussian stimuli sizes of 1 and 2, left and right, respectively. Red horizontal lines denote mapping according to the statistics ratio.

We found that the WTA mapping closely followed the probability density of the inputs for the two regions, with the percentage of map units approximately equal to the statistics ratio, see Fig. 5.7. However, some slight variation can be seen as a result of changing the size of the Gaussian stimuli, with the wider stimuli (Gauss size = 2) having a slight increase in representational allocation for the lower probability region.

The higher probability region is more magnified than a proportional mapping, according to its size. This approximately follows the probability of the sampling (statistics ratio), based on the mean allocation over all map simulations (see bar heights in Fig. 5.7). However, within each ratio, variations occur between each mapping (shown via the standard deviation and range of scatter points in Fig. 5.7). This variation, alongside differences due to the stimulus size, demonstrates the difficulty of finding a general magnification factor for the 2D SOM.

No inverse magnification for receptive fields

Previous work has assessed the receptive fields of the SOM units (Ritter et al., 1992; Obermayer et al., 1990b). The weights connections of each map unit denote regions of the input space to which the unit is responsive. The RFs have a strong response at their centre, and the weights to other receptors decay to zero outside the RF zone. There are several different ways to quantify receptive fields in self-organising models. Some previous work has not considered the extents of the receptive field, rather the location only (Ritter and Schulten, 1986; Obermayer et al., 1990b). Others have quantified the RF using each unit's 'scattering circle'. This is defined as the average distance between the unit's connection centre of gravity and each input receptor, weighted by their connection strength, and multiplied by π (assuming circular RFs) (Ritter et al., 1992). Units with larger receptive fields have more distant units from the centre of gravity, accompanied by stronger weight connections, resulting in a larger value for the RF size. However, this approach makes several assumptions. First, the RFs are always circular, and second, only one circular RF can occur for each map unit. There cannot be multiple centres of receptive fields within or between regions. In contrast, empirically measured somatosensory receptive fields can be of any shape, and some multi-digit RFs have been found (Merzenich et al., 1984). The form of the RFs can therefore be challenging to measure in models (Obermayer et al., 1990b).

Here we calculate the receptive fields without these previous assumptions. RFs can be of any size, location and can be made of multiple areas of activation both within and between regions. One thresholding parameter is used to determine whether a weight connection (and therefore an input receptor location on the sensory input sheet) is included within the RF. Those receptors with weight connections over a set threshold were included in the RF for each unit. The size of the RF is given by the extent of the region surface which has above threshold weights. Here we determine the receptive field extents by calculating the spatial convex hull (smallest possible polygon enclosing all points) of the receptor weights over the threshold for each map unit. An example of convex hull fitting for RFs can be seen in Fig. 5.8. See below for details on the setting of the threshold.

RFs in the model typically fall on a localised area of a single region; however, there were some occurrences of multiple within region RFs and between regions. For multiple, within region RFs, K-means clustering is first applied to create segregated groupings of points before calculating the convex hull for each cluster. Similarly, when the unit has RFs on both regions, these are calculated using separate convex-hull calculations.

The convex hull RF calculation can be slow for maps with many units. However, when the density of points is equal for both regions, the calculation of the full RF extent does not need to be made. Instead, the number of receptors over the threshold for each unit can be summed. We tested the correlation of this proxy measure with the full RF calculation as outlined above. For 30 randomly selected maps of each statistics ratio and Gaussian size, an r^2 of 0.98 ($R^2 = 0.77$), was found between the full measure and the threshold measure. This proxy measure was used for the following calculations. In the case of uneven densities, the complete convex hull calculation would

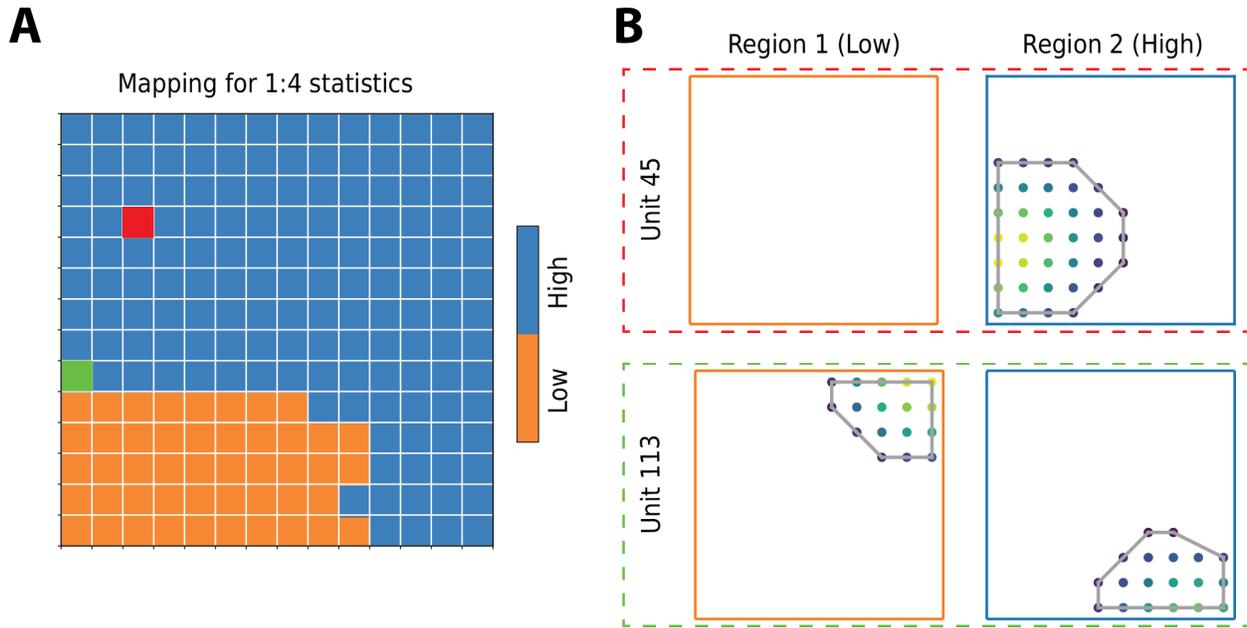


Figure 5.8: Example of cortical map and receptive fields. **A.** Cortical map example for statistics ratio of 1:4. Darker blue shows low statistics region, lighter blue, higher. **B. Example receptive fields.** Example receptive fields for two units, shown on the map as red and green squares, respectively. Top, map unit 45, showing a single receptive field over the higher statistics region (shown in red). Only units within the receptive field are shown; colour brightness denotes strength of weight to that receptor. Grey outline around receptors denotes convex hull. Bottom, RF for unit 113, shown in green on the map in A. Here the border unit demonstrates multiple receptive fields, covering both regions.

be required, or a correction to the proxy measure to account for the difference in receptor packing.

Choosing threshold of the receptive fields

The RF threshold determines whether a receptor is included in a map unit's RF. The threshold is applied to the z-scored map weights, with the same threshold applied to all maps. The receptive field measurement includes all receptors with a weight connection equal to or greater than the threshold. To decide a suitable threshold, we first assumed that all units have an RF. If the threshold is set too high, some units have no weights above the threshold, resulting in an RF of size zero (no RF). With this assumption, an upper limit was calculated for the threshold of approximately 1, see Fig. 5.9A. Similarly, we assume the RFs should not cover close to the whole region, especially for smaller sized Gaussian stimuli, providing a lower limit, Fig. 5.9B demonstrates average sizes of the RFs for increasing thresholds over two Gaussian stimuli sizes. As the thresholds increase, the RFs become smaller; however, the relative difference in RF extent between the stimuli sizes remains relatively constant. We, therefore, chose the threshold to be between 0 and 1, at 0.5. Based on Fig. 5.9B, we note that the size of the RFs is affected by choice of threshold. This affects the interpretation of the size of the RF relative to the original stimuli size (for example, whether it is a similar size, as suggested by Obermayer et al. (1990a)); however, relative differences in RF size between stimuli remain comparably similar.

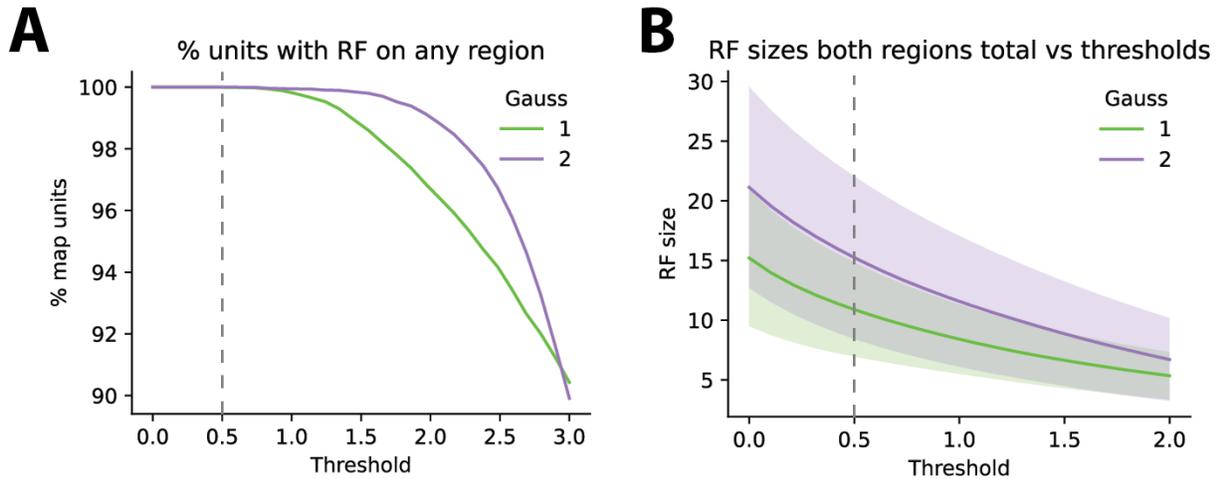


Figure 5.9: Threshold choice for receptive field inclusion. **A.** For different increasing thresholds, we examined how many of the map units have an RF on any region. Here, above the threshold of 1, some units no longer have receptive fields. We assume all units have an RF; therefore, the threshold was set at 0.5, where all units have an RF. Percentages are calculated for each threshold using 30 example maps of each statistics ratio (1:1-1:6). This was done separately for each Gaussian tested, either size 1 or 2. **B. Example RF sizes for increasing threshold.** For each threshold, the mean sizes and standard deviations were calculated, pooled over several statistics ratios. Here comparisons for Gaussian sizes of 1 and 2 are shown. Lines indicate the mean, and coloured surround, the standard deviation. RF sizes decay non-linearly with an increasing threshold; however, the relative difference between the two Gaussian stimuli sizes almost linearly decays.

Results for different sizes RFs over each Gaussian size and parameter set

Comparing between stimuli sizes, we found that larger sized stimuli result in larger sized receptive fields. This would be an expected pattern, given larger stimuli mean greater co-activation of neighbouring receptors. These are then likely to be updated to more similar strengths for that unit in the weight update.

When manipulating the input statistics ratio of the regions, two patterns emerge. Firstly, the mean RF size tends to increase with the increasing ratio for the lower probability stimulation region. However, the variation in sizes also widely increases, shown by the large standard deviations (for example, ratios 1:4 and 1:6 ratios in Fig. 5.10). For the higher probability region, the size of the receptive field has a minimal decrease with increasing probability; however, these are broadly similar across the statistics increase. Given the difference in number of samples for the RF calculation increases in line with the statistics, this may explain the large standard deviations for the lower region. Therefore, whether there is an actual difference between the RF size for the regions is difficult to establish.

It would be expected that the sizes of the RF do not follow the sizes of the stimuli exactly but should be underestimated, possibly due to border effects in the map (see the example of RFs in Fig. 5.8). The RFs are only calculated over the true region space; therefore, the edges may have smaller, non-circular RFs, which will lower the RF size.

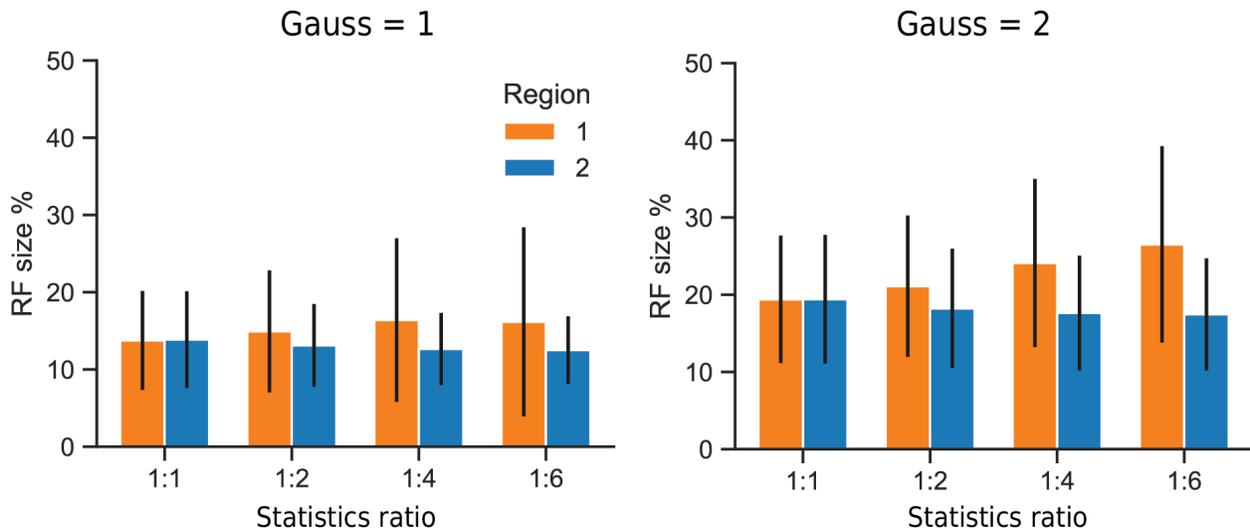


Figure 5.10: Receptive field sizes over different regions. Receptive field sizes as a percentage of region size for low stats and higher stats ratio regions (region 1, orange, and region 2, blue, respectively). Error bars denote the standard deviation of sizes. Each plot shows different stimuli sizes, left for the smaller size, and right plot, larger size. Map unit RF sizes are generally bigger for larger stimuli. For even sampling between the regions (1:1), RF sizes are approximately equal over both regions, as expected. When the statistics ratio increases, RF sizes for region 1 (lower statistics) increase, alongside the corresponding standard deviation. This pattern is consistent for both stimuli sizes. For the higher statistics region (region 2), sizes are more stable. There is little decrease in the RF size as the statistics ratio increases.

5.4 Discussion

In this chapter, we investigated the suitability of the Kohonen SOM as a model of somatosensory mappings. We studied changes in the maps when manipulating either the density ratio of receptors or the sampling between two independent 2D input spaces. Most previous modelling work has focused on replicating the effects of cortical plasticity on the mapping and has not studied factors contributing to the initial set-up of the mapping in great detail. Few considered that the densities of regions are different across the hand, even though this is a plausible mechanism behind the relationship between higher densities and greater cortical magnification. One study that implemented the difference did so as a change in sampling over the input space, with regions of typically higher densities sampled more (Ritter and Schulten, 1986). However, the effect of higher densities and higher sampling is not interchangeable. Furthermore, this study was focused on more qualitative aspects such as topography and changes under digit amputation; therefore, there was no systematic study of this assumption.

This work demonstrated that the Kohonen SOM mappings are affected by the sampling from different regions (via changes to the statistics ratio) but demonstrated little change when the input region density of receptors was adjusted. We found only a small deviation from equal-sized map representations of the two regions in this case. This was driven by the overall increase in the amount of activation. When there is a denser packing of input receptors, the total amount of

weight connections are increased, and more of these are active during a stimuli presentation than in less dense regions (assuming stimuli are a uniform size). The increase in magnification of the higher density region is not on the scale of the expected increase shown for empirical somatosensory cortex measures (under the assumption that the densities of receptors somewhat influence this). For example, the densely packed fingertips have a much larger representation than other lower density body parts, even though these are typically larger in physical size. Although magnification increases alongside the density ratio, even for a larger disparity of 1:6, there was only a 6% gain for the more dense region over a 1:1 density ratio.

For manipulating the statistics of region activation, we found a close match between the probabilities of contact and the number of units assigned to each region in the map. Previous work has suggested that Kohonen SOM under-represents the higher probability region in favour of lower regions. In 1D, under certain assumptions, this relationship is $M(w) \propto P(w)^{2/3}$ (Ritter and Schulten, 1986). Here we demonstrate the mapping was (on average) almost directly proportional; however, randomness in both the initial map weights and input presentation order could lead to non-proportional over or under-representation of regions. Given that other research has proposed adaptations and corrections to the proposed $p = 2/3$ exponent under certain assumptions, our results could be due to the setup of the input space (as two separate regions with no overlap), choice of BMU calculation, or dimensionality of the input and mapping space. As our regions were completely separate, co-activation of receptors only occurred within the regions but not between them. This may lead to a larger shared border region in the map, as weight connections into map units are more distributed between the two regions. A simulation with directly adjacent regions and overlapping stimuli may more closely follow the alternative exponents.

The receptive fields of each map unit demonstrated some properties previously seen in studies of the Kohonen SOM algorithm and real somatosensory maps, such as dual receptive fields covering multiple regions, particularly along the borders between representations (Obermayer et al., 1990a; Merzenich et al., 1984). When increasing the statistics ratio between regions, we found that the size of receptive fields for the lower statistics ‘baseline’ region increased, likely due to a smaller number of units allocated to represent the full space. However, the receptive field size does not shrink for the higher statistics region as the ratio between regions increases. Instead, it remains relatively constant despite the increase in the number of maps units representing the region. Therefore, it does not follow the magnification relationship whereby larger representations have smaller receptive fields. We found that the receptive field size in higher statistics regions depended on the stimuli’ size, and therefore the extent of the correlations in the inputs. The relationship between RF size and stimuli size has previously been demonstrated by Obermayer et al. (1990a); Ritter et al. (1992), who also found RF extents were approximately the width of the stimulus. Further consideration of the typical sizes of the stimuli may be required to produce this relationship in Kohonen SOM. Typically, higher density regions, such as the digit tips, offer increased spatial resolution and contact smaller stimuli. It may be possible to replicate this effect in SOMs by decreasing the size of the stimuli, alongside

an increase in the statistics ratio. However, it is clear that the increased statistics ratio would not be the driver of this effect but the spatial extent of the stimulation for each stimulus presentation. There are some differences between this assumption and empirical somatosensory findings. Firstly, correlations for the extent of stimulation across a region have not been well clarified. Although the digit tips are used for small stimuli, large objects can also stimulate them. For example, during grasping, the whole distal fingertip pad can be contacted (Gonzalez et al., 2014). The extent of co-activation for a mix of these stimuli sizes and a possible relationship between the decrease in co-activation and increase in statistics is unclear. We demonstrated that reductions to the size of the stimuli are also limited to ensure a smooth mapping. Uncorrelated receptor responses due to small stimuli can lead to massive under-representation with discontinuities in the mapping or complete abolition of region representations.

If RF sizes follow the stimuli size, then the smallest RF size is limited by the smallest applied stimuli size. This has consequences for interpreting plasticity manipulations. For example, (Jenkins et al., 1990) demonstrated that increased stimulation to a digit results in expansion of its cortical representation whilst shrinking the digit's corresponding RFs. The applied stimulation was confined to a small region, typically at the top of the digit tip. Given that the fingertip is used in everyday interactions, it is already considered a higher contact region. In this case, implementing the effect in Kohonen SOM would require the plasticity manipulation stimuli to be smaller than those used to train the initial map. The map would therefore show the relationship; however, the effect would be due to a reduction in stimuli size, not an increase in statistics. One alternative possibility could be proposed based on the map size; however, we did not directly test the effects of reducing the total number of units on receptive field size in this study. Given that the lower contact regions have fewer assigned map units and larger RFs, it is possible that a smaller overall map size would lead to larger RFs for all regions. In this case, the RF could become smaller if one region is stimulated and 'invades' the map units representing another region. However, the limit of the smallest RF is still dependent on the size of the stimuli. Similar changes would be required for an amputation manipulation. In Obermayer et al. (1990a); Ritter et al. (1992), an amputation manipulation was investigated, where the third digit was removed and the neighbouring digits 'invaded' the representation. However, receptive fields were only measured before the stimulation, and unfortunately, no comparison was made to test if the enlarged representations had shrunk their RFs, or if these remained constant.

In the maps, variation is seen in both the number of units representing each region and the topographic arrangement of the regions with respect to each other. Typically, two clusters of map units emerge, one representing each region. However, some variations can be seen, such as multiple discontinuous clusters of single regions. As we modelled the two regions as independent, we do not consider their topography in relation to each other in this study. Therefore, the map units along the border will randomly map the orientations of the two regions. There is no information on the orientation of regions with respect to each other, and therefore their representation could be aligned

along any of their sides. As a result, a larger variety of map organisations of the two clusters may have emerged than if some co-activation occurred between them. In somatotopic representations, ordering between regions is more stable; however, this may be influenced by early map development alongside refinements due to the input structure and co-activation of body parts.

Limitations of the SOM model and other considerations for a somatosensory model

In comparison to chapter 3, we found no effect of the bottleneck size (here number of map units) in the preliminary map size investigation. Regardless of the number of units being either smaller (a bottleneck) or larger than the number of receptors, the allocation of units between the regions was largely stable. The exact cost function being optimised in the Kohonen SOM is unclear, such that an analytic solution cannot be found for 2D; therefore the two methods are difficult to compare. For example, the differences may be due to constraints on the topography of organisation in Kohonen SOM. Given that key features of somatosensory cortex organisation could not be replicated (such as inverse magnification of receptive fields), it is perhaps unlikely that the cost function being optimised here is the same as that in sensory cortices.

The Kohonen SOM is an abstract model of the complexities of cortical mechanisms. In particular, network competition through the selection of a single winning unit to update can be likened to a global supervisor and is not biologically plausible. Synaptic weight updates occur locally, with neuronal activity mediated by interactions between lateral connections. In the SOM model, these are abstracted by using the neighbourhood size. Explicit lateral connections have since been implemented into models of somatosensory cortical plasticity (Grajski and Merzenich, 1990) and visual feature maps, such as ocular dominance and orientation tuning (Bednar and Wilson, 2016; Bednar and Miikkulainen, 2006). In the visual models ‘LISSOM’ (Sirosh and Miikkulainen, 1994) and ‘GCAL’ (Stevens et al., 2013), the competitive best matching unit of Kohonen SOM was replaced with lateral connections, with short-range excitation and longer range, weak inhibition to more distant units. This creates a localised bubble of activation on the map before unit updates. These lateral connections are also plastic and are pruned alongside the afferent connectivity from the input sheet.

The Kohonen SOM features one set of weights between the input and the map. It, therefore, does not consider processing along the pathway before reaching the cortical SOM map, for example, that topographic maps occur in both the cuneate nucleus and the thalamus (Qi et al., 2011). Finally, the SOM only considers the Hebbian style plasticity; however, cortical circuits are also affected by homeostatic plasticity mechanisms (Turrigiano et al., 1998), which act to regulate activity within the network with certain bounds. In the following chapter, we will further investigate possible contributions of homeostatic plasticity and patterns of cortical connectivity by replicating the results of a short-term plasticity manipulation in a mechanistic model of plasticity.

Chapter 6

Plasticity of the hand representation

6.1 Introduction

Cortical plasticity is the brain's ability to reorganise by forming new connections between neurons, modifying existing connections or changing neuron responsiveness (Mateos-Aparicio and Rodríguez-Moreno, 2019). It has been demonstrated as a ubiquitous property across sensory regions in the brain (Marik and Gilbert, 2017), including in the somatosensory (Florence et al., 1997; Merzenich et al., 1978), visual (Wiesel and Hubel, 1963, 1965; Karmarkar and Dan, 2006; Gilbert and Li, 2012), and auditory cortices (Recanzone et al., 1993; Rauschecker, 1999; Pantev et al., 1999; Schreiner and Winer, 2007). In somatosensory cortex, plasticity has implications for learning and development (Wang et al., 1995; Dempsey-Jones et al., 2016), modification of tactile experiences through human-machine interactions (Dominijanni et al., 2021), and clinical applications, including prosthesis use and development (Delhaye et al., 2016).

Plasticity occurs across the lifespan, with sensory inputs in early critical periods refining coarse maps set up through genetic processes (Flanagan, 2006; Wiesel and Hubel, 1963) (see background chapter, section 2.4.1). In adulthood, representations can be fine-tuned by retaining their ability to adapt to changing sensory input. In the somatosensory cortex, the general global topographic organisation (relative positioning of the body parts) stays mostly fixed over time (Penfield and Rasmussen, 1950). At the finer level, decades of work using animal models and human observations has demonstrated that neuron populations can change their selectivity, particularly to neighbouring body parts (Merzenich et al., 1978; Wang et al., 1995). Therefore, fine-grained body representations are not fixed but can be continuously modified to adapt to changes in input throughout life (Merzenich et al., 1983). It has been suggested that this continuous peripheral input may be an essential condition for maintaining normal somatotopic body organisation (Rossini et al., 1994).

In the previous chapter, we discussed the initial setup of the map and the possible driving factors. In this chapter, we explore changes to the somatosensory map under tactile sensory input manipulations. We consider potential contributions of the network connection setup, input structure and

different types of plasticity to a short-term anaesthetic digit block in humans. We model changes in somatosensory response before and after the block, replicating empirical measurements from fMRI. In the following sections, we provide an overview of empirical plasticity studies in somatosensory cortex, and discuss previously proposed mechanisms for these changes. Furthermore, most previous studies considered only Hebbian plasticity mechanisms, therefore we discuss possible contributions of homeostatic plasticity.

6.1.1 Plasticity in the somatosensory system

The classic textbook model of plasticity comes from studies of digit amputation and nerve resections by Merzenich and colleagues (Merzenich et al., 1984; Merzenich and Jenkins, 1993; Florence et al., 1997), where neighbouring digit representations invade cortex previously selective for the removed digit's input. Further study has demonstrated that the extent of changes, such as reorganisation, is dependent on both the type of manipulation and the timescales over which effects are studied.

Long-term tactile sensory input manipulations

Most plasticity studies using non-human primates have investigated longer-term effects of plasticity on timescales of weeks to years. The most common manipulations involve the removal of input through amputations and nerve deafferentation. However, some studies have also investigated digit syndactyly (merging of digits) and cortical lesioning (Merzenich and Jenkins, 1993).

In amputation studies, one or multiple digits of the hand are surgically removed, effectively silencing their inputs. In cases where individual digits are removed, neighbouring digits expand topographically into the territory of the amputated digit whilst maintaining their original representation (Merzenich et al., 1984). These invaded representations also maintain correspondence to typical properties of area 3b, with the enlarged representations demonstrating a shrinking of the receptive fields. Reorganisation from amputation is typically localised within cortical zones, not stretching far beyond the original territory. For example, neighbouring digits and the adjacent palm region can invade when a digit is amputated (Merzenich and Jenkins, 1993).

Studies of deafferentation have shown that the pattern of cortical reorganisation depends on the extent and pattern of input loss (Garraghty et al., 1994). When a nerve is cut or crushed to prevent regeneration, more extensive input regions are lost across the hand than for a local amputation. For example, deafferentation of the median nerve leads to loss of input from digits 1-3 of the glabrous skin. For the ulnar nerve, input from digits 3-5 of both the glabrous and hairy skin is lost. By transection of both the median and ulnar nerves, Garraghty and Kaas (1991) found that the entire region corresponding to the glabrous skin became responsive to the dorsum of the hand instead. In contrast, when multiple digits are amputated, areas of the deprived cortex do not fully recover (Merzenich et al., 1984). Differences between the two are not due to the method of input removal, but the pattern of removed inputs (Garraghty et al., 1994). In both cases, widespread inputs were removed, however, in the case of deafferentation, the input from the radial nerve was still intact, allowing remapping of the hairy skin. Deafferenting in a pattern that mimics full amputation

(removing both glabrous and hairy skin input) led to no remapping of still innervated hand regions (Garraghty et al., 1994). These findings suggest that there may be some limits to the extent of reorganisation.

Cortical lesions of digit representations can also lead to the reorganisation of the surrounding regions. For example, lesions of the third digit representation resulted in the adjacent intact regions, which are usually responsive to neighbouring digits, becoming responsive to the third (Jenkins and Merzenich, 1987). This effect is highly local, as larger lesions of multiple digit representations cannot reorganise the map to maintain topography.

Rather than complete removal of inputs, cortical changes also occur when the structure of input is changed. In digit syndactyly, multiple fingers are sutured together to create a ‘webbed finger’ (Allard et al., 1991). Following surgery, the digits are forced to work as one unit, meaning that inputs are delivered to both fingers near simultaneously. Three to seven months after surgical connection of digit three and four Allard et al. (1991) found the typical map discontinuity of the digit representation were abolished. This combined digit representation led to RFs extending across the syndactyly division line that previously separated the two digits.

Alternative manipulation of input is through changing the structure or frequency of stimulation over different digits. When increasing stimulation of a skin region, for example, the digit tips, corresponding increases in representation and decreasing RF sizes can be seen in cortical representations (Jenkins et al., 1990; Recanzone et al., 1992). This suggests that changing hand use and subsequent differences in the stimulation distribution may affect cortical magnification. The exact timing and combination of tactile stimulation across the digits may also affect representations. Wang et al. (1995) demonstrated that simultaneously stimulating distal or proximal phalanges of multiple digits resulted in more integration within these multi-digit representations, with multi-digit RFs. In contrast, more segregated digit representations were found after asynchronous stimulation. This is strikingly different to healthy monkeys, where there are some two-digit receptive fields, but virtually no three-digit ones (Merzenich and Jenkins, 1993). This pattern of changes suggests a relationship between the use of the hand in tactile interactions and the resulting cortical representations.

Studying the effects of plasticity in humans is more limited, as invasive manipulations are difficult or not possible. However, the development of fMRI has enabled investigation with amputees (Makin et al., 2015) and specialist hand users. For example, early studies in the 90’s found cortical magnification of certain digits after long term training in musicians and braille readers (Elbert et al., 1995; Pascual-Leone and Torres, 1993).

The large-scale cortical reorganisation found in NHPs has been challenged in human fMRI studies through the use of newer multivariate analysis techniques, particularly representational similarity analysis (RSA). This has demonstrated that somatotopic representations may be more persistent than previously thought, with underlying representations remaining stable regardless of peripheral input changes (See Makin and Bensaïa (2017) for a detailed review). RSA compares the similarity

of activation given different stimuli, enabling comparison over larger areas of cortex than traditional univariate approaches, which focus on individual digit representations. The production of a single ‘dissimilarity matrix’ comparing the different inputs enables direct comparison to other imaging modalities such as EEG, computational models and electrophysiology (Kriegeskorte et al., 2008). In somatosensory fMRI studies, RSA is used to compare voxel activation for the hand area given different tactile stimuli, for example tapping of individual digits. A canonical dissimilarity pattern is produced in non-clinical populations, with dissimilarity between pairs of digit representations increasing with distance (see Fig. 6.5B for an example). For example, the index and middle digit typically show higher similarity than the index and little digit. Recent studies using RSA have demonstrated that upper limb amputees may retain some representation of the digits in S1 even decades after amputation. This suggests that somatosensory representations may be more persistent than previously expected, and continual sensory input may not be required to maintain it (Kikkert et al., 2016). Instead, RSA has identified that the underlying representation of digits may persist even many years after loss of the limb in amputees (Wesselink et al., 2019). In comparison, congenital one-handers demonstrate no ordered topography or representation with RSA. In this case, it is possible that some input is needed to develop the representation in order for it to persist after limb loss, for example during foetal development or a critical period.

Short term tactile sensory input manipulations

Changes to somatosensory representations have also been found for shorter timescale manipulations. This includes behavioural manipulations (Dempsey-Jones et al., 2016), reversible digit syndactyly (Kolasinski et al., 2016) and nerve blocks to mimic amputation (Calford and Tweedale, 1991).

Amputation has been mimicked through short-term blocking of digit inputs using local anaesthetic in NHPs (Calford and Tweedale, 1991). Immediately after, receptive fields of neurons from the blocked digit demonstrated expanded RFs. Once the nerve block had worn off, the RF sizes decreased, akin to before the block. Although this demonstrated similar effects to full amputation, it is unlikely full reorganisation of connections occurred within the short time. Therefore it is possible that multiple timescales of plasticity mechanisms exist. Similar results have been reported in rats under lidocaine injections to the whiskers (Faggini et al., 1997). In both cases, it was suggested the effects are due to unmasking of non-dominant inputs from other regions rather than cortical reorganisation of neuron connections. This is further discussed in section 6.1.2.

Map changes have also been demonstrated 24 hours after reversible digit syndactyly, where digits two and three were glued together. This forces digits to work synchronously, and therefore changes the pattern of receptor activation across the hand, increasing co-activation of the glued digits. fMRI comparisons before and after syndactyly demonstrated a shifting of non-glued digit representation (Kolasinski et al., 2016).

6.1.2 Mechanisms of somatosensory plasticity

Connection patterns to and between cortical regions

Plasticity of representations, such as the ability of neighbouring neurons to reorganise, may depend partly on connections to and between different clusters of digit representation. More expansive feedforward connections may only reinforce more sparse connections to other cortical region representations rather than new connection development.

Feedforward connections from the thalamus to somatosensory representations are arranged in a topographic ordering, with diffuse projections into area 3b (Padberg et al., 2009; Garraghty et al., 1990). For example, digit-selective regions in the thalamus project to multiple digit representations in cortex, not just their corresponding cluster. For further discussion of thalamic cortical connections, see chapter 1, section 2.3.2.

There are clusters of neurons in somatosensory cortex that are typically selective for individual digits. For example, neurons in area 3b have small RFs restricted to single digits. However, lateral connectivity exists between these selective clusters, which may be the basis of the broad integration of information across the hand region. Within the 3b hand representation, digits are connected to all others and also parts of the wrist and palm, but more sparsely to other body part representations (Liao et al., 2013), such as the face (Fang et al., 2002). The connection strengths within the hand region are dependent on distance, with neighbouring digits preferentially connected with the strongest lateral connections, compared to those that are non-adjacent (Négyessy et al., 2013). Correlations between spiking patterns of neurons in different patches can occur at more than millimetre distances, demonstrating more broad functional connectivity of somatosensory representations than only single digits. Furthermore, although typical studies demonstrate small RFs using mapping techniques for minimal receptive fields, naturalistic stimuli may evoke wider more global areas of the digit representation (Reed et al., 2008). The processing of naturalistic stimuli may be aided through this lateral connectivity.

Lateral connections can perform lateral inhibition, which has been demonstrated predominately at the scale of single neurons and small patches. Inhibitory effects to surrounding neurons cause a sharpening of the response, increasing the contrast between responses to similar stimuli (Laskin and Spencer, 1979). This may be beneficial in enhancing spatial acuity and enabling low detection thresholds for stimuli. Nevertheless, the nature of the broad connections within somatosensory cortex has not been well-established. Lateral inhibition could be effective to enhance contrast; however, larger and multi-digit receptive fields exist, and the lateral connectivity could instead assist with the integration of information (Liao et al., 2013). Here representations might be connected in such a way to reciprocally enhance each other's activation rather than suppress it. Correlated activity enhancement to assist in the processing of similar stimuli has also been found in the olfactory bulb, implemented through excitatory signalling (Christie and Westbrook, 2006).

Distance and extent of changes- boundaries of the hand map

Given the typical pattern of lateral and thalamus connectivity being rather broad over the hand region, it is possible that local input manipulations can result in more widespread changes across the hand map.

What are the distances over which reorganisation typically occurs? Initial studies by Merzenich and colleagues gave a reorganisation distance limit of around 1-2mm (in NHPs), which represents the typical maximal expansion of representations and aborisation of single thalamocortical neuron axons (Merzenich et al., 1984). This leads to local-only remapping of neighbouring digits and limits more widespread reorganisation. This is supported by studies that show non-recovery for large deinnervations (Garraghty et al., 1994). However, others have demonstrated reorganisation over distance of approximately 4mm (Jones, 2003; Garraghty and Kaas, 1991), and even an order of magnitude higher than Merzenich suggested (10-14mm) (Pons et al., 1991). Here, this more widespread plasticity may have thalamic and subcortical targets such as in the brainstem, where the neural space of representations is more packed than higher regions.

Multiple timescales of changes - immediate unmasking?

Differences between the long and short-term manipulations described above suggest multiple possible timescales of reorganisation under input removal. Firstly a rapid ‘unmasking’ stage, where latent representations of non-dominant inputs can be seen. This is limited and occurs over short distances, typically affecting neighbouring digits, based on the thalamocortical and lateral connectivity pattern. A lack of lateral inhibition from the dominant inputs means that previously inhibited latent inputs are now uncovered. These representations are emergent from the normal circuitry when inputs are removed and do not involve synaptic plasticity mechanisms (Calford, 2002). Following this, a longer-term reorganisation stage may occur, which reinforces the latent activity to become more dominant. In this time, there is more extensive reorganisation followed by fine-tuning (Churchill et al., 1998), possibly driven by Hebbian plasticity mechanisms (Florence et al., 1997). Sprouting of new connections (Florence et al., 1998) or Hebbian based plasticity likely acts on longer timescales such as days to weeks (Rasmusson, 1982). It is difficult to distinguish between long- and short-term plasticity changes due to unmasking, as we can only assess cortical changes after possible plastic changes have taken place. Therefore it is hard to understand whether changes are only due to the removal of inputs and this normal circuitry or additional mechanisms.

6.1.3 Roles of the types of plasticity

Most somatosensory plasticity research has suggested reorganisation is due to synaptic (Hebbian) based plasticity. Hebb (1949) first postulated that neurons that tend to fire synchronously should form connections; this is more concisely known as ‘neurons that fire together, wire together’. Synaptic modification is the dominant idea of plasticity and the basis of changes in most plasticity models. The term synaptic plasticity encompasses several mechanisms which act at variable timescales, such as long-term potentiation (LTP) and long-term depression (LTP) (Lisman, 2017). A further form

of plasticity- Homeostatic plasticity was later identified by Turrigiano et al. (1998) and is focused on changes to neuronal firing thresholds (Lisman, 2017).

Homeostatic plasticity aims to balance the neural activity to maintain it at a stable level (Turrigiano and Nelson, 2000). This keeps neurons in a safely bounded state, forcing them to drift back to some stable activity level within the network. It is unclear whether the ‘set points’ that neurons return to are individual for a neuron or circuit wide, suggesting changes could occur over several spatial scales (Fox and Stryker, 2017). Homeostatic plasticity encompasses several mechanisms, such as synaptic scaling and alteration of intrinsic membrane properties (Lisman, 2017). Synaptic scaling is most frequently studied and enables neurons to regulate and stabilise their firing rate by adjusting their synapse strengths. This helps stabilise the activity within a network by offsetting any excessive excitation or inhibition. This can be through multiplicative scaling, where all synapses are scaled up or down by a factor to stabilise the neuron’s firing rate. In addition, current evidence suggests that neurons detect their own firing rates by detecting receptor traffic with calcium-dependent sensors. This can then be regulated through increasing or decreasing Glutamate receptors at synapses (Turrigiano, 2012). Alternatively, levels of inhibition can be altered to balance excitation and inhibition levels. This has been suggested as a mechanism immediately following input loss, where the amount of inhibition of excitatory neurons is reduced. Here a homeostatic balancing of activity is required (Orczyk and Garraghty, 2015).

The two plasticity mechanisms interact with homeostatic mechanisms regulating Hebbian (see Fox and Stryker (2017); Zenke et al. (2017) for a comprehensive review). Hebbian plasticity is in a positive state and feedback process, whereby activity increases lead to strengthening of synapses, increasing activity further. Once LTP is started, potentiated synapses undergo further potentiation more easily. This can lead to instability (Fox and Stryker, 2017; Vitureira and Goda, 2013; Turrigiano and Nelson, 2000), and runaway excitation if not controlled. Changes to synaptic efficacies may be kept in reasonable bounds by synaptic scaling. This means the differences in synaptic inputs are preserved, just scaled within reasonable bounds (see Turrigiano (2008); Pozo and Goda (2010) for rapid homeostasis).

Given the range of long-term homeostatic mechanisms, this may also occur at short timescales, with hours at the fastest (Turrigiano, 2008; Castaldi et al., 2020; Zenke et al., 2017). Therefore, there could be multiple timescales of plasticity in operation. Faster, rapid homeostatic changes are typically disinhibitory and may regulate short term changes to activity levels. A slower scaling of synapses may then occur in response to more persistent changes in thresholds. Homeostatic speeds may depend on the scale of changes, extensive changes to activity levels may be problematic for excitation inhibition balance, and faster homeostatic changes may be required.

6.1.4 Aims of this chapter

In this chapter, we investigate the plasticity of somatosensory representations of the hand to a short-term digit block, akin to previous studies in NHPs (Calford and Tweedale, 1991). fMRI

demonstrated that global changes occurred across the hand map, rather than localised to only neighbouring digits. Here we produce a mechanistic model to understand possible mechanisms that could drive the global changes. We consider the need for distributed incoming connections from upstream regions (for example, thalamic input) and lateral connections between neighbouring cortical digit representations. A set of realistic tactile input responses before and after digit block were created using a previously developed model of afferent responses to stimuli (Saal et al., 2017). This enabled comparison before and after block assuming no cortical changes, which demonstrated that additional mechanisms were required to reproduce the empirically measured effects. Due to the short timescales, we consider homeostatic rather than Hebbian mechanisms could drive the effects.

6.2 Methods

In this section, we first describe the collection of the fMRI data both at baseline (without digit anaesthetic) and after the lidocaine block for the index digit (D2). In the second part, we aim to reproduce this result in a computational model. We first use a toy model to investigate required connection patterns from the hand to cortex that can reproduce the typical canonical hand structure. Following this, we create a full model of the lidocaine block manipulation, fitted on the fMRI dataset. We identify the required changes in the model that capture the difference in responses after digit block. Somatosensory cortex responses before and after lidocaine block are compared using representational similarity analysis, a technique which is further explained below.

6.2.1 Empirical fMRI data: baseline and post-lidocaine

Baseline fMRI methods

The fMRI and spectroscopy data were collected and analysed by collaborators at University College London, and the University of Oxford, as referenced in Wesselink et al. (2020), here we briefly summarise the findings. Somatosensory responses of fifteen participants¹ were measured during tactile stimulation before and after a nerve block was applied to the index digit. Applied stimulation could be active, where the participant presses a key, or passive, where the experimenter manually stimulates the digits. As the computational model of afferent responses to stimuli can model passive stimulation only, we focus on the results of the passive condition only. However, as noted in Wesselink et al. (2020) the results were similar between the two conditions. We first discuss baseline somatosensory responses to digit stimulation before block.

fMRI scanning was conducted using a 7T-machine, providing high spatial resolution, with voxel sizes of 1mm^3 . Individual dominant cortical digit clusters (denoted C1-C5) were established using a localiser task. This enables the identification of clusters that are more selective to one digit than all other four digits and is used in both the univariate and multivariate analyses. In the passive

¹For further details on sample demographics see (Wesselink et al., 2020).

task, the experimenter stimulated each digit by tapping the distal phalanx of the finger with a plastic probe. Participant’s concentration on the digit tapping was encouraged by self-report of occasional double-tapping. 93.3% and 94.2% accuracy of the double taps for baseline and block respectively were reported, demonstrating that participants were focused on the task. Univariate analyses of the digit responses involved averaging the activity response of each digit stimulation in each cluster (established from the localiser task). Although the subdivisions of S1 could not be delineated precisely (eg. 3a, 3b etc), probabilistic mapping of these regions demonstrated that most of the activity was likely in areas 3b and 1.

Anesthetic block methods

To test somatosensory cortex responses to a reversible short-term amputation-like removal of inputs, participants received a nerve block to the index digit- hence called ‘block session’. This was compared to the baseline session results, which were described earlier. Task-based fMRI scanning was measured, alongside the tactile acuity of the digits, resting-state scans, and magnetic resonance spectroscopy (MRS) of the hand region.

In the block session, lidocaine anaesthetic was administered to form a ring block around the base of the right-hand index finger. Perceptual testing was conducted after the anaesthetic was administered to ensure it had taken effect and after all scanning and spectroscopy to confirm that the block had remained stable. Tests comparing participants’ detection of horizontal versus vertical gratings demonstrated a chance level performance for the blocked index digit (53%), down from 95% accuracy, which persisted after scanning. The block did not affect the neighbouring digits, with the middle digit remaining at high detection accuracy.

Magnetic Resonance Spectroscopy (MRS) data of 10 participants were analysed post fMRI scans in both sessions. Spectroscopy measures the concentration of metabolites, therefore imaging the neurochemical profile of a region. Here the profile of GABA and Glutamate spectra were measured (inhibitory and excitatory neurotransmitters, respectively), and absolute concentrations were extracted from a voxel of size 2x1x1cm around the S1 hand knob. The ratios of Glutamate/GABA were calculated for each participant and compared across baseline and block sessions.

6.2.2 Computational modelling methods

Modelling of the fMRI data was split into two parts. First, we used a toy model to understand the required connectivity patterns between the hand and cortex. Following this, we calculated the full model based on the collected fMRI data. Both models use afferent response data to digit stimulation calculated with a peripheral sensory model— TouchSim.

Peripheral sensory model: TouchSim

The TouchSim model (Saal et al., 2017) recreates realistic responses of three classes of afferents in the hand to modelled stimuli- SA1, RA and PC². These afferents are placed across the 2D model

²SA2 afferents were not included in the original development of the model due to lack of evidence about their

hand surface according to their realistic innervation densities. TouchSim incorporates biological details taken from experimental work with both humans and non-human primates, enabling reproduction of the afferent spiking activity over time in response to the stimuli placed across any location on the hand. The responses are determined using a Leaky-integrate and fire (LIF) model incorporating realistic skin dynamics that occur under deformation from an object on the skin. Stimuli parameters can be manipulated to increase the accuracy of the responses, including the frequency, indentation depth, size and duration of stimulation. Given the inclusion of multiple afferent types, TouchSim enables the accurate generation of typical patterns of responses across the whole hand surface to stimuli. When a stimulus is pressed against the skin, responses are not localised to the directly stimulated skin surface. Instead, ripples of activation occur across the skin surface, typically from vibrations (Manfredi et al., 2012; Shao et al., 2020, 2016), see Fig. 6.1A. These are picked up by RA afferents, which terminate in Pacinian corpuscles (see 2.1.1 for more details). The peripheral receptive fields of these afferents integrate information from stimuli placed over the hand (Johansson, 1978).

We generated a TouchSim hand model with 1600 afferents (the total number of afferents was scaled down from the full model for speed of computation). Afferent classes were placed in their relative densities across the hand according to the standard model parameters. To replicate the fMRI inputs, all five digits were stimulated on the distal phalanx, with 1000 stimuli responses calculated for each digit. A circle stimulus probe was centred in the middle of the distal phalanx with some random jitter (around 2mm) applied to produce variation in individual responses, which may naturally occur during manual stimulation. Parameters were chosen within a likely range, with an indentation depth of 1mm and a probe radius of 15mm.

Following the stimuli generation, we obtained the rate based responses generated by the TouchSim model. Afferent responses of each digit were pooled, leading to a 5x5 response matrix of responses from each digit to stimulation of each of the other digits, see Fig. 6.1B. Activation from the palm regions was excluded as there was no direct testing of this region in the fMRI data.

Toy model: Reproducing the canonical hand structure

First, we investigated the model architecture required to reproduce the canonical representation seen in many somatosensory mapping studies, whereby cluster similarities decay with digit distance (for example, D1 and D3 cortical responses are less similar than D1 and D2). We tested different feedforward weight connection patterns to identify whether these should be local only (between each digit and its corresponding cluster) or more distributed, with connections to all clusters from all digits. We establish the requirement for widespread cortical connections from this model, which are then later implemented in the full model reproducing the fMRI lidocaine block result.

We modelled somatosensory cortical representations as five units, each representing a finger-selective population cluster for each digit (denoted C1-C5, as in empirical data). We then tested two ideas

responses.

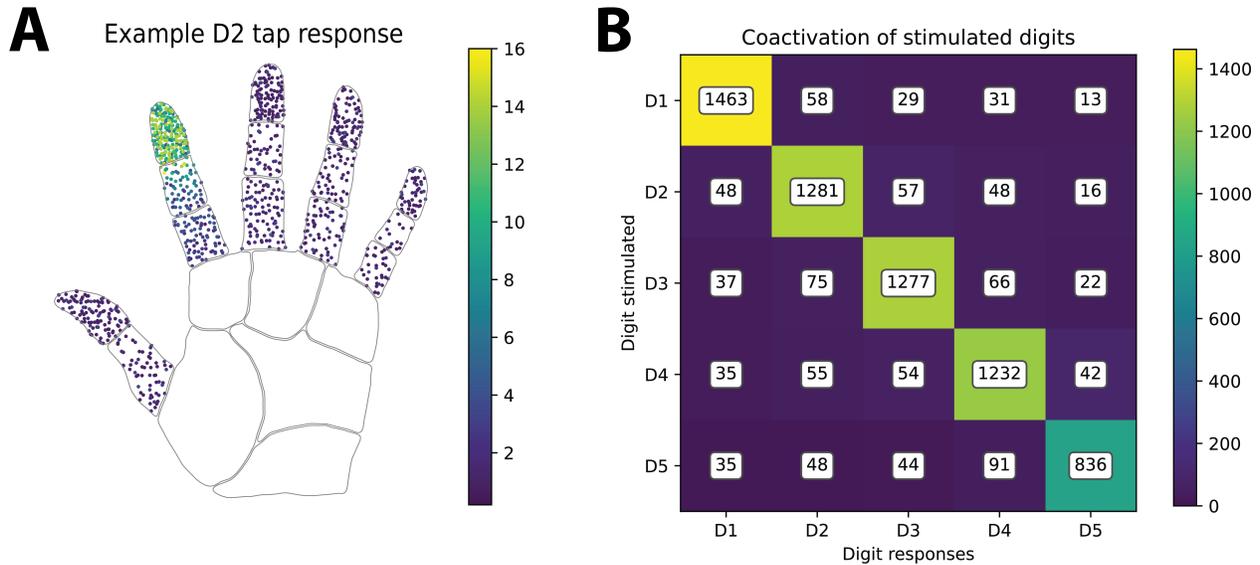


Figure 6.1: **A.** Example of a TouchSim response pattern for tapping the index digit distal phalanx. Afferent responses are strongest for skin regions where stimulation is directly applied, here, the index digit fingertip. However, activity also spreads to other digits on the hand. This typically occurs through the rippling of vibrations across the skin following deformation. **B.** Typical pooled responses across the digits for all afferents during stimulation. Rows indicate a change in the digit tip being stimulated; columns indicate the response of a particular digit to each digit stimulation. For any stimulation, the digit being stimulated has the strongest response, as expected (see diagonal). However, other digits also respond to stimulation. In the ‘co-activation’ condition, these off-diagonals are considered. In the no co-activation conditions, off-diagonals are zeroed out, and it is assumed that only the stimulated digit is active.

of how the distributed somatotopy could occur, through either ‘co-activation’ of digits during stimulation, or through divergent connectivity from the digits to cortex.

First, during passive digit stimulation, mechanoreceptors can be activated beyond the focus of the digit stimulation, due to rippling of the skin (as previously discussed in section 6.2.2). Therefore, during the stimulation of a single digit tip, other mechanoreceptors are also more weakly activated on the other digits and palm (see Fig. 6.1A). In the simulation we test whether this distributed peripheral activation can evoke responses across all digit clusters during a single digit stimulation, accounting for the broad activation seen across the S1 region. This is done by either retaining or removing the mechanoreceptor activation on the non-stimulated digits, and are termed the ‘input co-activation’ and ‘no input co-activation’ conditions, respectively. For example, when the index digit (D2) is stimulated, the activation on D1 and D3-D5 is zeroed in the ‘no input co-activation’ condition.

Second, as discussed in section 6.1.2, connections from upstream regions to somatosensory cortex are more distributed, meaning that some connections are beyond the peak response area or selective cluster (Calford and Tweedale, 1991; Garraghty et al., 1990; Padberg et al., 2009). Therefore we tested whether one-to-one connections between each digit and its corresponding cluster could reproduce the effect, or if more divergent connections are required. In the latter case, we set each

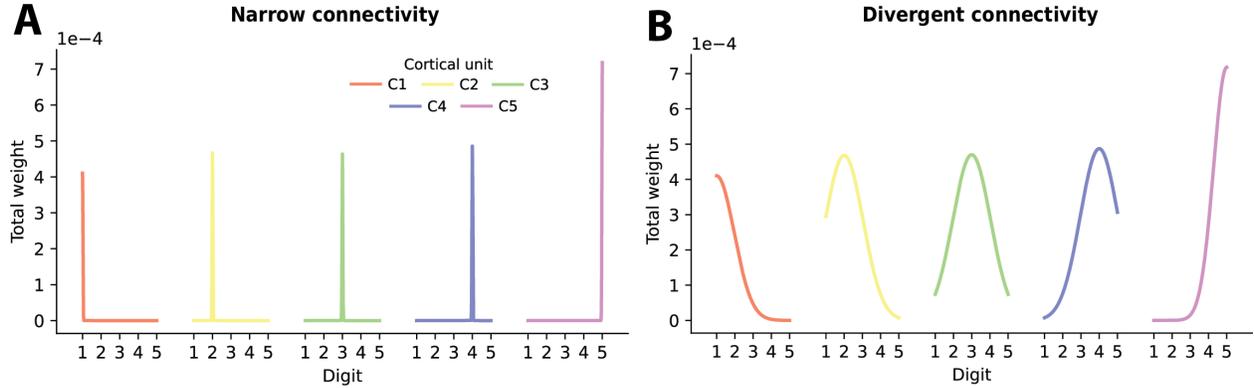


Figure 6.2: **A.** Example weight connections from the digits into each model cortical cluster. Cortical units denoted by colors. Connections peak for the corresponding digit cluster to the digit for both patterns. However, there are no divergent connections in the narrow pattern into other clusters. **B.** For the divergent pattern, digits are connected to neighbours and more distant units with smaller weight.

digit to have some weighted connection to neighbouring clusters. Furthermore, lateral connectivity can enhance digit clusters' responses and cause distributed changes to responses across the map. We therefore also tested the effects of lateral connectivity, and whether single digit to cluster patterns with lateral dynamics is enough to produce the canonical effect.

Input connections were modelled as Gaussian-shaped profiles for which the width and peak were adjusted. Due to the small number of afferents and thus lower activation of the fifth digit, we increased the amplitude of this Gaussian to give approximately equal activation to each cluster. The widths for each Gaussian were therefore either narrow with digits only connected to their corresponding clusters (see Fig. 6.2A), or wide, where neighbouring clusters are also connected—'divergent' connectivity case. Here the strength of the connection decays with distance from the digit (see Fig. 6.2B).

The cortical cluster response c_i to each digit stimulation was calculated separately, as follows:

$$c_i = \mathbf{w}_i^\top \mathbf{p} \quad (6.1)$$

where c_i is the activity of cortical cluster i , w is a weight vector specifying each digit's connection strength to the cortical cluster and p is the vector of peripheral inputs from each digit.

In somatosensory cortex, lateral connections exist between digit representations. This lateral strength is lower than the afferent activity connections and typically involves shorter range excitation and longer range inhibitory action of units further away. We tested adding lateral connectivity to the fully distributed input patterns, which enabled a sharpening effect of the response, where responses within the dominant, corresponding cluster for that digit are reinforced, and responses in other digit clusters are reduced. With the lateral connectivity the model specification is:

$$c_{i(t)} = \mathbf{w}_i^\top \mathbf{p} + \mathbf{l}_i^\top \mathbf{c}_{t-1} \quad (6.2)$$

where $c_{i(t)}$ is the activity of cortical cluster i at timestep t , w and p are the weight vectors and peripheral inputs respectively, l is the lateral connection field and c_{t-1} is the vector of cortical activity from the previous time step.

Full model of lidocaine block

To reproduce the effects of the block, we modelled the cortex with five units as above, representing five cortical digit-selective clusters akin to the baseline mapping of the fMRI. Based on the above finding, which requires full connectivity from each digit to each cluster, each cortical unit was connected to the five peripheral digit units. Cortical units were also connected laterally to all other units. We first demonstrate the cortical response without additional plasticity mechanisms to a set of modelled lidocaine digit responses. We then demonstrate additional mechanisms required to reproduce the empirical effect of the ‘block’ fMRI condition.

We utilise the TouchSim model as in the previous investigation. We modelled similar passive stimulation to the fMRI task for the baseline condition, generating simulated input to tapping the five fingertips. The anaesthetic block of the index digit, D2, was modelled by reducing the activity of the receptors in D2 under all digit tapping simulations to 20% of the baseline inputs (see Fig. 6.3A). This assumes some small residual activation, potentially from the movement of the digit or proprioceptive input. As before, peripheral input to each cortical cluster was pooled over each digit to produce five activity values for each digit during each stimulation.

First, we calculated the baseline model, which represents somatosensory cortex before any block manipulation. Based on the results of the toy model (see section 6.3.2, we note that lateral connectivity alone is not enough to produce the gradient of activation across clusters when there is no divergent connection. Each digit and the cortical cluster is therefore connected with a set of feedforward weight connections, which are stronger for a digit’s corresponding cluster and weaker for others (see Fig. 6.3B). Cortical units are connected by lateral weights, such that each cluster selectively either excites or inhibits itself and other clusters. The strength of the lateral connection is partly determined by the activation of each cluster when stimulated. The lateral part is dynamic and must be iterated over some timesteps before activity settles. Additional parameters are introduced from the toy model including the scalar offset for fitting the fMRI data, and a gain parameter which is manipulated to reproduce cortical changes. The complete model can be described as:

$$c_{i(t)} = \alpha_i \mathbf{w}_i^\top \mathbf{p} + b + \mathbf{l}_i^\top \mathbf{c}_{t-1} \quad (6.3)$$

where $c_{i(t)}$ is the activity of cortical cluster i at timestep t , α is a gain factor, w is a weight vector specifying each digit’s connection strength to the cortical cluster, p is the vector of peripheral inputs from each digit, b is a scalar offset. As before, l is the lateral connection field and c_{t-1}

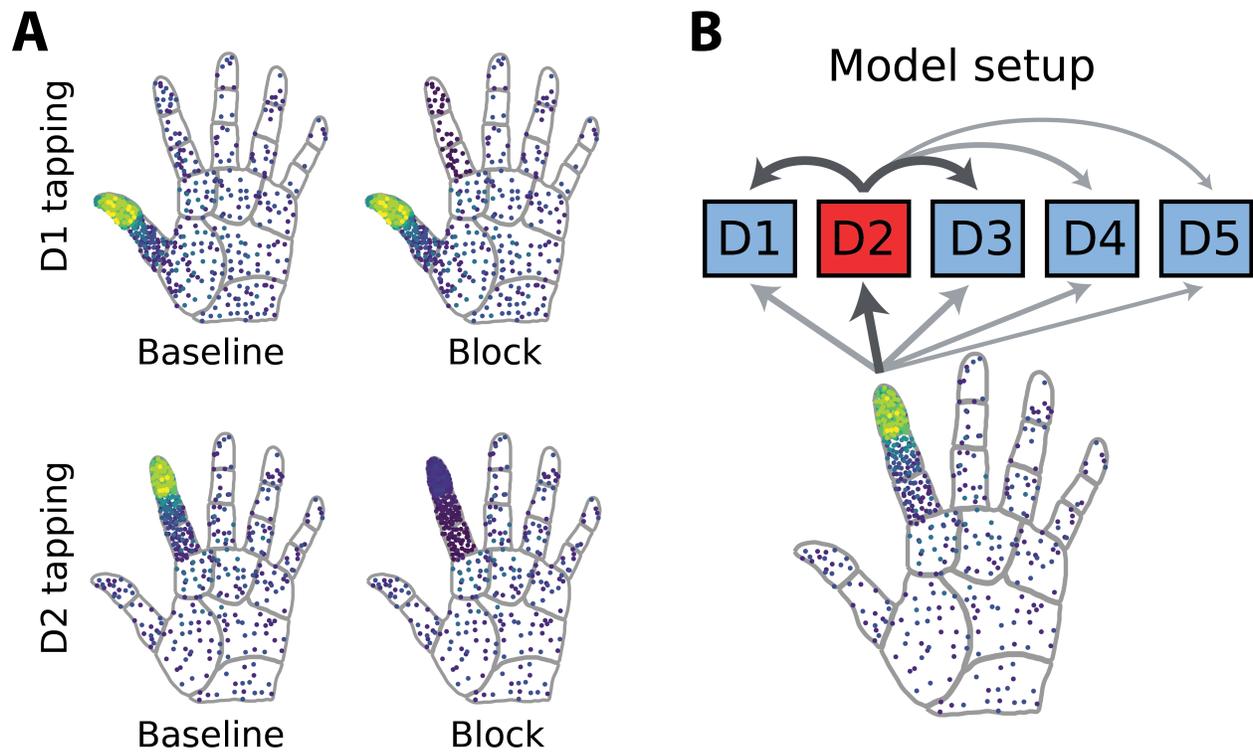


Figure 6.3: **A.** Typical responses to the model, with and without modelled lidocaine block for simulated tapping of D1 or D2 (blocked). Lighter coloured points indicate higher firing rates of afferents. The activity of the blocked digit is reduced under the tapping of both digits. **B.** Model setup. Each digit is connected to all clusters. Darker lines indicate stronger weighted connections occur between the matching cluster of a digit. Clusters are also connected laterally, with stronger connections to neighbouring digits than those further away.

is the vector of cortical activity from the previous time step. The lateral connections l_{ij} between clusters i and j depend on the distance between the cortical clusters alone. Therefore each cortical cluster has the same lateral connection pattern, but this is shifted spatially for each. The exact pattern of the lateral weights between cortical clusters depends on the model's baseline. Here we increased the baseline of the fMRI results fitted, using lateral connections where there is short range excitation and longer range inhibition for more distant clusters. The exact values for the pattern were manually calculated based on balancing inhibition and excitation to reproduce the fMRI activation seen across clusters. The choice of the rest-state will affect the precise pattern of lateral connections chosen; if the baseline is changed, the model may require a different pattern of connection. Here we use short-range excitation, however if the value of rest state changes, a different pattern may be required to reproduce the baseline results, for example, shorter-range inhibition. The full model setup with lateral connections is shown in Fig. 6.3B.

The feedforward weights w_i between digit activation and all clusters were fit using multiple regression. A global model regression intercept was fit to all weights of 0.7. The weight fitting was then calculated as an iterative process to consider the effects of the lateral connections, l . The first step of this process was to fit the weights without the lateral connectivity; here, we used the TouchSim generated responses before the digit block. Following this, the lateral contribution was added, and the error between this and the fMRI baseline data is calculated. Using the weights w_i calculated with multiple regression, equation 6.3 is then iterated until the activity settles in each cluster c , usually within 6 timesteps. All gain factors α were set to one for the baseline model. The fMRI values to be fitted were then adjusted based on the error, and this process was then repeated until the error was low. These weights and lateral connections provided the simulation result before lidocaine block.

6.2.3 Comparing before and after block: Representational similarity analysis

Representational similarity analysis (RSA, (Kriegeskorte et al., 2008)) enables comparison across different techniques (such as between fMRI and modelling) or experimental conditions. This is done by comparing variations in the evoked responses to different inputs. In this study, we used RSA to compare global responsivity in voxels/clusters across the S1 region during stimulation of each digit. This also enabled the comparison of changes across the map before and after lidocaine digit block. The full procedure is detailed in Fig. 6.4. First, the patterns of activity across all clusters (or imaged voxels) are measured in response to a certain input or stimuli. Here, the different inputs correspond to the stimulation of each digit (see Fig. 6.4, left). The evoked activity across all imaged voxels or model clusters to each of these stimuli is a vector of values, where each value corresponds to the activity in one cluster (see Fig. 6.4, middle). Using each vector, the distances between this and all other vectors are computed. Calculating all possible combination pairs, a symmetric dissimilarity matrix can be created (see Fig. 6.4, right). Pairs of stimuli that evoke more different patterns of activity across the response vectors will have a higher dissimilarity. For example, the distance (and resulting dissimilarity) between cortical responses to stimulation of the thumb and middle digit

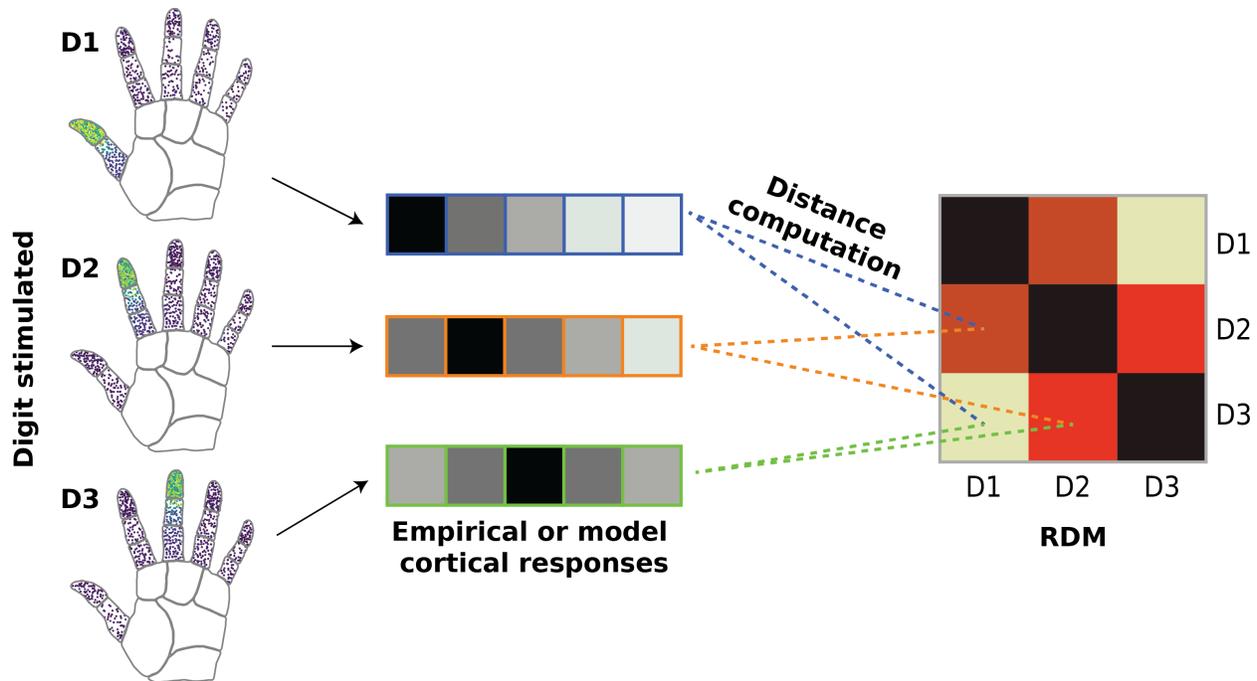


Figure 6.4: RSA procedure. Stimulation focused on the digit tips leads to different afferent responses across the hand. The cortical responses to these patterns can be thought of as vectors (for example, a vector of imaged fMRI voxels or modelled cortical cluster responses). These contain the responses of all voxels/clusters to a single stimulus, regardless of each voxel/cluster’s typical digit selectivity. Some of these voxels may have a larger response than others, indicated by a darker grey colour. Different patterns of response emerge based on the digit stimulated. Each response vector is then compared with every other using a chosen distance metric (see dashed lines). This results in the representational dissimilarity matrix (RDM). In this example, lighter colours indicate larger distances between the responses. The RDM is symmetric, with off-diagonals containing the dissimilarity value of cortical responses to pairs of different stimuli.

(D1 vs D3) is larger than that of stimulating neighbouring digits (compare the pattern of response across the top and bottom vectors in Fig. 6.4, versus top and middle). The resulting RDM provides a global representation of the digit responses across the whole hand map (considering all voxels or clusters), rather than typical univariate methods which focus only on maximally selective regions for each digit.

In the calculation of the RDM, different distance metrics can be used. For the fMRI data in this study, the cross-nobis (cross-validated Mahalanobis) distance matrix was selected (Wesselink et al., 2020). Distances were calculated between each digit pair, averaged over all imaging runs. Under this metric, larger values indicate greater distances (larger dissimilarities). These inter-digit dissimilarities can be further visualised using 2D multi-dimensional scaling (MDS). The dimensions of the dissimilarity matrix are reduced to the two which explain the largest between finger variance, enabling the relative distances to be visualised. For multiple conditions, such as visualising the difference between before and after lidocaine block, the 2D dissimilarities can be transformed using Procrustes alignment. This is a form of coordinate alignment where the data is rotated and mirrored (but not scaled) such that the two sets of points have the lowest alignment

error.

For the model, the RDM was calculated using the Euclidean distance metric. Classical metric MDS (non-ranked data) was then calculated from the pre-computed RDM using Python's Sklearn MDS algorithm (Pedregosa et al., 2011).

6.3 Results

6.3.1 Baseline fMRI results

Baseline somatosensory responses demonstrated that each corresponding cluster was selective of the corresponding digit (Fig. 6.5A). The activity of other digits followed a topographic pattern, with neighbours typically exhibiting stronger responses to stimulation than those digits further away. Representational similarity analysis for the baseline condition showed that the resulting representational dissimilarity matrix (RDM) demonstrated global structure across the hand region, reproducing the typical canonical pattern of somatotopy (see Fig. 6.5B). Dimensionality reduction of the RDM to two dimensions for visualisation purposes is shown in Fig. 6.5C.

6.3.2 Toy model results

We found that narrow digit connections to only the corresponding cluster led to disordered mapping results with no gradient of activation or a very weak one (Fig. 6.6A, top left and right, respectively). Conversely, the canonical representation emerged for both conditions with and without input digit co-activation with divergent connectivity to all clusters. However, having input co-activation enabled the representation of more distal digits, for example, stronger activation of the C5 cluster when D1 is stimulated.

These results were reflected in the RSA dissimilarity matrices, which showed the canonical pattern emerged strongly for both simulations where full digit cluster connectivity was used (Fig. 6.7A, bottom row). For single connectivity, the canonical representation did not emerge (no co-activation of other digits) or was only weak (with co-activation). MDS fittings show orderly representations matching these canonical representations, see Fig. 6.7B. For single connectivity with co-activation of digits, orderly representations were occasionally returned by the MDS fitting.

The same pattern is obtained with an input set where the stimulus radius was larger and indentation further, suggesting that even if the stimulation was stronger, the connections to all clusters are required to obtain the canonical representation. When the Gaussian was set to be wider, more distant digits on the hand are also connected to a cluster. However, this resulted in a larger activation of the immediately neighbouring clusters and those further out. This leads to a weaker gradient of activation, as responses across all clusters become more similar. Therefore there is less distinction between the activation of each cluster, particularly for the immediate neighbours.

We tested adding lateral connectivity to the fully distributed input patterns, which enabled a

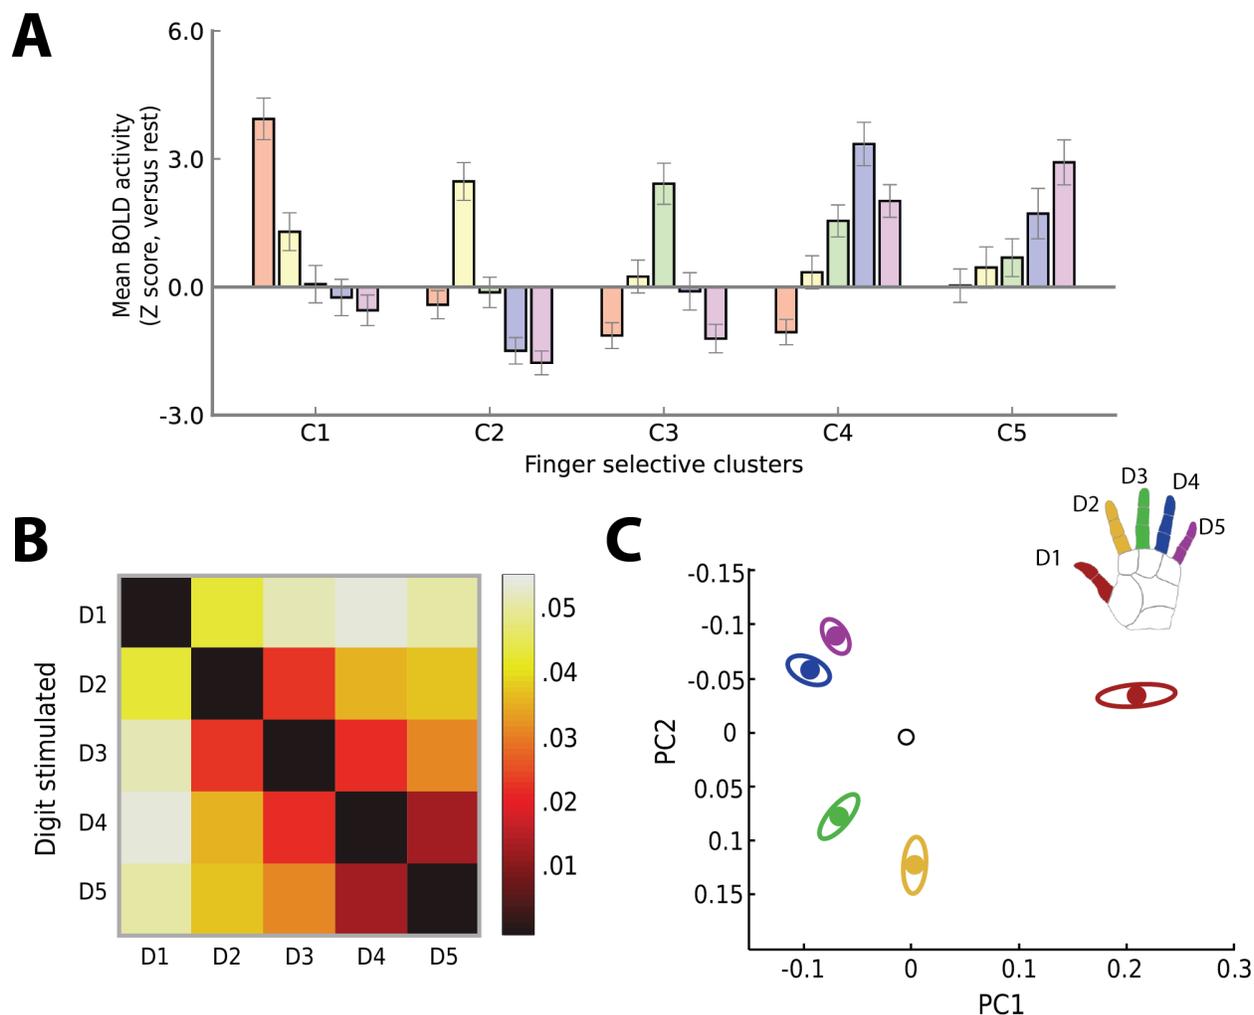


Figure 6.5: **A** Univariate fMRI results for group-level baseline activation during digit tapping stimulation versus rest. Each colour corresponds to responses during tapping of a single digit over each cluster C1-C5. Error bars are standard error of the mean. Responses in each cluster are greatest for the corresponding digit, for example C1 has greatest BOLD activity for tapping of D1. **B** Representational dissimilarity matrix during stimulation of each digit. Brighter colours denote larger distances (greater dissimilarity). Neighbouring digits are typically less dissimilar (darker colors) than those further away. Colourbar arbitrary units. Dissimilarity of responses typically increases with distance. This canonical shape has been reported in many previous studies (Ejaz et al., 2015; Mehring et al., 2019; Kieliba et al., 2021). **C** 2D MDS fitting of the RDM. Center circle denotes baseline mean and coloured circles the mean for each digit. Ellipses show between subject standard error. Distance between points denotes difference of representation across the somatosensory hand area. Points that are closer together have more similar representations, such as D4 and D5. Figures reproduced from Wesselink et al. (2020).

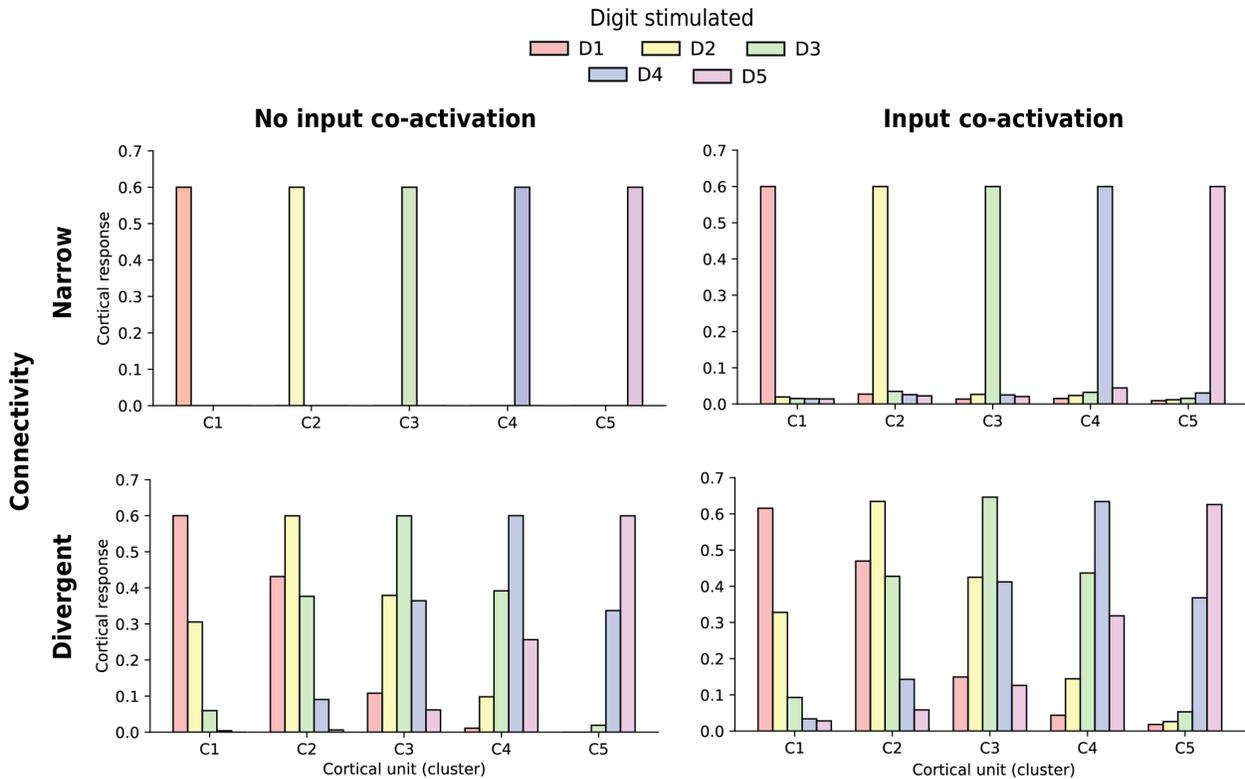


Figure 6.6: Responses to each cluster from the canonical model. Top row shows responses when digits are only connected to their respective clusters. Bottom row has divergent feedforward connectivity, where neighbour clusters receive some input from neighbouring digits. Left column indicates no input co-activation (afferent activation from neighbouring digits is removed), and right, including the afferent activation from neighbouring digits.

sharpening effect of the response, where responses within the dominant, corresponding cluster for that digit are reinforced, and responses in other digit clusters are reduced. We note that lateral connectivity alone is not enough to produce the gradient of activation across clusters when there is no divergent connection. Based on these results and past research, we therefore assume divergent connectivity is required and combine this with lateral connectivity, which is discussed further in the full model.

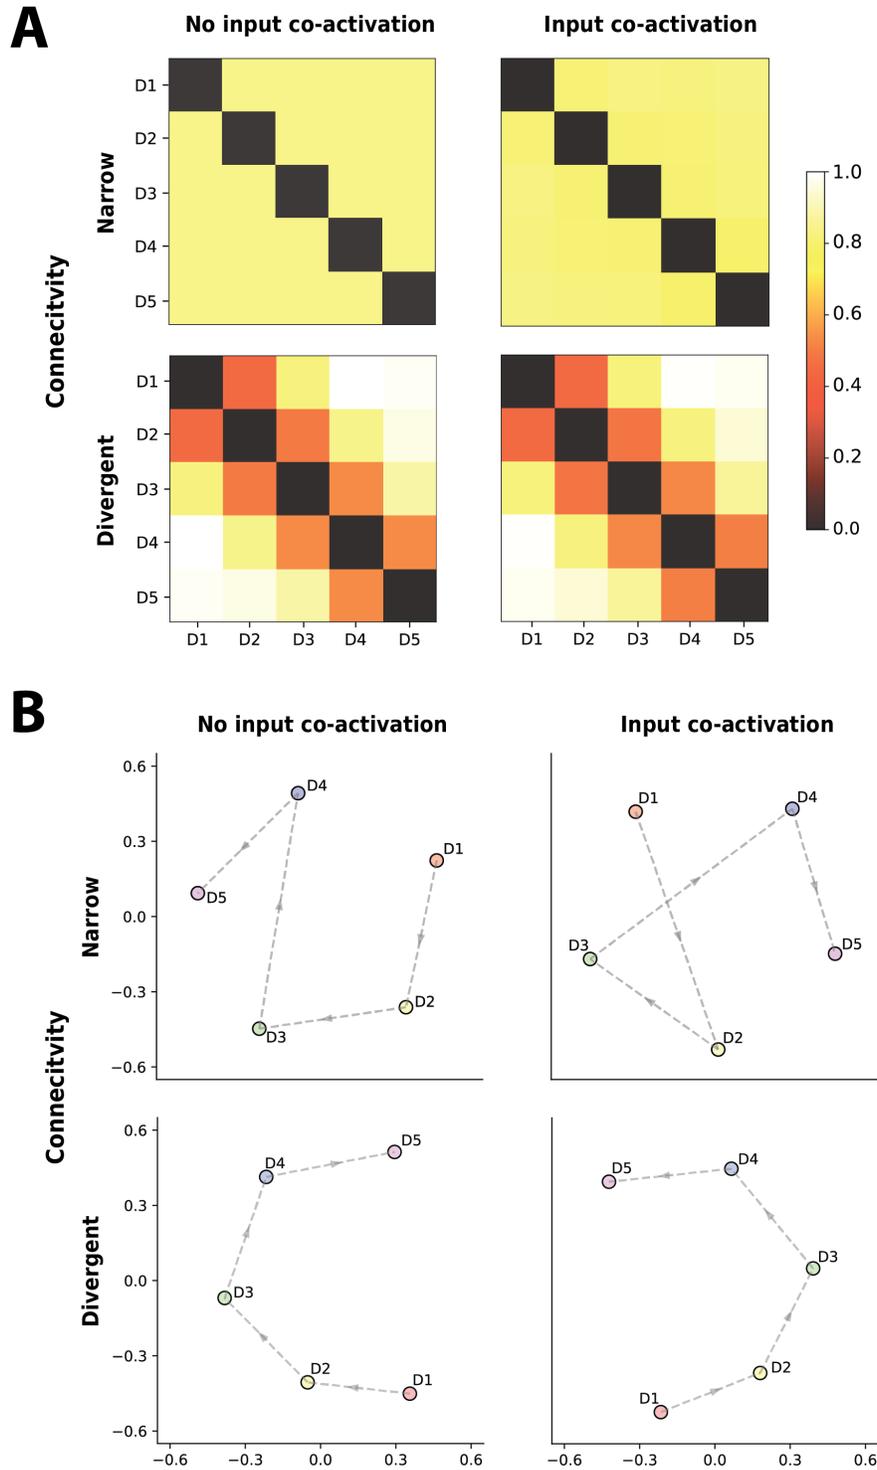


Figure 6.7: **A.** RSA dissimilarity matrices for each simulation considering different connectivity between the hand and cortical cluster (rows) and whether activation occurs across the hand (columns). Brighter colours indicate less similar responses. Only when connectivity is more divergent does the canonical hand pattern emerge. **B.** MDS reconstructions for dissimilarity matrices. Labels denote each digit, and directional lines demonstrate ordering if somatotopic structure is preserved. When lines cross or zigzag, this indicates disordered mapping, for example, in the top right, D3 and D4 representations are more similar to D1 than D2.

6.3.3 fMRI block results

Overall the results demonstrated that the block affected cortical responses to D2 stimulation (see Fig. 6.8A, yellow coloured bars across the clusters C1-C5). Post block, activation could not be elicited in the corresponding C2 cluster during stimulation of the blocked digit, D2, compared to rest ($\mu=-0.04$; $t_{14} = -0.10$, $p = .919$, Bayes' Factor=.266), and was significantly decreased from the baseline session ($t_{14}=-4.17$, $p = .001$), as assessed by t-tests corrected for multiple comparisons. Responses in other digit clusters under D2 stimulation were also reduced versus baseline ($F_{1,112}=9.90, p = .002$)³. Furthermore, changes were seen more globally across the hand map. Activity was not just reduced in the C2 cluster, but also attenuated in the other digit clusters versus baseline.

Global changes across the hand map were demonstrated by comparing inter-digit dissimilarity between conditions using RSA. At baseline, RSA demonstrates the typical canonical pattern, such that digit representations that are more similar to each other are positioned closer together. Post-block, the general shape of the canonical structure was preserved. Representations remained mostly stable between digits; however, all digits became more similar to one another (global scaling down of the pattern). The effect of the scaling is homogeneous between the digits, where all were scaled by a similar amount.

The MRS results demonstrated a decrease in the Glutamate/GABA ratio, potentially suggesting larger inhibitory GABA concentration across the region. However, due to the spatial imprecision of MRS, it is not possible to show fine-grain changes within the measured region. Therefore, it is not possible to distinguish whether inhibition is widespread or more localised, such as whether only certain digits are inhibited.

6.3.4 Full model: Somatosensory responses to lidocaine block

We first tested whether it is possible that the blocked peripheral input alone plus the feedforward connectivity and lateral dynamics were enough to reproduce the activity profiles of the fMRI without further changes; we call this the 'static model'. Here we use the previously calculated feedforward weights and the blocked input. RDMS were calculated, and the fit between the baseline and blocked inputs was compared using MDS and Procrustes alignment (see Fig. 6.9A). Given that all digits are connected to all clusters, it was possible this would lead to some consistent activation of D2 and the surrounding clusters based on the spread of activity across the hand and lateral connectivity. Therefore, this stimulation's input and resulting activation would not be zero for the other digits.

Given the reduction in D2 input, the static model showed reduced activation for the D2 cluster (see Fig. 6.9A, Procrustes alignment). However, unlike the empirical fMRI results, shrinking of activation was only seen for the D2 cluster and activation for other digit simulations remained

³Under the active condition, there was some activity in C2 versus rest, but significantly decreased versus baseline. Some residual C2 activation in the active condition is likely due to activation of proprioceptive pathways during movement or possible motor system efferent connectivity.

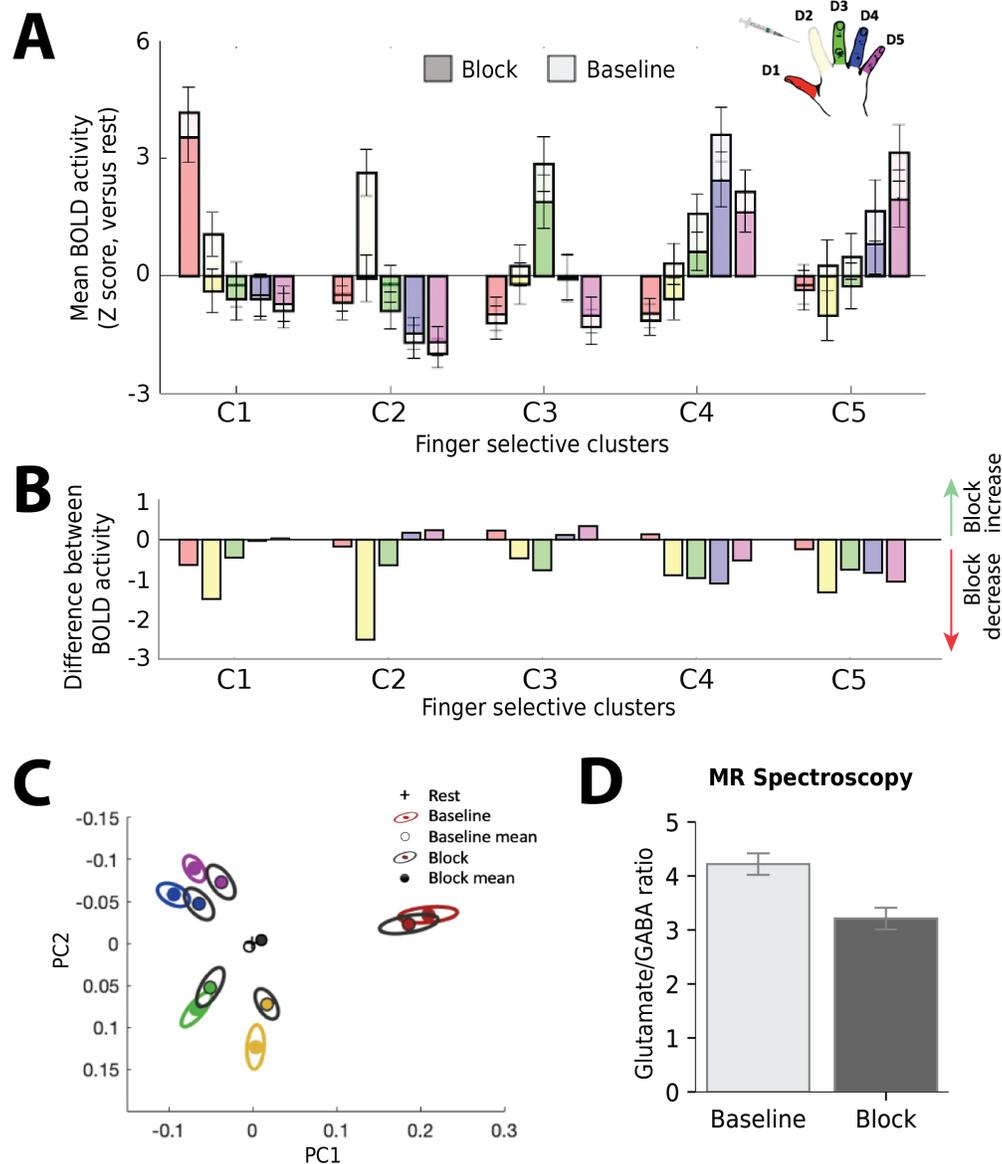


Figure 6.8: **A.** Univariate fMRI results for group-level baseline versus block digit activation. Shaded bars indicate blocked condition. Errors bars show the standard error of the mean. Change in colours denote the stimulation of each digit; for example, red shading denotes activation in the cluster C1-C5 for stimulation of the thumb. Overall, the general topographic gradient of digit activation is seen at baseline, with stimulated digits most strongly active in the respective clusters and neighbours showing more similar responses than those further away. **B.** Difference between BOLD activity for the baseline and block conditions. Bars below zero indicate *decrease* in activation for the block condition. Largest reductions in activity can be seen for the blocked digit (yellow bars). Furthermore, there is a global effect across all clusters. Under the stimulation of any digit, most clusters have a reduction of activation. Some small increases in activity are also present, typically for digits that are more distal to those of the corresponding cluster (see activation of D4 and D5 in cluster C2, for example). **C.** RSA results for baseline versus block condition activation. Coloured ellipses denote baseline, and black outlines denote block. There is a global reduction in the dissimilarity of all digits; digit representations become more similar. **D.** MRS results. Glut/GABA ratio decreases after lidocaine block across the scanned region. This suggests an increase in GABA inhibition for the whole area. All images reproduced from Wesselink et al. (2020).

stable (see square markers versus circle in Fig. 6.9A). The model suggests that the empirical results are unlikely to be due to changes in D2 activation alone. Therefore, other mechanisms are likely required to enable global shrinking of all digit activation at a similar level. We tested whether a central mechanism could mediate the peripheral loss and reproduce the global map-wide changes.

We implemented the central plasticity mechanisms by manipulating each cluster's gain parameter, α . We found that reproduction of the empirical data scaling required two changes. First, the effect could be replicated with a proportional reduction of digit activation by around 25% in all clusters, except C2. This global inhibition factor was fit via regression between responses to blocked inputs with baseline feedforward weights and the empirical block response. The regression slope was used as the global scaling factor, see Fig. 6.9B. The scaling was applied across all clusters except D2, which required boosting instead. As D2 activity was reduced in the static model beyond the empirical global scaling, we found that boosting of around 25% was required to produce a proportional scaling across digits. Small feedback effects through the lateral connections additionally help stabilise the representation of D2 across the clusters. Therefore the gain parameter for D2 was $\alpha_2 = 1.25$ and for all other digits $\alpha = 0.75$. The resulting homeostatic model demonstrated reduction across digits, and a global proportional scaling of the RDM, see Fig. 6.9C, left and right. Between the 'static' and 'homeostatic' RDM, both D2 dissimilarities and all others are reduced in the homeostatic case, but only D2 in the 'static' model.

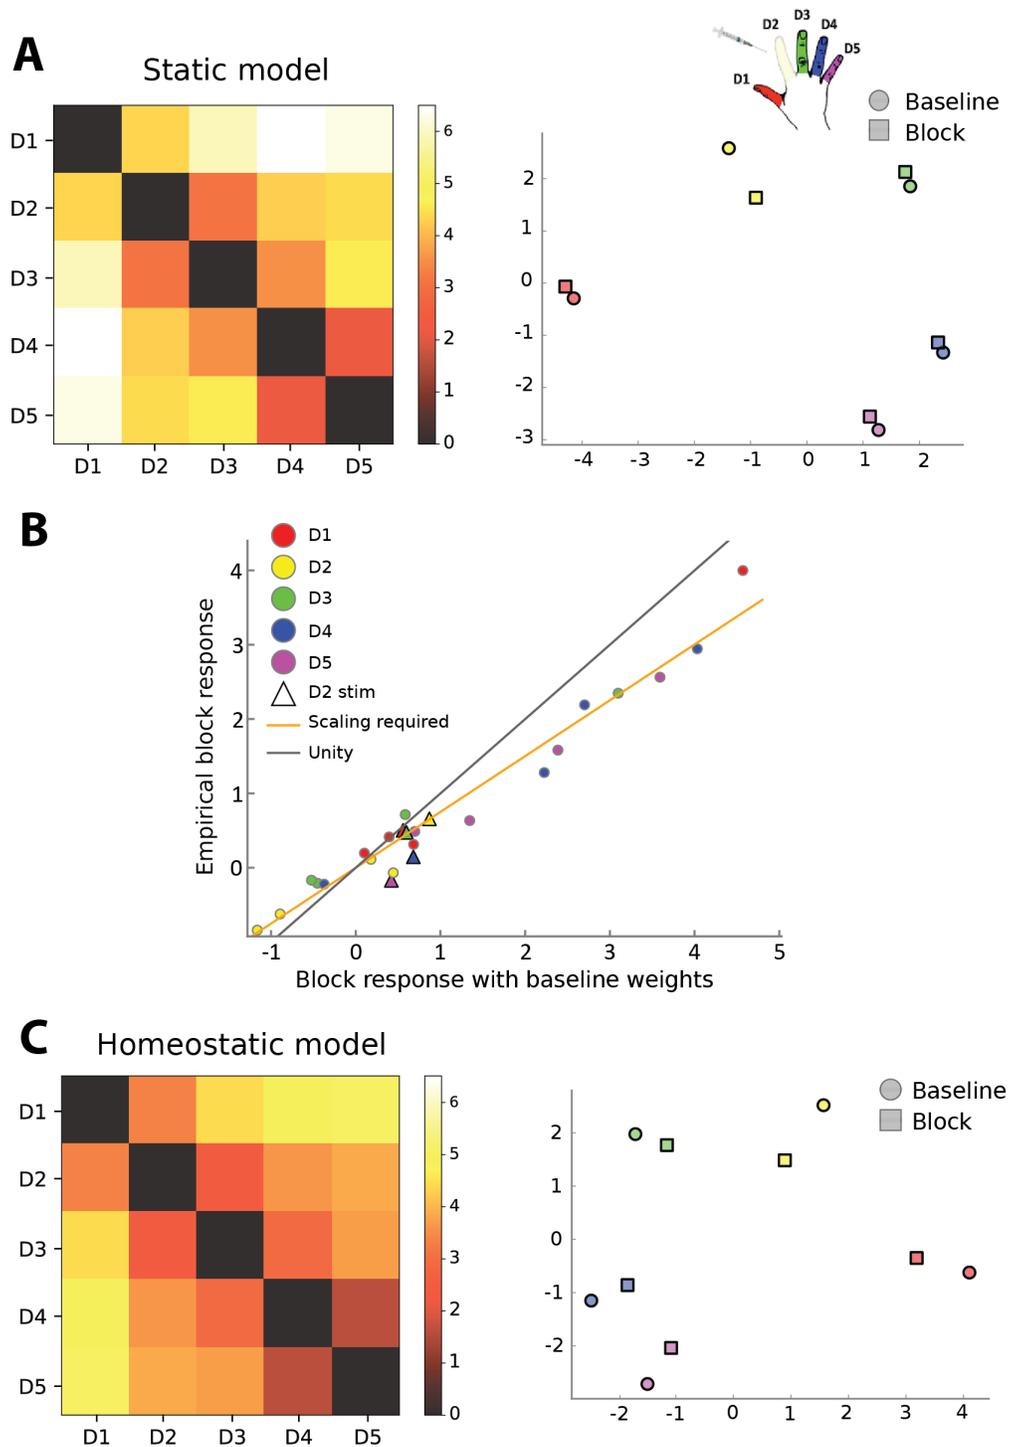


Figure 6.9: Model comparison for static model and homeostatic model. **A.** Dissimilarity matrix for the model with blocked inputs and no further changes (left), with corresponding Procrustes alignment between baseline inputs and static model (right). **B.** Comparison of predicted univariate activity with D2 block for baseline weight (static model) compared to empirical results from Fig. 6.8. **C.** The effect of digit block, as seen in the empirical results, can be obtained with a global decrease in gain, coupled with a change in D2 gain, here increase. Left, dissimilarity pattern, right, Procrustes alignment with baseline (non-block).

6.4 Discussion

In this chapter, we modelled somatosensory cortex responses to a short-term lidocaine block of the index digit, which had been empirically measured using fMRI in human participants (Wesselink et al., 2020). Pre- and post-fMRI scans, perceptual analyses demonstrated that tactile sensations from the index digit were abolished after nerve block, such that there was no longer accurate perception of stimuli. Analysis of S1 cortical responses demonstrated that post-block activity was reduced for the blocked digit and across the whole hand map. This extended across stimulation of other digits, even for those not anaesthetised. Furthermore, activation of the whole hand representation was reduced by a relatively constant amount. This was reflected in multivariate analyses, which demonstrated a global reduction of dissimilarity across the map, such that activation of digits became more similar regardless of digit stimulated. The topographic gradient representation of D2 did not collapse, and the representation was preserved across cortical clusters. This demonstrates that a missing input can cause global changes across the entire hand map, suggesting the effects of central plasticity mechanisms and more distributed representations of digits.

We aimed to understand the mechanisms required to reproduce the empirically measured cortical changes. We first examined the connectivity patterns required to reproduce the topographic gradient of activation seen across clusters and found that broad connectivity is required to obtain the canonical hand representation previously reported in empirical studies (Ejaz et al., 2015; Kieliba et al., 2021). We then reproduced the effect of digit block on tactile afferent responses in a realistic model of tactile activation under stimuli, TouchSim (Saal et al., 2017). Without additional plasticity mechanisms acting on the cortical representations, the global scaling of activation across the hand region could not be replicated under the model’s block reduction of the peripheral activation. Instead, we propose global mechanisms that could occur in reaction to unbalanced excitation and inhibition following a nerve block. We implemented this via a change in activation gain, which scales activation across the hand map. We found that the fMRI changes could be reproduced by a global inhibition across the whole hand map, with a selective increase in the activation of D2. These ideas were supported by magnetic resonance spectroscopy, which demonstrated increased GABA concentration across the hand map (suggesting an increase in inhibition). A change in gain could be considered a homeostatic-based mechanism acting on a shorter timescale to rebalance activity.

6.4.1 All-to-all map connectivity is required for activation of all clusters

When part of the hand is stimulated, receptor responses are not limited to the stimulation location, and instead, responses occur across the hand (Shao et al., 2020; Manfredi et al., 2012). This typically activates Pacinian Corpuscles and corresponding RAI afferents, which lay deep in the skin and have wide, not specific RFs (Johansson and Vallbo, 1983). With only localised connections between digits and cortical clusters, it is possible that activation occurring across the hand could elicit responses in all cortical digit clusters, therefore reproducing the topographic pattern of activation. However, as demonstrated in the initial model, this activation was not enough to produce the typical canonical

pattern of the RDM. Instead, more distributed connectivity was required for each cluster. Empirical studies suggest that connections into somatosensory cortex from the immediate upstream thalamus VPL region are divergent across the hand map (Qi et al., 2011). Furthermore, fMRI studies have recently demonstrated that the mapping may be more distributed in somatosensory cortex, with the representation of body parts found across the S1 map; for example, information about the hand is found in the typical representation of the foot (Muret et al., 2021).

Although the empirical data was not segmented into distinct subregions-Brodmann areas 3a, 3b, 1, analysis of probabilistic mapping demonstrated that most of the selected voxels fell into likely 3b or area 1. In area 1, representations of individual digits are less selective, with some multi-digit RFs. However, it is unlikely that the gradient across digit clusters is caused by potentially less individual selectivity in area 1, as previous studies where subregions were delineated show a similar pattern of activation in area 3b, with changes in activation across all clusters to stimulation of a single-digit (Martuzzi et al., 2014).

The addition of lateral connectivity enables a sharpening effect of the selectivity, meaning that corresponding cortical clusters to the stimulated digit are more active than surrounding clusters. Lateral connectivity in models is typically weaker than the feedforward afferent connections to cortex units; such that it is not enough alone to reproduce the activation gradient seen without more distributed connectivity.

6.4.2 Global and local changes are required to reproduce the effects of digit-block

Here we discuss possible causes of the global map changes, with activation attenuated across the hand region. We first cover potential reasons which could not explain the global changes in this case, then discuss possible causes emergent from the modelling.

Firstly, the fMRI effect in cluster C2 is unlikely due to residual peripheral activation from an unsuccessful nerve block, as there was no large residual activation of the digit. Furthermore, behavioural tests indicated that thresholds for detecting orientated gratings were at chance level (see Wesselink et al. (2020)). Other digits were perceptually unaffected by the block, suggesting that peripheral input was similar for both the block and baseline conditions, except for the blocked D2 digit.

Second, previous work has suggested that plasticity effects may be explained at short time scales by ‘unmasking’, where latent connections from other digits become uncovered by removing a dominant input. Further, the dominant input’s inhibition effects on other clusters are also removed, enabling them to increase their activation. The suggested unmasking effects demonstrated in previous work cannot explain the effects of the current study. With unmasking, residual input from other digits would be represented in the now-removed D2 region. However, C2 still had the greatest activation relative to other clusters when D2 was stimulated. It was not the case that voxels selective to the

lidocaine digit became selective to the neighbouring digits, as suggested in earlier works (Merzenich et al., 1984; Calford and Tweedale, 1991). Using a winner-takes-all mechanism would still lead to D2 being most strongly represented in the D2 cluster; other digits would not take-over this representation.

Modelling of the empirical results demonstrated that both global map-wide changes and local changes are required to reproduce the shrinking of the canonical RSA pattern seen under lidocaine block. We found that a global decrease, or inhibitory effect across the whole region, combined with an increase in the D2 cluster, was required to reproduce the effects. The global changes required for other digits were at a similar level, but lateral connectivity added distance-dependent interactions across the map. We suggest that one mechanism for these changes could be a homeostatic effect, aiming to regulate the network activity, given the drop-out of the D2 digit activity. Further investigation is required, but the model supports a non-specific homeostatic plasticity mechanism.

Homeostatic plasticity pushes the network activity towards a stable state, controlling any additional excitation or inhibition to maintain balance. Several mechanisms of homeostatic plasticity are possible, which all have some effect on a neuron's firing rate. Here we do not model or suggest a specific homeostatic mechanism but utilise a gain parameter to adjust and scale the overall activity of each cluster. This overall effect could be a completion of the topographic gradient seen as baseline level, with the excitation and inhibition balanced to maintain the typical pattern of activation across the clusters given D2 residual input. Although homeostasis is typically studied at longer timescales in relation to controlling Hebbian synaptic mechanisms (Turrigiano and Nelson, 2000), there is evidence this could occur at the rapid timescales in line with this study (Zenke et al., 2017; Castaldi et al., 2020).

We suggest that the result is unlikely a synaptic mechanism as this would involve mass rewiring of all connections, a process that seems unlikely given the short timescale the fMRI study was conducted. Reorganisation of large areas of the map (through re-sprouting of connections, for example) are typically not seen immediately and are instead studied weeks to years after amputation or deinnervation (Merzenich et al., 1984; Pons et al., 1985; Kambi et al., 2014; Florence et al., 1988). Although we do not suggest an exact homeostatic mechanism for this study, this may be some preparatory stage before long-term changes.

An increase in C2 activation via boosting the corresponding D2 digit activation, coupled with global inhibition, could stabilise the map by bringing responses of other digits within closer range of the response to the attenuated D2 inputs. The global inhibitory effect required has some support from the measured MR spectroscopy. This revealed a reduction in the Glutamate/GABA ratio, suggesting greater inhibition across the region. However, the lack of spatial precision for MRS means we cannot distinguish whether there was any increased activation (excitation) for the C2 cluster, as these effects are mixed with inhibition across the measured hand region.

6.4.3 Limitations and future work

S1 can be divided into several sub-regions, such as 3a, 3b, 1 and 2, which each have varying response properties. Most imaged voxels in this study could be probabilistically attributed to areas 3b and 1. However, other regions, such as 3a, are known to have interconnections with these regions, and as a result, some proprioceptive input may be present (see Delhaye et al. (2018) for a review). There could also be small induced finger movements from the stimulation in the passive condition that cannot be included in the peripheral model. We assumed some residual activation from D2 after the block to compensate for this. Therefore, the activation was reduced to only 20% of the initial level.

There is some evidence of plasticity at a sub-cortical level in the brainstem and the thalamus (Jones, 2003). It is possible that plastic changes could occur in these upstream targets, which are simply reflected in S1. Furthermore, our model is general, and it is possible that some effects occur at the lower levels, for example the cuneate nucleus or thalamus, and then propagate up to cortex. The model may therefore capture the effects at a conceptual level. We cannot disambiguate these possible changes in other upstream regions. However, most subcortical changes noted have been found on longer timescales, which may not be apparent at the short-timescales of this study. For example, widespread changes due to re-sprouting of connections in the subcortical regions were only seen years after amputation (Jain et al., 2000). However, other studies suggest plasticity mechanisms are cortical in origin. For example, after weeks of stimulating multiple digits synchronously, only cortical changes occurred, and no corresponding thalamic changes were found (Wang et al., 1995).

Lateral connectivity exists between representations of each digit in the hand map (Liao et al., 2013), however, the precise details of this have not been extensively characterised. For example, the types of populations mediating these and whether they are excitatory or inhibitory. Similarly, the strengths of these connections are relatively unknown, except that they are likely to fine-tune or boost the signals but not cause large changes in responsiveness over the afferent connectivity. In our model we assume shorter range excitatory connections and longer range inhibitory connections to other more distant clusters. We set the same lateral connectivity pattern for all clusters.

In our model, we do not consider potentially more widespread connections between other somatosensory regions, such as the neighbouring face region or upper limb, which borders the thumb and D5, respectively. Between the thumb and face, it is unlikely that there is a large contribution as these are separated by a cell-poor septa boundary, with relatively sparse connections (Fang et al., 2002). Changes across these have been noted, but at long term, scales utilising Hebbian plasticity mechanisms (Pons et al., 1991). Furthermore, some activation could have occurred on the palm region of the hand, an area we did not consider. In humans, recent 7T imaging found the centre of the palm to be located medial to D5 (Akselrod et al., 2021). Some activation is likely across the palm when simulating the digit responses in the peripheral model due to the wide-scale rippling effect. However, given the palm's position, activation in this region is more likely to affect C5 than the removed digit cluster. Furthermore, we demonstrated that residual activation across the hand

was not a driving force in the difference in representation post-block, and additional mechanisms were required.

In the model, we assume one cortical cluster for each digit; however, in somatosensory cortex, digit representations can be different sizes and vary in their distances with respect to each other; for example, the thumb is larger in its representation than other digits (Nelson et al., 1980; Martuzzi et al., 2014; Schweisfurth et al., 2018; Penfield and Boldrey, 1937). A more precise model may consider using multiple clusters with different total sizes for each digit and the extent of overlaps between representations. This may also lead to subtle changes to lateral connectivity is based on the distance between representations. Although a more complex model is unlikely to suggest different mechanisms (a global and local effect would still be required), this may allow better localisation of the extent of the C2 specific changes.

To conclude, our model required divergent afferent and lateral connectivity to reproduce the canonical hand map frequently observed in somatosensory mapping studies. Considering the digit block manipulation, we found that activation from other digits alone during D2 stimulation could not reproduce the pattern of responses in the map. Furthermore, in order to model the global attenuation of activation across the hand region, we found that an additional cortical mechanism was required. This could be due to the balancing of excitation and inhibition across the map to maintain the topographic gradients of digit representations.

Chapter 7

General discussion

7.1 Overview

In this thesis, we investigated the factors driving the organisation and representations in somatosensory maps. As discussed in chapter 2, representations are topographic such that neighbouring regions in somatosensory cortex are of adjacent body parts. Furthermore, certain body regions are allocated more resources than others, leading to larger representations than would be expected given their physical size. Of particular interest is the hand representation, which is magnified in cortex. What is driving the differences in body part representation? It has been suggested that the density of afferents could be one factor, as regions with a higher number of afferents tend to have greater magnification, such as the fingertips. However, regions with similar density and size can also vary in their cortical representations. Therefore, a role of body use has been implicated, such that regions that are more heavily involved in our tactile interactions should have greater representation, reflecting their importance. For example, those behaviourally more used tend to be the hands (for manipulating objects) and regions of the face (for communicating and eating behaviours). What is the contribution of these factors, and how does the setup of somatosensory system lead to the representations that emerge?

In chapter 3, we considered that the somatosensory system has limited resources available to process all incoming information from over 230,000 receptors on the body. We utilised ideas from the efficient coding framework, specifically decorrelation, to extract information from the inputs in terms of the amount of variance explained. We demonstrated that different ratios of receptor densities and activation, which vary across body regions, can lead to different optimal allocations. Crucially, we found that the allocation of resources to different regions depends on the bottleneck's width- or the amount of resources available in the system. This can lead to both over and under representation of different regions, depending on their densities and activation, which is a prominent feature of sensory cortical representation. Unlike visual sensing, the typical structure of tactile inputs for primates has not been well established. However, in the specialist sensing system of

the star-nosed mole, the statistics of tactile contact, receptors densities and cortical sizes have been extensively characterised. Therefore we tested the prediction of the efficient coding model in chapter 4 and predicted a close fit to cortical allocations when considering both the distribution of receptors and contact across the mole's nose. Comparing between models considering density or touch statistics only or both factors combined, we found models considering the tactile statistics best predicted empirical allocations. Adding density information marginally improved this prediction. From this, we suggest that the tactile system could be utilising decorrelation techniques to find optimal representations of inputs.

In chapter 5, we explored a further model utilising correlation between inputs to enable topographic representations of input surfaces using Hebbian based learning mechanisms- the Kohonen SOM. This was previously used to model the primary sensory cortex, focusing on plasticity. Whilst this model could accurately represent regions in the model map according to their probabilities of stimulation (the statistics of inputs), we found that there was little effect of different receptor densities in the map. Furthermore, empirical studies demonstrate that a principle of somatosensory mapping is the inverse magnification and receptive field rule, where over-magnified regions tend to have smaller corresponding RFs. We found that the map algorithm could not demonstrate this effect, and the receptive fields' size was largely dependent on the extent of the stimulation (stimuli size). Furthermore, unlike the abstract model of chapter 3, the exact cost function of the Kohonen SOM algorithm is unclear. Given that several empirically identified features of maps could not be explained, we suggest that this therefore may not be the same cost function as used by sensory maps in the brain.

Finally, we considered other map mechanisms that could be important in determining allocations, such as afferent to cortical connections and lateral connectivity patterns, which the previous model abstracted. In chapter 6, we demonstrated that certain conditions were required to reproduce the canonical hand arrangement, such as distributed connections between cortical clusters of each digit. Previous somatosensory models also focused on Hebbian based mechanisms which involve a rewiring of the connections. We instead considered the effects of a homeostatic type plasticity mechanism, which aims to re-balance excessive or over-dampened activation across cortical sensory regions. We demonstrated that this mechanism could reproduce empirical fMRI results of a short term plasticity manipulation, whereby the whole hand map activation was attenuated.

Overall, several considerations should be made for the representations of somatosensory maps. First, the statistics and structure of stimulation across the sheet, in part, determine the amount of information received, but this is ultimately limited by the spatial sampling across the input space- the density of receptors. Correlations in the inputs can be used to determine optimal representations; however, allocations are also influenced by the amount of resources available to a system. Finally, considerations should be made concerning the setup of the network structure from both upstream sensory regions and connectivity of representations within regions which suggest that inputs are representations in a somewhat more distributed way across the cortical maps. Finally,

the balance of network excitation and inhibition may also determine representations when inputs are perturbed, and this homeostatic type plasticity mechanism could be an early step in the more long-term reorganisation of maps.

In the next sections, we address wider implications of these ideas and applications to relevant domains. We explore possible influences on neuroscience on robotics domains, considerations for clinical applications and future technologies incorporating sensory processing, such as body augmentation, which have been popularised in recent years.

7.2 Considerations for future work

Natural tactile scenes

In vision and audition, natural images (Kersten, 1987) and sounds (Nelken et al., 1999; Moerel et al., 2012; Lewicki, 2002) have been used to investigate neural coding, as these stimuli have a typical structure that would be encountered by our sensory systems. In vision, this has been expanded to the use of videos with natural content (Berkes et al., 2009), enabling analysis of temporal aspects of sensing and ideas of predictive coding to be tested (Singer et al., 2018; Chalk et al., 2018; Salisbury and Palmer, 2016). Our work suggests that tactile statistics and the structure of inputs that impact the extent of resulting correlations would have an effect on resulting representations. Collected data of contact regions and forces could also be utilised with peripheral afferent models such as TouchSim (Saal et al., 2017), to create and understand afferent responses to natural stimuli, and, as in chapter 6 of this thesis, generate realistic inputs for somatosensory modelling. However, natural tactile scenes are more difficult to calculate due to the active nature of touch. Gathering and estimating the contacts of skin and typical forces or vibrations applied to the skin is therefore complex and involves significant engineering considerations. The skin is a 2D deformable sheet covering a 3D structure, with complex soft-tissue mechanics and vibratory wave propagation across the skin (Manfredi et al., 2012; Sripathi et al., 2006). Furthermore, natural tactile scenes encompass a large range of possible hand interactions with both active movements and passively applied stimulation. Which digits or regions of skin are contacted, material qualities of interacting objects, task demands and the goals of the user should be considered (Cini et al., 2019; Cutkosky, 1989).

One reason the structure of the tactile inputs has not been extensively addressed is due to the difficulty of accessing this information. Typical human hand grasps (Feix et al., 2009, 2016), and exploratory actions (Lederman and Klatzky, 1987) have been classified; however, few studies have extensively quantified these in a naturalistic environment (Gonzalez et al., 2014). Ways to accurately characterise tactile data is currently being approached using multiple methods. One is the use of tactile gloves and wearable sensors, which capture typical contact areas and forces during object interaction (Sundaram et al., 2019; Sagisaka et al., 2011; Büscher et al., 2015). Nevertheless, gloves are not often worn when manipulating objects in realistic environments, and as a result, this may lead to unnatural variations in hand use. An alternative approach that does not require the user to wear potentially bulky hardware is the use of motion sensing and camera recording.

Coupled with recent advances in computer vision technology, this can overcome previous problems such as digit occlusion (where digit position is out of view or covered by objects and other digits). Developments in algorithms have been able to create surface models of hand shape and form and recreate typical contacts across the hand in response to different object interactions (Grady et al., 2021; Brahmhatt et al., 2020; Cao et al., 2020). Nevertheless, many of these studies aim to solve the problem of accessing the data, and therefore there is still much work to be done to fully characterise natural hand statistics in humans.

The ability to quantify hand use in expert hand users may advance our understanding of the contributions of specific usage patterns and individual differences. For example, fMRI studies of musicians have found more greatly magnified representations of the digits in somatosensory cortex (Elbert et al., 1995). Long-term use of the hands in specialised patterns could be quantified, and the resulting maps compared to non-experts to further understand the impact of tactile input statistics.

7.3 Implications for robotics and sensing applications

Tactile sensing has many robotic applications, from improving manipulation and grasping in industrial work to developing responsive affective, social robots (Silvera-Tawil et al., 2015; Ege Cansev et al., 2021), and biomedical applications, such as enabling high accuracy tactile feedback in surgery (Tiwana et al., 2012). Most early robots have relied on visual sensors, lasers and motion encoders, to name a few (Li and Liu, 2019). However, direct touch-based sensing has been gaining interest through recent hardware and machine learning advances. High-resolution tactile sensing is more difficult, and most state-of-the-art sensors still utilise vision for this problem (Ward-Cherrier et al., 2018; Lambeta et al., 2020). The placement of sensors to allow high spatial acuity, for example, in robotic hands, is complex due to the size of sensors, power constraints and interference of other hardware for movement of the device (Dahiya et al., 2010; Dahiya, 2019). In the human hand, there are over 17,000 receptors (Johansson and Vallbo, 1979) that convey tactile information. However, as shown in chapter 3, how receptors (or in sensors in robotic devices) should be allocated to different input surfaces partly depends on how much those regions are contacted. This idea could be harnessed in robotic applications where the number of sensors is limited to determine optimal placements in order to maximise information.

7.4 Considerations for neuroscience: clinical applications and body augmentation

Clinical applications: prosthetics

Prosthetic devices can help restore sensorimotor function after limb loss. These can come in various forms, from non-functional cosmetic only limbs to robotic hands, with the ability to control individual digits and manipulate objects. It is widely accepted that successful interaction with the environment requires tactile feedback (Witney et al., 2004; Delhayet et al., 2016), and methods of

integrating this are being developed (Raspopovic et al., 2021; Bensmaia et al., 2020). Given the hardware restrictions and need to provide online real-time feedback, ideas from efficient coding theory may help to optimise information transfer. By considering how the prosthesis is used and the placement of its sensors, we can calculate which signals carry the most information.

Furthermore, evidence from plasticity studies has demonstrated that somatosensory cortex representations can be altered when tactile inputs are manipulated (Merzenich et al., 1984). However, more recent studies of amputees have demonstrated that some underlying representations may remain long after the loss of limb, which could be harnessed to reestablish tactile inputs of the prosthetic device (Makin and Bensmaia, 2017). In chapter 6 we investigated short-term plasticity and demonstrated that widespread changes after a short-term digit block could be modelled by considering more distributed feedforward connectivity and balancing of excitation and inhibition during input loss. Reorganisation effects could therefore be partly due to homeostatic changes (see Muret and Makin (2021) for further discussion on this). Although the scales of reorganisation may be more long-term for amputees and establishing prosthetic use, initial cortical changes may first be homeostasis-based and, therefore, contribute to later Hebbian synaptic reorganisation. Indeed, extensive works on plasticity mechanisms have demonstrated that homeostatic mechanisms may act to stabilise Hebbian changes and restore balance (Keck et al., 2017). Future models of somatosensory remapping under changing inputs should therefore consider these factors.

Neural resource allocation and capacity

The development of new consumer technologies such as wearables and virtual reality in gaming and telecommunications (Giri et al., 2021; Yang et al., 2021) has led to growing interest in the effects of sensory augmentation. For the tactile domain, this has included additional thumbs and digits (Kieliba et al., 2021; Baldi et al., 2020; Hussain et al., 2016; Shafti et al., 2021), to assistive devices such as exoskeletons (Dollar and Herr, 2008). Augmenting the body leads to changes in tactile inputs, particularly the statistics of our interactions. For example, the devices could directly stimulate parts of our skin or alter our usual patterns of hand use to accommodate them. How might cortical representations change to accommodate new, removed or manipulated inputs? If somatosensory maps have limited representational capacity, considerations should be made for how new inputs can be processed without negatively affecting current inputs or trading off some tactile processing abilities for improvements in other areas (for further discussion on this, see the recent review by Dominijanni et al. (2021)). Our work suggests that representations may change and be affected by the system's processing capacity. Developments of our findings may propose ways to harness existing representations or utilise the setup of the sensory system. For example, considering how mechanoreceptors are distributed in designing devices to enable higher resolution augmentation.

Bibliography

- Aguilar, J., Rivadulla, C., Soto, C., and Canedo, A. (2003). New corticocuneate cellular mechanisms underlying the modulation of cutaneous ascending transmission in anesthetized cats. *J. Neurophysiol.*, 89(6):3328–3339.
- Akselrod, M., Martuzzi, R., van der Zwaag, W., Blanke, O., and Serino, A. (2021). Relation between palm and finger cortical representations in primary somatosensory cortex: A 7T fMRI study. *Hum. Brain Mapp.*, 42(7):2262–2277.
- Allard, T., Clark, S. A., Jenkins, W. M., and Merzenich, M. M. (1991). Reorganization of somatosensory area 3b representations in adult owl monkeys after digital syndactyly. *J. Neurophysiol.*, 66(3):1048–1058.
- Amari, S.-I. (1980). Topographic organization of nerve fields. *Bull. Math. Biol.*, 42(3):339–364.
- Ames, III, A. (2000). CNS energy metabolism as related to function. *Brain Res. Rev.*, 34:42–68.
- Anderson, S. J., Mullen, K. T., and Hess, R. F. (1991). Human peripheral spatial resolution for achromatic and chromatic stimuli: limits imposed by optical and retinal factors. *J. Physiol.*, 442:47–64.
- Arcaro, M. J., Schade, P. F., and Livingstone, M. S. (2019). Body map proto-organization in newborn macaques. *Proc. Natl. Acad. Sci. U.S.A.*, 116(49):24861–24871.
- Armentrout, S. L., Reggia, J. A., and Weinrich, M. (1994). A neural model of cortical map reorganization following a focal lesion. *Artif. Intell. Med.*, 6(5):383–400.
- Asanuma, H., Larsen, K., and Yumiya, H. (1980). Peripheral input pathways to the monkey motor cortex. *Exp. Brain Res.*, 38(3):349–355.
- Ashaber, M., Pálfi, E., Friedman, R. M., Palmer, C., Jáklí, B., Chen, L. M., Kántor, O., Roe, A. W., and Négyessy, L. (2014). Connectivity of somatosensory cortical area 1 forms an anatomical substrate for the emergence of multifinger receptive fields and complex feature selectivity in the squirrel monkey (*saimiri sciureus*). *J. Comp. Neurol.*, 522(8):1769–1785.
- Atick, J. J. (1992). Could information theory provide an ecological theory of sensory processing? *Network: Computation in neural systems*, 3(2):213–251.

- Atick, J. J. and Redlich, A. N. (1990). Towards a theory of early visual processing. *Neural Comput.*, 2(3):308–320.
- Atick, J. J. and Redlich, A. N. (1992). What does the retina know about natural scenes? *Neural Comput.*, 4(2):196–210.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychol. Rev.*, 61(3):183–193.
- Ayers, H. (1884). *On the structure and development of the nasal rays in Condylura cristata*, volume 4-356-360. Biologisches Centralblatt.
- Azevedo, F. A. C., Carvalho, L. R. B., Grinberg, L. T., Farfel, J. M., Ferretti, R. E. L., Leite, R. E. P., Jacob Filho, W., Lent, R., and Herculano-Houzel, S. (2009). Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *J. Comp. Neurol.*, 513(5):532–541.
- Baldi, T. L., D’Aurizio, N., Gaudeni, C., Gurgone, S., Borzelli, D., D’Avella, A., and Prattichizzo, D. (2020). Exploiting implicit kinematic kernel for controlling a wearable robotic extra-finger. *arXiv preprint arXiv:2012.03600*.
- Barlow, H. (1980). The ferrier lecture, 1980: Critical limiting factors in the design of the eye and visual cortex. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 212(1186):1–34.
- Barlow, H. (2001). Redundancy reduction revisited. *Network: Computation in Neural Systems*, 12(3):241–253.
- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1:217–234.
- Bednar, J. A. and Miikkulainen, R. (2006). Joint maps for orientation, eye, and direction preference in a self-organizing model of V1. *Neurocomputing*, 69(10-12):1272–1276.
- Bednar, J. A. and Wilson, S. P. (2016). Cortical maps. *Neuroscientist*, 22(6):604–617.
- Bell, A. J. and Sejnowski, T. J. (1997). Edges are the independent components of natural scenes. In *Advances in Neural Information Processing Systems 9*, pages 831–837.
- Bensmaia, S. J., Tyler, D. J., and Micera, S. (2020). Restoration of sensory information via bionic hands. *Nat Biomed Eng*, pages 1–13.
- Berkes, P., Turner, R. E., and Sahani, M. (2009). A structured model of video reproduces primary visual cortical organisation. *PLoS Comput. Biol.*, 5(9):e1000495.
- Besle, J., Sánchez-Panchuelo, R.-M., Bowtell, R., Francis, S., and Schluppeck, D. (2014). Event-related fMRI at 7T reveals overlapping cortical representations for adjacent fingertips in S1 of individual subjects. *Hum. Brain Mapp.*, 35(5):2027–2043.

- Bhand, M., Mudur, R., Suresh, B., Saxe, A., and Ng, A. Y. (2011). Unsupervised learning models of primary cortical receptive fields and receptive field plasticity. In Shawe-Taylor, J., Zemel, R. S., Bartlett, P. L., Pereira, F., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems 24*, pages 1971–1979. Curran Associates, Inc.
- Bialek, W., Nemenman, I., and Tishby, N. (2001). Predictability, complexity, and learning. *Neural Comput.*, 13(11):2409–2463.
- Blankenburg, F., Ruben, J., Meyer, R., Schwiemann, J., and Villringer, A. (2003). Evidence for a rostral-to-caudal somatotopic organization in human primary somatosensory cortex with mirror-reversal in areas 3b and 1. *Cereb. Cortex*, 13(9):987–993.
- Boillat, Y., Bazin, P.-L., and van der Zwaag, W. (2020). Whole-body somatotopic maps in the cerebellum revealed with 7T fMRI. *Neuroimage*, 211:116624.
- Bouton, C. and Pagès, G. (1994). Convergence in distribution of the One-Dimensional kohonen algorithms when the stimuli are not uniform. *Adv. Appl. Probab.*, 26(1):80–103.
- Brahmbhatt, S., Ham, C., Kemp, C. C., and Hays, J. (2019). ContactDB: Analyzing and predicting grasp contact via thermal imaging. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Brahmbhatt, S., Tang, C., Twigg, C. D., Kemp, C. C., and Hays, J. (2020). ContactPose: A dataset of grasps with object contact and hand pose. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16.*, pages 361–378. Springer International Publishing.
- Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Barth.
- Brunel, N. and Nadal, J. P. (1998). Mutual information, fisher information, and population coding. *Neural Comput.*, 10(7):1731–1757.
- Bullock, I. M., Ma, R. R., and Dollar, A. M. (2013). A hand-centric classification of human and robot dexterous manipulation. *IEEE Trans. Haptics*, 6(2):129–144.
- Büscher, G. H., Kõiva, R., Schürmann, C., Haschke, R., and Ritter, H. J. (2015). Flexible and stretchable fabric-based tactile sensor. *Rob. Auton. Syst.*, 63:244–252.
- Calford, M. B. (2002). Dynamic representational plasticity in sensory cortex. *Neuroscience*, 111(4):709–738.
- Calford, M. B. and Tweedale, R. (1991). Immediate expansion of receptive fields of neurons in area 3b of macaque monkeys after digit denervation. *Somatosens. Mot. Res.*, 8(3):249–260.

- Camarillo, L., Luna, R., Nácher, V., and Romo, R. (2012). Coding perceptual discrimination in the somatosensory thalamus. *Proc. Natl. Acad. Sci. U.S.A.*, 109(51):21093–21098.
- Cao, Z., Radosavovic, I., Kanazawa, A., and Malik, J. (2020). Reconstructing Hand-Object interactions in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12417–12426.
- Carlson, M. (1981). Characteristics of sensory deficits following lesions of brodmann’s areas 1 and 2 in the postcentral gyrus of macaca mulatta. *Brain Res.*, 204(2):424–430.
- Castaldi, E., Lunghi, C., and Morrone, M. C. (2020). Neuroplasticity in adult human visual cortex. *Neurosci. Biobehav. Rev.*, 112:542–552.
- Catani, M. (2017). A little man of some importance. *Brain*, 140(11):3055–3061.
- Catania, K. C. (1995a). Magnified cortex in star-nosed moles. *Nature*, 375(6531):453–454.
- Catania, K. C. (1995b). Structure and innervation of the sensory organs on the snout of the star-nosed mole. *J. Comp. Neurol.*, 351(4):536–548.
- Catania, K. C. (1996). Ultrastructure of the eimer’s organ of the star-nosed mole. *J. Comp. Neurol.*, 365(3):343–354.
- Catania, K. C. (2020). All in the family - touch versus olfaction in moles. *Anat. Rec.*, 303(1):65–76.
- Catania, K. C. and Kaas, J. H. (1995). Organization of the somatosensory cortex of the star-nosed mole. *J. Comp. Neurol.*, 351(4):549–567.
- Catania, K. C. and Kaas, J. H. (1996). The unusual nose and brain of the Star-Nosed mole. *Bioscience*, 46(8):578–586.
- Catania, K. C. and Kaas, J. H. (1997). Somatosensory fovea in the star-nosed mole: behavioral use of the star in relation to innervation patterns and cortical representation. *J. Comp. Neurol.*, 387(2):215–233.
- Catania, K. C., Leitch, D. B., and Gauthier, D. (2011). A star in the brainstem reveals the first step of cortical magnification. *PLoS One*, 6(7):e22406.
- Catania, K. C., Northcutt, R. G., Kaas, J. H., and Beck, P. D. (1993). Nose stars and brain stripes. *Nature*, 364(6437):493.
- Catania, K. C. and Remple, F. E. (2004). Tactile foveation in the star-nosed mole. *Brain Behav. Evol.*, 63(1):1–12.
- Chalk, M., Marre, O., and Tkačik, G. (2018). Toward a unified theory of efficient, predictive, and sparse coding. *Proc. Natl. Acad. Sci. U.S.A.*, 115(1):186–191.

- Chaplin, T. A., Yu, H. H., and Rosa, M. G. P. P. (2013). Representation of the visual field in the primary visual area of the marmoset monkey: Magnification factors, point-image size, and proportionality to retinal ganglion cell density. *J. Comp. Neurol.*, 521(5):1001–1019.
- Christie, J. M. and Westbrook, G. L. (2006). Lateral excitation within the olfactory bulb. *J. Neurosci.*, 26(8):2269–2277.
- Churchill, J. D., Muja, N., Myers, W. A., Besheer, J., and Garraghty, P. E. (1998). Somatotopic consolidation: a third phase of reorganization after peripheral nerve injury in adult squirrel monkeys. *Exp. Brain Res.*, 118(2):189–196.
- Cini, F., Ortenzi, V., Corke, P., and Controzzi, M. (2019). On the choice of grasp type and location when handing over an object. *Sci Robot*, 4(27).
- Clarke, D. D. and Sokoloff, L. (1994). *Circulation and energy metabolism of the brain. In: Basic Neurochemistry.*
- Collins, C. E., Turner, E. C., Sawyer, E. K., Reed, J. L., Young, N. A., Flaherty, D. K., and Kaas, J. H. (2016). Cortical cell and neuron density estimates in one chimpanzee hemisphere. *Proc. Natl. Acad. Sci. U. S. A.*, 113(3):740–745.
- Connolly, M. and Van Essen, D. (1984). The representation of the visual field in parvicellular and magnocellular layers of the lateral geniculate nucleus in the macaque monkey. *J. Comp. Neurol.*, 226(4):544–564.
- Coq, J. O. and Xerri, C. (1998). Environmental enrichment alters organizational features of the forepaw representation in the primary somatosensory cortex of adult rats. *Exp. Brain Res.*, 121(2):191–204.
- Corniani, G. and Saal, H. P. (2020). Tactile innervation densities across the whole body. *J. Neurophysiol.*, 124(4):1229–1240.
- Cottrell, M., Olteanu, M., Rossi, F., and Villa-Vialaneix, N. (2018). Self-Organizing maps, theory and applications. *Revista de Investigacion Operacional*, 39(1):1–22.
- Cowey, A. and Rolls, E. T. (1974). Human cortical magnification factor and its relation to visual acuity. *Exp. Brain Res.*, 21(5):447–454.
- Cunningham, D. A., Machado, A., Yue, G. H., Carey, J. R., and Plow, E. B. (2013). Functional somatotopy revealed across multiple cortical regions using a model of complex motor task. *Brain Res.*, 1531:25–36.
- Curcio, C. A. and Allen, K. A. (1990). Topography of ganglion cells in human retina. *J. Comp. Neurol.*, 300(1):5–25.

- Curcio, C. A., Sloan, K. R., Kalina, R. E., and Hendrickson, A. E. (1990). Human photoreceptor topography. *Journal of comparative neurology*, 292(4):497–523.
- Cutkosky and Wright, P. (1986). Modeling manufacturing grips and correlations with the design of robotic hands. *IEEE Int. Conf. Robot. Autom.*, 3:1533–1539.
- Cutkosky, M. R. (1989). On grasp choice, grasp models, and the design of hands for manufacturing tasks. *IEEE Trans. Rob. Autom.*, 5(3):269–279.
- Dahiya, R. (2019). E-Skin: From humanoids to humans [point of view]. *Proc. IEEE*, 107(2):247–252.
- Dahiya, R. S., Metta, G., Valle, M., and Sandini, G. (2010). Tactile Sensing—From humans to humanoids. *IEEE Trans. Rob.*, 26(1):1–20.
- Dall’Orso, S., Steinweg, J., Allievi, A. G., Edwards, A. D., Burdet, E., and Arichi, T. (2018). Somatotopic mapping of the developing sensorimotor cortex in the preterm human brain. *Cereb. Cortex*, 28(7):2507–2515.
- Darian-Smith, C., Darian-Smith, I., Burman, K., and Ratcliffe, N. (1993). Ipsilateral cortical projections to areas 3a, 3b, and 4 in the macaque monkey. *J. Comp. Neurol.*, 335(2):200–213.
- Delhaye, B. P., Long, K. H., and Bensmaia, S. J. (2018). Neural basis of touch and proprioception in primate cortex. *Compr. Physiol.*, 8:1575–1602.
- Delhaye, B. P., Saal, H. P., and Bensmaia, S. J. (2016). Key considerations in designing a somatosensory neuroprosthesis. *J. Physiol. Paris*, 110(4):402–408.
- Dempsey-Jones, H., Harrar, V., Oliver, J., Johansen-Berg, H., Spence, C., and Makin, T. R. (2016). Transfer of tactile perceptual learning to untrained neighboring fingers reflects natural use relationships. *J. Neurophysiol.*, 115(3):1088–1097.
- Deuchert, M., Ruben, J., Schwiemann, J., Meyer, R., Thees, S., Krause, T., Blankenburg, F., Villringer, K., Kurth, R., Curio, G., and Villringer, A. (2002). Event-related fMRI of the somatosensory system using electrical finger stimulation. *Neuroreport*, 13(3):365–369.
- Doi, E., Gauthier, J. L., Field, G. D., Shlens, J., Sher, A., Greschner, M., Machado, T. A., Jepson, L. H., Mathieson, K., Gunning, D. E., Litke, A. M., Paninski, L., Chichilnisky, E. J., and Simoncelli, E. P. (2012). Efficient coding of spatial information in the primate retina. *J. Neurosci.*, 32(46):16256–16264.
- Doi, E. and Lewicki, M. S. (2014). A simple model of optimal population coding for sensory systems. *PLoS Comput. Biol.*, 10(8):e1003761.
- Dollar, A. M. and Herr, H. (2008). Lower extremity exoskeletons and active orthoses: Challenges and State-of-the-Art. *IEEE Trans. Rob.*, 24(1):144–158.

- Dolman, C. L., McCormick, A. Q., and Drance, S. M. (1980). Aging of the optic nerve. *Arch. Ophthalmol.*, 98(11):2053–2058.
- Dominijanni, G., Shokur, S., Salvietti, G., Buehler, S., Palmerini, E., Rossi, S., De Vignemont, F., d’Avella, A., Makin, T. R., Prattichizzo, D., and Micera, S. (2021). The neural resource allocation problem when enhancing human bodies with extra robotic limbs. *Nature Machine Intelligence*, 3(10):850–860.
- Duncan, R. O. and Boynton, G. M. (2003). Cortical magnification within human primary visual cortex correlates with acuity thresholds. *Neuron*, 38(4):659–671.
- Duncan, R. O. and Boynton, G. M. (2007). Tactile hyperacuity thresholds correlate with finger maps in primary somatosensory cortex (s1). *Cereb. Cortex*, 17(12):2878–2891.
- Dutta, A., Kambi, N., Raghunathan, P., Khushu, S., and Jain, N. (2014). Large-scale reorganization of the somatosensory cortex of adult macaque monkeys revealed by fMRI. *Brain Struct. Funct.*, 219(4):1305–1320.
- Edelman, G. M. (1981). Group selection as the basis for higher brain function. *The organization of the cerebral cortex*, pages 535–563.
- Edmondson, L. R., Jimenez Rodriguez, A., and Saal, H. P. (2019). Nonlinear scaling of resource allocation in sensory bottlenecks. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alche Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32*, pages 7543–7552. Curran Associates, Inc.
- Ege Cansev, M., Nordheimer, D., Andrea Kirchner, E., and Beckerle, P. (2021). Feel-Good requirements: Neurophysiological and psychological design criteria of affective touch for (assistive) robots. *Front. Neurorobot.*, 15:661207.
- Ejaz, N., Hamada, M., and Diedrichsen, J. O. r. (2015). Hand use predicts the structure of representations in sensorimotor cortex. *Nat. Neurosci.*, 103(June):1–10.
- Elbert, T., Pantev, C., Wienbruch, C., Rockstroh, B., and Taub, E. (1995). Increased cortical representation of the fingers of the left hand in string players. *Science*, 270(5234):305–307.
- Faggin, B. M., Nguyen, K. T., and Nicoletti, M. A. (1997). Immediate and simultaneous sensory reorganization at cortical and subcortical levels of the somatosensory system. *Proc. Natl. Acad. Sci. U.S.A.*, 94(17):9428–9433.
- Fang, P.-C., Jain, N., and Kaas, J. H. (2002). Few intrinsic connections cross the hand-face border of area 3b of new world monkeys. *J. Comp. Neurol.*, 454(3):310–319.
- Fang, P.-C., Stepniewska, I., and Kaas, J. H. (2006). The thalamic connections of motor, premotor, and prefrontal areas of cortex in a prosimian primate (otolemur garnetti). *Neuroscience*, 143(4):987–1020.

- Feix, T., Pawlik, R., Schmiedmayer, H.-B., Romero, J., and Kragic, D. (2009). A comprehensive grasp taxonomy. *Robotics, Science and Systems Conference: Workshop on Understanding the Human Hand for Advancing Robotic Manipulation*, pages 2–3.
- Feix, T., Romero, J., Schmiedmayer, H. B., Dollar, A. M., and Kragic, D. (2016). The GRASP taxonomy of human grasp types. *IEEE Transactions on Human-Machine Systems*, 46(1):66–77.
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am.*, 4(12):2379.
- Finlay, B. L., Franco, E. C. S., Yamada, E. S., Crowley, J. C., Parsons, M., Muniz, J. A. P. C., and Silveira, L. C. L. (2008). Number and topography of cones, rods and optic nerve axons in new and old world primates. *Vis. Neurosci.*, 25(3):289–299.
- Flanagan, J. G. (2006). Neural map specification by gradients. *Curr. Opin. Neurobiol.*, 16(1):59–66.
- Florence, S. L., Jain, N., and Kaas, J. H. (1997). Plasticity of somatosensory cortex in primates. *Semin. Neurosci.*, 9(1-2):3–12.
- Florence, S. L., Taub, H. B., and Kaas, J. H. (1998). Large-scale sprouting of cortical connections after peripheral injury in adult macaque monkeys. *Science*, 282(5391):1117–1121.
- Florence, S. L., Wall, J. T., and Kaas, J. H. (1988). The somatotopic pattern of afferent projections from the digits to the spinal cord and cuneate nucleus in macaque monkeys. *Brain Res.*, 452(1-2):388–392.
- Fox, K. and Stryker, M. (2017). Integrating hebbian and homeostatic plasticity: introduction. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, 372(1715).
- Fragaszy, D. M. and Crast, J. (2016). Functions of the hand in primates. In Kivell, T. L., Lemelin, P., Richmond, B. G., and Schmitt, D., editors, *The Evolution of the Primate Hand: Anatomical, Developmental, Functional, and Paleontological Evidence*, pages 313–344. Springer New York, New York, NY.
- Francis, S. T., Kelly, E. F., Bowtell, R., Dunseath, W. J. R., Folger, S. E., and McGlone, F. (2000). fMRI of the responses to vibratory stimulation of digit tips. *Neuroimage*, 11(3):188–202.
- Friedman, D. P. and Jones, E. G. (1981). Thalamic input to areas 3a and 2 in monkeys. *J. Neurophysiol.*, 45(1):59–85.
- Friedman, R. M., Chen, L. M., and Roe, A. W. (2008). Responses of areas 3b and 1 in anesthetized squirrel monkeys to single- and dual-site stimulation of the digits. *J. Neurophysiol.*, 100(6):3185–3196.
- Ganguli, D. and Simoncelli, E. P. (2010). Implicit encoding of prior probabilities in optimal neural populations. *Adv. Neural Inf. Process. Syst.*, 2010:658–666.

- Ganguli, D. and Simoncelli, E. P. (2014). Efficient sensory encoding and bayesian inference with heterogeneous neural populations. *Neural Comput.*, 26(10):2103–2134.
- Ganguli, D. and Simoncelli, E. P. (2016). Neural and perceptual signatures of efficient sensory coding. *arXiv preprint arXiv:1603.00058*.
- García-Piqueras, J., García-Mesa, Y., Cárcaba, L., Feito, J., Torres-Parejo, I., Martín-Biedma, B., Cobo, J., García-Suárez, O., and Vega, J. A. (2019). Ageing of the somatosensory system at the periphery: age-related changes in cutaneous mechanoreceptors. *J. Anat.*, 234(6):839–852.
- Garraghty, P. E., Florence, S. L., and Kaas, J. H. (1990). Ablations of areas 3a and 3b of monkey somatosensory cortex abolish cutaneous responsivity in area 1. *Brain Res.*, 528(1):165–169.
- Garraghty, P. E., Hanes, D. P., Florence, S. L., and Kaas, J. H. (1994). Pattern of peripheral deafferentation predicts reorganizational limits in adult primate somatosensory cortex. *Somatosens. Mot. Res.*, 11(2):109–117.
- Garraghty, P. E. and Kaas, J. H. (1991). Large-scale functional reorganization in adult monkey cortex after peripheral nerve injury. *Proc. Natl. Acad. Sci. U.S.A.*, 88(16):6976–6980.
- Garraghty, P. E., Pons, T. P., Sur, M., and Kaas, J. H. (1989). The arbors of axons terminating in middle cortical layers of somatosensory area 3b in owl monkeys. *Somatosens. Mot. Res.*, 6(4):401–411.
- Gelnar, P. A., Krauss, B. R., Szeverenyi, N. M., and Apkarian, A. V. (1998). Fingertip representation in the human somatosensory cortex: an fMRI study. *Neuroimage*, 7(4 Pt 1):261–283.
- Gilbert, C. D. and Li, W. (2012). Adult visual cortical plasticity. *Neuron*, 75(2):250–264.
- Giri, G. S., Maddahi, Y., and Zareinia, K. (2021). An Application-Based review of haptics technology. *Robotics*, 10(1):29.
- Godde, B., Stauffenberg, B., Spengler, F., and Dinse, H. R. (2000). Tactile coactivation-induced changes in spatial discrimination performance. *J. Neurosci.*, 20(4):1597–1604.
- Gonzalez, F., Gosselin, F., and Bachtá, W. (2014). Analysis of hand contact areas and interaction capabilities during manipulation and exploration. *IEEE Trans. Haptics*, 7(4):415–429.
- Goodchild, A. K., Ghosh, K. K., and Martin, P. R. (1996). Comparison of photoreceptor spatial density and ganglion cell morphology in the retina of human, macaque monkey, cat, and the marmoset callithrix jacchus. *J. Comp. Neurol.*, 366(1):55–75.
- Goodhill, G. J. (2007). Contributions of theoretical modeling to the understanding of neural map development. *Neuron*, 56(2):301–311.

- Grady, P., Tang, C., Twigg, C. D., Vo, M., Brahmbhatt, S., and Kemp, C. C. (2021). ContactOpt: Optimizing contact to improve grasps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1471–1481.
- Graham, D. J., Chandler, D. M., and Field, D. J. (2006). Can the theory of “whitening” explain the center-surround properties of retinal ganglion cell receptive fields? *Vision Res.*, 46(18):2901–2913.
- Graham, D. J. and Field, D. J. (2010). Natural images: Coding efficiency. In *Encyclopedia of Neuroscience*, pages 19–27.
- Grajski, K. A. (2016). Emergent spatial patterns of excitatory and inhibitory synaptic strengths drive somatotopic representational discontinuities and their plasticity in a computational model of primary sensory cortical area 3b. *Front. Comput. Neurosci.*, 10:72.
- Grajski, K. A. and Merzenich, M. M. (1990). Hebb-Type dynamics is sufficient to account for the inverse magnification rule in cortical somatotopy. *Neural Comput.*, 2(1):71–84.
- Grant, A. C., Fernandez, R., Shilian, P., Yanni, E., and Hill, M. A. (2006). Tactile spatial acuity differs between fingers: a study comparing two testing paradigms. *Percept. Psychophys.*, 68(8):1359–1362.
- Hämäläinen, H., Hiltunen, J., and Titievskaja, I. (2000). fMRI activations of SI and SII cortices during tactile stimulation depend on attention. *Neuroreport*, 11(8):1673–1676.
- Harding-Forrester, S. and Feldman, D. E. (2018). Chapter 4 - somatosensory maps. In Vallar, G. and Coslett, H. B., editors, *Handbook of Clinical Neurology*, volume 151, pages 73–102. Elsevier.
- Härtner, J., Strauss, S., Pfannmöller, J., and Lotze, M. (2021). Tactile acuity of fingertips and hand representation size in human area 3b and area 1 of the primary somatosensory cortex. *Neuroimage*, 232:117912.
- Harvey, B. M. and Dumoulin, S. O. (2011). The relationship between cortical magnification factor and population receptive field size in human visual cortex: constancies in cortical architecture. *J. Neurosci.*, 31(38):13604–13612.
- Hebb, D. O. (1949). *The Organization of Behaviour: A Neuropsychological Theory*, volume 44. New York: Science Editions.
- Herculano-Houzel, S. and Lent, R. (2005). Isotropic fractionator: a simple, rapid method for the quantification of total cell and neuron numbers in the brain. *J. Neurosci.*, 25(10):2518–2521.
- Hermundstad, A. M., Briguglio, J. J., Conte, M. M., Victor, J. D., Balasubramanian, V., and Tkačik, G. (2014). Variance predicts salience in central sensory processing. *Elife*, 3.

- Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.*, 160:106–154.
- Huerta, M. F. and Pons, T. P. (1990). Primary motor cortex receives input from area 3a in macaques. *Brain Res.*, 537(1-2):367–371.
- Huffman, K. J. and Krubitzer, L. (2001). Area 3a: Topographic organization and cortical connections in marmoset monkeys. *Cereb. Cortex*, 11(9):849–867.
- Huffman, K. J., Molnár, Z., Van Dellen, A., Kahn, D. M., Blakemore, C., and Krubitzer, L. (1999). Formation of cortical fields on a reduced cortical sheet. *J. Neurosci.*, 19(22):9939–9952.
- Hussain, I., Salvietti, G., Spagnoletti, G., and Prattichizzo, D. (2016). The Soft-SixthFinger: a wearable EMG controlled robotic Extra-Finger for grasp compensation in chronic stroke patients. *IEEE Robotics and Automation Letters*, 1(2):1000–1006.
- Hyvärinen, A., Hurri, J., and Hoyer, P. O. (2009). *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision*. Springer Science & Business Media.
- Hyvärinen, A. and Oja, E. (2000). Independent component analysis: algorithms and applications. *Neural Netw.*, 13(4-5):411–430.
- Hyvärinen, J. and Poranen, A. (1978). Receptive field integration and submodality convergence in the hand area of the post-central gyrus of the alert monkey. *J. Physiol.*, 283:539–556.
- Iwamura, Y., Tanaka, M., Sakamoto, M., and Hikosaka, O. (1983). Functional subdivisions representing different finger regions in area 3 of the first somatosensory cortex of the conscious monkey. *Exp. Brain Res.*, 51(3):315–326.
- Iwamura, Y., Tanaka, M., Sakamoto, M., and Hikosaka, O. (1985). Diversity in receptive field properties of vertical neuronal arrays in the crown of the postcentral gyrus of the conscious monkey. *Exp. Brain Res.*, 58(2):400–411.
- Iwamura, Y., Tanaka, M., Sakamoto, M., and Hikosaka, O. (1993). Rostrocaudal gradients in the neuronal receptive field complexity in the finger region of the alert monkey's postcentral gyrus. *Exp. Brain Res.*, 92(3):360–368.
- Jain, N., Florence, S. L., Qi, H. X., and Kaas, J. H. (2000). Growth of new brainstem connections in adult monkeys with massive sensory loss. *Proc. Natl. Acad. Sci. U.S.A.*, 97(10):5546–5550.
- Jenkins, W. M. and Merzenich, M. M. (1987). Reorganization of neocortical representations after brain injury: a neurophysiological model of the bases of recovery from stroke. *Prog. Brain Res.*, 71:249–266.

- Jenkins, W. M., Merzenich, M. M., Ochs, M. T., Allard, T., and Guíc-Robles, E. (1990). Functional reorganization of primary somatosensory cortex in adult owl monkeys after behaviorally controlled tactile stimulation. *J. Neurophysiol.*, 63(1):82–104.
- Jimenez Rodriguez, A. (2018). *The shape of the PCA trajectories and the population neural coding of movement initiation in the basal ganglia*. PhD thesis, University of Sheffield.
- Johansson, R. S. (1978). Tactile sensibility in the human hand: receptive field characteristics of mechanoreceptive units in the glabrous skin area. *J. Physiol.*, 281(1):101–125.
- Johansson, R. S. and Vallbo, A. B. (1979). Tactile sensibility in the human hand: relative and absolute densities of four types of mechanoreceptive units in glabrous skin. *J. Physiol.*, 286(1):283–300.
- Johansson, R. S. and Vallbo, Å. B. (1983). Tactile sensory coding in the glabrous skin of the human hand. *Trends Neurosci.*, 6:27–32.
- Johnson, K. O. (2000). Neural coding. *Neuron*, 26(3):563–566.
- Johnson, K. O. and Phillips, J. R. (1981). Tactile spatial resolution. i. two-point discrimination, gap detection, grating resolution, and letter recognition. *J. Neurophysiol.*, 46(6):1177–1192.
- Jonas, J. B., Müller-Bergh, J. A., Schlötzer-Schrehardt, U. M., and Naumann, G. O. (1990). Histomorphometry of the human optic nerve. *Invest. Ophthalmol. Vis. Sci.*, 31(4):736–744.
- Jones, E. G. (1975). Lamination and differential distribution of thalamic afferents within the sensory-motor cortex of the squirrel monkey. *J. Comp. Neurol.*, 160(2):167–203.
- Jones, E. G. (2003). Cortical and subcortical contributions to activity-dependent plasticity in primate somatosensory cortex. *Annu. Rev. Neurosci.*, 23:1–37.
- Jones, E. G. and Porter, R. (1980). What is area 3a? *Brain Res.*, 203(1):1–43.
- Jones, L. A. and Lederman, S. J. (2006). *Human hand function*. Oxford University Press.
- Kaas, J. H. (1983). What, if anything, is SI? organization of first somatosensory area of cortex. *Physiol. Rev.*, 63(1):206–231.
- Kaas, J. H., Nelson, R. J., Sur, M., Lin, C. S., and Merzenich, M. M. (1979). Multiple representations of the body within the primary somatosensory cortex of primates. *Science*, 204(4392):521–523.
- Kambi, N., Halder, P., Rajan, R., Arora, V., Chand, P., Arora, M., and Jain, N. (2014). Large-scale reorganization of the somatosensory cortex following spinal cord injuries is due to brainstem plasticity. *Nat. Commun.*, 5:3602.
- Karklin, Y. and Simoncelli, E. P. (2011). Efficient coding of natural images with a population of noisy Linear-Nonlinear neurons. *Adv. Neural Inf. Process. Syst.*, 24:999–1007.

- Karlsen, A. S. and Pakkenberg, B. (2011). Total numbers of neurons and glial cells in cortex and basal ganglia of aged brains with down syndrome—a stereological study. *Cereb. Cortex*, 21(11):2519–2524.
- Karmarkar, U. R. and Dan, Y. (2006). Experience-dependent plasticity in adult visual cortex. *Neuron*, 52(4):577–585.
- Keck, T., Toyozumi, T., Chen, L., Doiron, B., Feldman, D. E., Fox, K., Gerstner, W., Haydon, P. G., Hübener, M., Lee, H.-K., Lisman, J. E., Rose, T., Sengpiel, F., Stellwagen, D., Stryker, M. P., Turrigiano, G. G., and van Rossum, M. C. (2017). Integrating hebbian and homeostatic plasticity: the current state of the field and future research directions. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, 372(1715):20160158.
- Kersten, D. (1987). Predictability and redundancy of natural images. *J. Opt. Soc. Am.*, 4(12):2395.
- Kieliba, P., Clode, D., Maimon-Mor, R. O., and Makin, T. R. (2021). Robotic hand augmentation drives changes in neural body representation. *Sci Robot*, 6(54):2020.06.16.151944.
- Kikkert, S., Kolasinski, J., Jbabdi, S., Tracey, I., Beckmann, C. F., Johansen-Berg, H., and Makin, T. R. (2016). Revealing the neural fingerprints of a missing hand. *Elife*, 5:1–19.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biol. Cybern.*, 43(1):59–69.
- Kohonen, T. (1984). *Self-Organization and Associative Memory*. Springer, Berlin, Heidelberg.
- Kohonen, T. (1990). The self-organizing map. *Proc. IEEE*, 78(9):1464–1480.
- Kolasinski, J., Makin, T. R., Logan, J. P., Jbabdi, S., Clare, S., Stagg, C. J., and Johansen-Berg, H. (2016). Perceptually relevant remapping of human somatotopy in 24 hours. *Elife*, 5(digit 4):1–15.
- Krause, T., Kurth, R., Ruben, J., Schwiemann, J., Villringer, K., Deuchert, M., Moosmann, M., Brandt, S., Wolf, K., Curio, G., and Villringer, A. (2001). Representational overlap of adjacent fingers in multiple areas of human primary somatosensory cortex depends on electrical stimulus intensity: an fMRI study. *Brain Res.*, 899(1-2):36–46.
- Kriegeskorte, N., Mur, M., and Bandettini, P. (2008). Representational similarity analysis - connecting the branches of systems neuroscience. *Front. Syst. Neurosci.*, 2:4.
- Krubitzer, L., Huffman, K. J., Disbrow, E., and Recanzone, G. (2004). Organization of area 3a in macaque monkeys: Contributions to the cortical phenotype. *J. Comp. Neurol.*, 471(1):97–111.
- Krubitzer, L. a. and Kaas, J. H. (1988). Responsiveness and somatotopic organization of anterior parietal field 3b and adjoining cortex in newborn and infant monkeys. *Somatosens. Mot. Res.*, 6(2):179–205.

- Krubitzer, L. A. and Kaas, J. H. (1990). The organization and connections of somatosensory cortex in marmosets. *J. Neurosci.*, 10(3):952–974.
- Kruger, L., Siminoff, R., and Witkovsky, P. (1961). Single neuron analysis of dorsal column nuclei and spinal nucleus of trigeminal in cat. *J. Neurophysiol.*, 24:333–349.
- Kuehn, E. and Pleger, B. (2020). Encoding schemes in somatosensation: From micro- to metatopography. *Neuroimage*, 223:117255.
- Kurth, R., Villringer, K., Curio, G., Wolf, K. J., Krause, T., Repenthin, J., Schwiemann, J., Deuchert, M., and Villringer, A. (2000). fMRI shows multiple somatotopic digit representations in human primary somatosensory cortex. *Neuroreport*, 11(7):1487–1491.
- Kurth, R., Villringer, K., Curio, G., Wolf, K. J., Krause, T., Repenthin, J., Schwiemann, J., Deuchert, M., and Villringer, A. (2005). fMRI shows multiple somatotopic digit representations in human primary somatosensory cortex. *Neuroreport*, 11(7):1–5.
- Lambeta, M., Chou, P.-W., Tian, S., Yang, B., Maloon, B., Most, V. R., Stroud, D., Santos, R., Byagowi, A., Kammerer, G., et al. (2020). Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *IEEE Robotics and Automation Letters*, 5(3):3838–3845.
- Laskin, S. E. and Spencer, W. A. (1979). Cutaneous masking. II. geometry of excitatory and-inhibitory receptive fields of single units in somatosensory cortex of the cat. *J. Neurophysiol.*, 42(4):1061–1082.
- Lederman, S. J. and Klatzky, R. L. (1987). Hand movements: A window into haptic object recognition. *Cogn. Psychol.*, 19:342–368.
- Lederman, S. J. and Klatzky, R. L. (2009). Haptic perception: a tutorial. *Atten. Percept. Psychophys.*, 71(7):1439–1459.
- Lee, K. J. and Woolsey, T. A. (1975). A proportional relationship between peripheral innervation density and cortical neuron number in the somatosensory system of the mouse. *Brain Res.*, 99(2):349–353.
- Lee, M., Lee, S., Kim, J., Lim, J., and Yang, S. (2021). Graphene-electrode array for brain map remodeling of the cortical surface. *NPG Asia Materials*, 13(1).
- Lehnert, B. P., Santiago, C., Huey, E. L., Emanuel, A. J., Renauld, S., Africawala, N., Alkisar, I., Zheng, Y., Bai, L., Koutsioumpa, C., Hong, J. T., Magee, A. R., Harvey, C. D., and Ginty, D. D. (2021). Mechanoreceptor synapses in the brainstem shape the central representation of touch. *Cell*, 184(22):5608–5621.e18.

- Lenz, F. A., Dostrovsky, J. O., Tasker, R. R., Yamashiro, K., Kwan, H. C., and Murphy, J. T. (1988). Single-unit analysis of the human ventral thalamic nuclear group: somatosensory responses. *J. Neurophysiol.*, 59(2):299–316.
- Lewicki, M. S. (2002). Efficient coding of natural sounds. *Nat. Neurosci.*, 5(4):356–363.
- Li, P. and Liu, X. (2019). Common sensors in industrial robots: A review. *J. Phys. Conf. Ser.*, 1267(1):012036.
- Liao, C.-C., Gharbawie, O. A., Qi, H., and Kaas, J. H. (2013). Cortical connections to single digit representations in area 3b of somatosensory cortex in squirrel monkeys and prosimian galagos. *J. Comp. Neurol.*, 521(16):3768–3790.
- Lindsey, J., Ocko, S. A., Ganguli, S., and Deny, S. (2019). A unified theory of early visual representations from retina to cortex through anatomically constrained deep CNNs. *arXiv preprint arXiv:1901.00945*.
- Lisman, J. (2017). Glutamatergic synapses are structurally and biochemically complex because of multiple plasticity processes: long-term potentiation, long-term depression, short-term potentiation and scaling. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, 372(1715).
- Löken, L. S., Wessberg, J., Morrison, I., McGlone, F., and Olausson, H. (2009). Coding of pleasant touch by unmyelinated afferents in humans. *Nat. Neurosci.*, 12(5):547–548.
- Loutit, A. J., Vickery, R. M., and Potas, J. R. (2021). Functional organization and connectivity of the dorsal column nuclei complex reveals a sensorimotor integration and distribution hub. *J. Comp. Neurol.*, 529(1):187–220.
- Lutz, O. J. and Bensmaia, S. J. (2021). Proprioceptive representations of the hand in somatosensory cortex. *Current Opinion in Physiology*, 21:9–16.
- Machens, C. K., Gollisch, T., Kolesnikova, O., and Herz, A. V. M. (2005). Testing the efficiency of sensory coding with optimal stimulus ensembles. *Neuron*, 47(3):447–456.
- Makin, T. R. and Bensmaia, S. J. (2017). Stability of sensory topographies in adult cortex. *Trends Cogn. Sci.*, 21(3):195–204.
- Makin, T. R., Scholz, J., Henderson Slater, D., Johansen-Berg, H., and Tracey, I. (2015). Reassessing cortical reorganization in the primary sensorimotor cortex following arm amputation. *Brain*, 138(Pt 8):2140–2146.
- Maldjian, J. A., Gottschalk, A., Patel, R. S., Detre, J. A., and Alsop, D. C. (1999). The sensory somatotopic map of the human hand demonstrated at 4 tesla. *Neuroimage*, 10(1):55–62.
- Mancini, F., Bauleo, A., Cole, J., Lui, F., Porro, C. A., Haggard, P., and Iannetti, G. D. (2014). Whole-body mapping of spatial acuity for pain and touch. *Ann. Neurol.*, 75(6):917–924.

- Manfredi, L. R., Baker, A. T., Elias, D. O., Dammann, 3rd, J. F., Zielinski, M. C., Polashock, V. S., and Bensmaia, S. J. (2012). The effect of surface wave propagation on neural responses to vibration in primate glabrous skin. *PLoS One*, 7(2):e31203.
- Marik, S. A. and Gilbert, C. D. (2017). Cortical plasticity. In Chao, M. V., editor, *The Oxford Handbook of Developmental Neural Plasticity*. Oxford University Press.
- Martin, R. and Obermayer, K. (2009). Self-Organizing maps. In *Encyclopedia of Neuroscience*, pages 551–560. Elsevier.
- Martuzzi, R., van der Zwaag, W., Farthouat, J., Gruetter, R., and Blanke, O. (2014). Human finger somatotopy in areas 3b, 1, and 2: A 7T fMRI study using a natural stimulus. *Hum. Brain Mapp.*, 35(1):213–226.
- Mateos-Aparicio, P. and Rodríguez-Moreno, A. (2019). The impact of studying brain plasticity. *Front. Cell. Neurosci.*, 13:66.
- Mazza, M., de Pinho, M., Piqueira, J. R. C., and Roque, A. C. (2004). A dynamical model of fast cortical reorganization. *J. Comput. Neurosci.*, 16(2):177–201.
- Mazza, M. B. and Roque-da Silva, A. C. (1999). Computational model of topographic reorganization in somatosensory cortex in response to digit lesions. *Neurocomputing*, 26-27:435–441.
- McGlone, F., Kelly, E. F., Trulsson, M., Francis, S. T., Westling, G., and Bowtell, R. (2002). Functional neuroimaging studies of human somatosensory cortex. In *Behavioural Brain Research*, volume 135, pages 147–158.
- McLaughlin, T. and O’Leary, D. D. M. (2005). Molecular gradients and development of retinotopic maps. *Annu. Rev. Neurosci.*, 28(1):327–355.
- Mehring, C., Akselrod, M., Bashford, L., Mace, M., Choi, H., Blüher, M., Buschhoff, A.-S., Pistoohl, T., Salomon, R., Cheah, A., Blanke, O., Serino, A., and Burdet, E. (2019). Augmented manipulation ability in humans with six-fingered hands. *Nat. Commun.*, 10(1):2401.
- Meltzoff, A. N., Saby, J. N., and Marshall, P. J. (2019). Neural representations of the body in 60-day-old human infants. *Dev. Sci.*, 22(1):e12698.
- Merzenich, M. M. and Brugge, J. F. (1973). Representation of the cochlear partition on the superior temporal plane of the macaque monkey. *Brain Res.*, 50(2):275–296.
- Merzenich, M. M. and Jenkins, W. M. (1993). Reorganization of cortical representations of the hand following alterations of skin inputs induced by nerve injury, skin island transfers, and experience. *J. Hand Ther.*, 6(2):89–104.

- Merzenich, M. M., Kaas, J. H., Sur, M., and Lin, C. S. (1978). Double representation of the body surface within cytoarchitectonic areas 3b and 1 in “SI” in the owl monkey (*aotus trivirgatus*). *J. Comp. Neurol.*, 181(1):41–73.
- Merzenich, M. M., Kaas, J. H., Wall, J., Nelson, R. J., Sur, M., and Felleman, D. (1983). Topographic reorganization of somatosensory cortical areas 3b and 1 in adult monkeys following restricted deafferentation. *Neuroscience*, 8(1):33–55.
- Merzenich, M. M., Nelson, R. J., Kaas, J. H., Stryker, M. P., Jenkins, W. M., Zook, J. M., Cynader, M. S., and Schoppmann, A. (1987). Variability in hand surface representations in areas 3b and 1 in adult owl and squirrel monkeys. *J. Comp. Neurol.*, 258(2):281–296.
- Merzenich, M. M., Nelson, R. J., Stryker, M. P., Cynader, M. S., Schoppmann, A., and Zook, J. M. (1984). Somatosensory cortical map changes following digit amputation in adult monkeys. *J. Comp. Neurol.*, 224(4):591–605.
- Miikkulainen, R. (1991). Self-organizing process based on lateral inhibition and synaptic resource redistribution. In *Proceedings of the International Conference on Artificial Neural Networks*, pages 415–420.
- Moerel, M., De Martino, F., and Formisano, E. (2012). Processing of natural sounds in human auditory cortex: tonotopy, spectral tuning, and relation to voice sensitivity. *J. Neurosci.*, 32(41):14205–14216.
- Muret, D. and Makin, T. R. (2021). The homeostatic homunculus: rethinking deprivation-triggered reorganisation. *Curr. Opin. Neurobiol.*, 67:115–122.
- Muret, D., Root, V., Kieliba, P., Clode, D., and Makin, T. R. (2021). Beyond body maps: information content of specific body parts is distributed across the somatosensory homunculus. *Available at SSRN 3931628*.
- Napier, J. R. (1956). The prehensile movements of the human hand. *J. Bone Joint Surg. Br.*, 38-B(4):902–913.
- Négyessy, L., Pálfi, E., Ashaber, M., Palmer, C., Jákli, B., Friedman, R. M., Chen, L. M., and Roe, A. W. (2013). Intrinsic horizontal connections process global tactile features in the primary somatosensory cortex: neuroanatomical evidence. *J. Comp. Neurol.*, 521(12):2798–2817.
- Nelken, I., Rotman, Y., and Bar Yosef, O. (1999). Responses of auditory-cortex neurons to structural features of natural sounds. *Nature*, 397(6715):154–157.
- Nelson, A. J. and Chen, R. (2008). Digit somatotopy within cortical areas of the postcentral gyrus in humans. *Cereb. Cortex*, 18(10):2341–2351.
- Nelson, R. J., Sur, M., Felleman, D. J., and Kaas, J. H. (1980). Representations of the body surface in postcentral parietal cortex of macaca fascicularis. *J. Comp. Neurol.*, 192(4):611–643.

- Obermayer, K., Ritter, H., and Schulten, K. (1990a). Large-scale simulation of a self-organizing neural network: Formation of a somatotopic map. *Parallel Processing in Neural Systems and Computers*, pages 71–74.
- Obermayer, K., Ritter, H., and Schulten, K. (1990b). Large-scale simulations of self-organizing neural networks on parallel computers: application to biological modelling. *Parallel Comput.*, 14(3):381–404.
- Olshausen, B. A., , and Field, D. J. (1996). Wavelet-like receptive fields emerge from a network that learns sparse codes for natural images. *Nature*, 381:607–609.
- Olshausen, B. A. and Field, D. J. (1997). Sparse coding with an overcomplete basis set: a strategy employed by v1? *Vision Res.*, 37(23):3311–3325.
- Olshausen, B. A. and Field, D. J. (2004). Sparse coding of sensory inputs. *Curr. Opin. Neurobiol.*, 14(4):481–487.
- Orczyk, J. J. and Garraghty, P. E. (2015). Reconciling homeostatic and use-dependent plasticity in the context of somatosensory deprivation. *Neural Plast.*, 2015:290819.
- Overduin, S. A. and Servos, P. (2004). Distributed digit somatotopy in primary somatosensory cortex. *Neuroimage*, 23(2):462–472.
- Padberg, J., Cerkevich, C., Engle, J., Rajan, A. T., Recanzone, G., Kaas, J., and Krubitzer, L. (2009). Thalamocortical connections of parietal somatosensory cortical fields in macaque monkeys are highly divergent and convergent. *Cereb. Cortex*, 19(9):2038–2064.
- Palmer, S. E., Marre, O., Berry, 2nd, M. J., and Bialek, W. (2015). Predictive information in a sensory population. *Proc. Natl. Acad. Sci. U.S.A.*, 112(22):6908–6913.
- Pantev, C., Wollbrink, A., Roberts, L. E., Engelen, A., and Lütkenhöner, B. (1999). Short-term plasticity of the human auditory cortex. *Brain Res.*, 842(1):192–199.
- Paré, M., Smith, A. M., and Rice, F. L. (2002). Distribution and terminal arborizations of cutaneous mechanoreceptors in the glabrous finger pads of the monkey. *J. Comp. Neurol.*, 445(4):347–359.
- Pascual-Leone, A. and Torres, F. (1993). Plasticity of the sensorimotor cortex representation of the reading finger in braille readers. *Brain*, 116 (Pt 1):39–52.
- Pawling, R., Cannon, P. R., McGlone, F. P., and Walker, S. C. (2017). C-tactile afferent stimulating touch carries a positive affective value. *PLoS One*, 12(3):e0173457.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830.

- Penfield, W. and Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain*, 60(4):389–443.
- Penfield, W. and Rasmussen, T. (1950). *The cerebral cortex of man; a clinical study of localization of function*, volume 248. Macmillan The cerebral cortex of man; a clinical study of localization of function., Oxford, England.
- Perry, V. H., Oehler, R., and Cowey, A. (1984). Retinal ganglion cells that project to the dorsal lateral geniculate nucleus in the macaque monkey. *Neuroscience*, 12(4):1101–1123.
- Peters, R. M. and Goldreich, D. (2013). Tactile spatial acuity in childhood: effects of age and fingertip size. *PLoS One*, 8(12):e84650.
- Peters, R. M., Hackeman, E., and Goldreich, D. (2009). Diminutive digits discern delicate details: fingertip size and the sex difference in tactile spatial acuity. *J. Neurosci.*, 29(50):15756–15761.
- Petersen, C. C. H. (2007). The functional organization of the barrel cortex. *Neuron*, 56(2):339–355.
- Plumbley, M. D. (1999). Do cortical maps adapt to optimize information density? *Network: Computation in Neural Systems*, 10(1):41–58.
- Pons, T. P., Garraghty, P. E., Cusick, C. G., and Kaas, J. H. (1985). The somatotopic organization of area 2 in macaque monkeys. *J. Comp. Neurol.*, 241(4):445–466.
- Pons, T. P., Garraghty, P. E., Ommaya, A. K., Kaas, J. H., Taub, E., and Mishkin, M. (1991). Massive cortical reorganization after sensory deafferentation in adult macaques. *Science*, 252(5014):1857–1860.
- Pons, T. P., Wall, J. T., Garraghty, P. E., Cusick, C. G., and Kaas, J. H. (1987). Consistent features of the representation of the hand in area 3b of macaque monkeys. *Somatosens. Res.*, 4(4):309–331.
- Pozo, K. and Goda, Y. (2010). Unraveling mechanisms of homeostatic synaptic plasticity. *Neuron*, 66(3):337–351.
- Provis, J. M., van Driel, D., Billson, F. A., and Russell, P. (1985). Human fetal optic nerve: overproduction and elimination of retinal axons during development. *J. Comp. Neurol.*, 238(1):92–100.
- Purves, D., Augustine, G. J., Fitzpatrick, D., Hall, W. C., Lamantia, A.-S., and White, L. E. (2018). *Neuroscience*. Sinauer Associates, 6 edition.
- Qi, H.-X., Gharbawie, O. A., Wong, P., and Kaas, J. H. (2011). Cell-poor septa separate representations of digits in the ventroposterior nucleus of the thalamus in monkeys and prosimian galagos. *J. Comp. Neurol.*, 519(4):738–758.
- Qi, H.-X. and Kaas, J. H. (2006). Organization of primary afferent projections to the gracile nucleus of the dorsal column system of primates. *J. Comp. Neurol.*, 499(2):183–217.

- Qi, H.-X., Preuss, T. M., and Kaas, J. H. (2008). Somatosensory areas of the cerebral cortex: architectonic characteristics and modular organization.
- Ralph, F., Large, D. R., Burnett, G., Lang, A., and Morris, A. (2021). U can't touch this! face touching behaviour whilst driving: Implications for health, hygiene and human factors. *Ergonomics*, pages 1–29.
- Randolph, M. and Semmes, J. (1974). Behavioral consequences of selective subtotal ablations in the postcentral gyrus of macaca mulatta. *Brain Res.*, 70(1):55–70.
- Rasmusson, D. D. (1982). Reorganization of raccoon somatosensory cortex following removal of the fifth digit. *J. Comp. Neurol.*, 205(4):313–326.
- Raspopovic, S., Valle, G., and Petrini, F. M. (2021). Sensory feedback for limb prostheses in amputees. *Nat. Mater.*, 20(7):925–939.
- Rauschecker, J. P. (1999). Auditory cortical plasticity: a comparison with other sensory systems. *Trends Neurosci.*, 22(2):74–80.
- Recanzone, G. H., Merzenich, M. M., and Schreiner, C. E. (1992). Changes in the distributed temporal response properties of SI cortical neurons reflect improvements in performance on a temporally based tactile discrimination task. *J. Neurophysiol.*, 67(5):1071–1091.
- Recanzone, G. H., Schreiner, C. E., and Merzenich, M. M. (1993). Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. *J. Neurosci.*, 13(1):87–103.
- Reed, J. L., Pouget, P., Qi, H.-X., Zhou, Z., Bernard, M. R., Burish, M. J., Haitas, J., Bonds, A. B., and Kaas, J. H. (2008). Widespread spatial integration in primary somatosensory cortex. *Proc. Natl. Acad. Sci. U.S.A.*, 105(29):10233–10237.
- Rigato, S., Begum Ali, J., Van Velzen, J., and Bremner, A. J. (2014). The neural basis of somatosensory remapping develops in human infancy. *Curr. Biol.*, 24(11):1222–1226.
- Ritter, H. (1990). Self-organizing maps for internal representations. *Psychol. Res.*, 52(2-3):128–136.
- Ritter, H., Martinetz, T., Schulden, K., Barsky, D., and Tesch, M. (1992). *Neural computation and self-organizing maps: an introduction*. Reading, MA: Addison-Wesley.
- Ritter, H. and Schulden, K. (1986). On the stationary state of kohonen's self-organizing sensory mapping. *Biol. Cybern.*, 54(2):99–106.
- Rojas, R. (1996). Kohonen networks. In Rojas, R., editor, *Neural Networks: A Systematic Introduction*, pages 389–410. Springer Berlin Heidelberg, Berlin, Heidelberg.

- Rossini, P. M., Martino, G., Narici, L., Pasquarelli, A., Peresson, M., Pizzella, V., Tecchio, F., Torrioli, G., and Romani, G. L. (1994). Short-term brain ‘plasticity’ in humans: transient finger representation changes in sensory cortex somatotopy following ischemic anesthesia. *Brain Res.*, 642(1):169–177.
- Roux, F.-E., Djidjeli, I., and Durand, J.-B. (2018). Functional architecture of the somatosensory homunculus detected by electrostimulation. *J. Physiol.*, 596(5):941–956.
- Rubel, E. W. (1971). A comparison of somatotopic organization in sensory neocortex of newborn kittens and adult cats. *J. Comp. Neurol.*, 143(4):447–480.
- Saadon-Grosman, N., Loewenstein, Y., and Arzy, S. (2020). The ‘creatures’ of the human cortical somatosensory system. *Brain Commun*, 2(1).
- Saal, H. P. and Bensmaia, S. J. (2014). Touch is a team effort: interplay of submodalities in cutaneous sensibility. *Trends Neurosci.*, 37(12):689–697.
- Saal, H. P., Delhaye, B. P., Rayhaun, B. C., and Bensmaia, S. J. (2017). Simulating tactile signals from the whole hand with millisecond precision. *Proc. Natl. Acad. Sci. U.S.A.*, 114(28):E5693–E5702.
- Saby, J. N., Meltzoff, A. N., and Marshall, P. J. (2015). Neural body maps in human infants: Somatotopic responses to tactile stimulation in 7-month-olds. *Neuroimage*, 118:74–78.
- Sachdev, R. N. S. and Catania, K. C. (2002). Receptive fields and response properties of neurons in the star-nosed mole’s somatosensory fovea. *J. Neurophysiol.*, 87(5):2602–2611.
- Sagisaka, T., Ohmura, Y., Kuniyoshi, Y., Nagakubo, A., and Ozaki, K. (2011). High-density conformable tactile sensing glove. In *2011 11th IEEE-RAS International Conference on Humanoid Robots*, pages 537–542.
- Sakai, K., Watanabe, E., Onodera, Y., Itagaki, H., Yamamoto, E., Koizumi, H., and Miyashita, Y. (1995). Functional mapping of the human somatosensory cortex with echo-planar MRI. *Magn. Reson. Med.*, 33(5):736–743.
- Salisbury, J. M. and Palmer, S. E. (2016). Optimal prediction in the retina and natural motion statistics. *J. Stat. Phys.*, 162(5):1309–1323.
- Sanchez-Panchuelo, R. M., Besle, J., Beckett, A., Bowtell, R., Schluppeck, D., and Francis, S. (2012). Within-digit functional parcellation of brodmann areas of the human primary somatosensory cortex using functional magnetic resonance imaging at 7 tesla. *J. Neurosci.*, 32(45):15815–15822.
- Sánchez-Panchuelo, R. M., Besle, J., Mougin, O., Gowland, P., Bowtell, R., Schluppeck, D., and Francis, S. (2014). Regional structural differences across functionally parcellated brodmann areas of human primary somatosensory cortex. *Neuroimage*, 93:221–230.

- Sanchez Panchuelo, R. M., Besle, J., Schluppeck, D., Humberstone, M., and Francis, S. (2018). Somatotopy in the human somatosensory system. *Front. Hum. Neurosci.*, 12:235.
- Sanchez-Panchuelo, R. M., Francis, S., Bowtell, R., and Schluppeck, D. (2010). Mapping human somatosensory cortex in individual subjects with 7T functional MRI. *J. Neurophysiol.*, 103(5):2544–2556.
- Sanchez Panchuelo, R. M., Schluppeck, D., Francis, S. T., and Bowtell, R. (2008). Mapping human somatosensory cortex with fMRI at 1 mm isotropic resolution. *Proc. Intl. Soc. Mag. Reson. Med.* 16, 462(1):2522–2522.
- Sandell, J. H. and Peters, A. (2001). Effects of age on nerve fibers in the rhesus monkey optic nerve. *J. Comp. Neurol.*, 429(4):541–553.
- Sawyer, E. K. and Catania, K. C. (2016). Somatosensory organ topography across the star of the star-nosed mole (*condylura cristata*). *J. Comp. Neurol.*, 524(5):917–929.
- Saxe, A. M., Bansal, Y., Dapello, J., Advani, M., Kolchinsky, A., Tracey, B. D., and Cox, D. D. (2019). On the information bottleneck theory of deep learning*. *J. Stat. Mech.*, 2019(12):124020.
- Schluppeck, D., Sanchez-Panchuelo, R.-M., and Francis, S. T. (2018). Exploring structure and function of sensory cortex with 7T MRI. *Neuroimage*, 164:10–17.
- Schreiner, C. E. and Winer, J. A. (2007). Auditory cortex mapmaking: principles, projections, and plasticity. *Neuron*, 56(2):356–365.
- Schweisfurth, M. A., Frahm, J., Farina, D., and Schweizer, R. (2018). Comparison of fMRI digit representations of the dominant and non-dominant hand in the human primary somatosensory cortex. *Front. Hum. Neurosci.*, 12:492.
- Schweisfurth, M. a., Frahm, J., and Schweizer, R. (2014). Individual fMRI maps of all phalanges and digit bases of all fingers in human primary somatosensory cortex. *Front. Hum. Neurosci.*, 8(September):658–658.
- Schweisfurth, M. A., Schweizer, R., and Frahm, J. (2011). Functional MRI indicates consistent intra-digit topographic maps in the little but not the index finger within the human primary somatosensory cortex. *Neuroimage*, 56(4):2138–2143.
- Schweizer, R., Voit, D., and Frahm, J. (2008). Finger representations in human primary somatosensory cortex as revealed by high-resolution functional MRI of tactile stimulation. *Neuroimage*, 42(1):28–35.
- Seelke, A. M. H., Padberg, J. J., Disbrow, E., Purnell, S. M., Recanzone, G., and Krubitzer, L. (2012). Topographic maps within brodmann’s area 5 of macaque monkeys. *Cereb. Cortex*, 22(8):1834–1850.

- Sefton, A. J., Horsburgh, G. M., and Lam, K. (1985). The development of the optic nerve in rodents. *Aust. N. Z. J. Ophthalmol.*, 13(2):135–145.
- Shafti, A., Haar, S., Mio, R., Guilleminot, P., and Faisal, A. A. (2021). Playing the piano with a robotic third thumb: assessing constraints of human augmentation. *Sci. Rep.*, 11(1):21375.
- Shan, D., Geng, J., Shu, M., and Fouhey, D. F. (2020). Understanding human hands in contact at internet scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9869–9878.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423.
- Shao, Y., Hayward, V., and Visell, Y. (2016). Spatial patterns of cutaneous vibration during whole-hand haptic interactions. *Proc. Natl. Acad. Sci. U.S.A.*, 113(15):4188–4193.
- Shao, Y., Hayward, V., and Visell, Y. (2020). Compression of dynamic tactile information in the human hand. *Sci Adv*, 6(16):eaaz1158.
- Shively, C. and Mitchell, G. (1986). Perinatal behavior of anthropoid primates. *Comparative primate biology.*, 2(Part A):245–294.
- Shouno, H. and Kurata, K. (2001). Formation of a direction map by projection learning using kohonen’s self-organization map. *Biol. Cybern.*, 85(4):241–246.
- Shwartz-Ziv, R. and Tishby, N. (2017). Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*.
- Silvera-Tawil, D., Rye, D., and Velonaki, M. (2015). Artificial skin and tactile sensing for socially interactive robots: A review. *Rob. Auton. Syst.*, 63:230–243.
- Simoncelli, E. P. (2003). Vision and the statistics of the visual environment. *Curr. Opin. Neurobiol.*, 13(2):144–149.
- Simoncelli, E. P. and Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annu. Rev. Neurosci.*, 24(1):1193–1216.
- Simoncelli, E. P. and Schwartz, O. (1999). Modeling surround suppression in V1 neurons with a statistically-derived normalization model. In *Advances in neural information processing systems 11: Proceedings of the 1998 conference*, volume 11, page 153. books.google.com.
- Singer, Y., Teramoto, Y., Willmore, B. D., Schnupp, J. W., King, A. J., and Harper, N. S. (2018). Sensory cortex is optimized for prediction of future input. *Elife*, 7.
- Sirosh, J. and Miikkulainen, R. (1994). Cooperative self-organization of afferent and lateral connections in cortical maps. *Biol. Cybern.*, 71(1):65–78.

- Smith, M. C. and Deacon, P. (1984). Topographical anatomy of the posterior columns of the spinal cord in man. the long ascending fibres. *Brain*, 107 (Pt 3):671–698.
- Sripati, A. P., Bensmaia, S. J., and Johnson, K. O. (2006). A continuum mechanical model of mechanoreceptive afferent responses to indented spatial patterns. *J. Neurophysiol.*, 95(6):3852–3864.
- Stafford, T. and Wilson, S. P. (2007). Self-organisation can generate the discontinuities in the somatosensory map. *Neurocomputing*, 70(10):1932–1937.
- Stevens, J.-L. R., Law, J. S., Antolík, J., and Bednar, J. A. (2013). Mechanisms for stable, robust, and adaptive development of orientation maps in the primary visual cortex. *J. Neurosci.*, 33(40):15747–15766.
- Strauss, W. A. (2007). *Partial Differential Equations: An Introduction*. John Wiley & Sons.
- Stringer, Ann, E., Qiao, P. G., Friedman, R. M., Holroyd, L., Newton, A. T., Gore, J. C., and Chen, L. M. (2014). Distinct fine-scale fMRI activation patterns of contra- and ipsilateral somatosensory areas 3b and 1 in humans. *Hum. Brain Mapp.*, 35(9):4841–4857.
- Stringer, E. A., Chen, L. M., Friedman, R. M., Gatenby, C., and Gore, J. C. (2011). Differentiation of somatosensory cortices by high-resolution fMRI at 7T. *Neuroimage*, 54(2):1012–1020.
- Sundaram, S., Kellnhofer, P., Li, Y., Zhu, J.-Y., Torralba, A., and Matusik, W. (2019). Learning the signatures of the human grasp using a scalable tactile glove. *Nature*, 569(7758):698–702.
- Sur, M., Garraghty, P. E., and Bruce, C. J. (1985). Somatosensory cortex in macaque monkeys: laminar differences in receptive field size in areas 3b and 1. *Brain Res.*, 342(2):391–395.
- Sur, M., Merzenich, M. M., and Kaas, J. H. (1980). Magnification, receptive-field area, and “hypercolumn” size in areas 3b and 1 of somatosensory cortex in owl monkeys. *J. Neurophysiol.*, 44(2):295–311.
- Sur, M., Nelson, R. J., and Kaas, J. H. (1982). Representations of the body surface in cortical areas 3b and 1 of squirrel monkeys: comparisons with other primates. *J. Comp. Neurol.*, 211(2):177–192.
- Sur, M., Wall, J. T., and Kaas, J. H. (1984). Modular distribution of neurons with slowly adapting and rapidly adapting responses in area 3b of somatosensory cortex in monkeys. *J. Neurophysiol.*, 51(4):724–744.
- Suresh, A. K., Winberry, J. E., Versteeg, C., Chowdhury, R., Tomlinson, T., Rosenow, J. M., Miller, L. E., and Bensmaia, S. J. (2017). Methodological considerations for a chronic neural interface with the cuneate nucleus of macaques. *J. Neurophysiol.*, 118(6):3271–3281.

- Sutton, III, G. G., Reggia, J. A., Armentrout, S. L., and D'Autrechy, C. L. (1994). Cortical map reorganization as a competitive process. *Neural Comput.*, 6(1):1–13.
- Swindale, N. V. and Bauer, H.-U. (1998). Application of kohonen's self-organizing feature map algorithm to cortical maps of orientation and direction preference. *Proceedings of the Royal Society B: Biological Sciences*, 265(1398):827.
- Takeuchi, A. and Amari, S. (1979). Formation of topographic maps and columnar microstructures in nerve fields. *Biol. Cybern.*, 35(2):63–72.
- Taylor, C. L. and Schwarz, R. J. (1955). The anatomy and mechanics of the human hand. *Artif. Limbs*, 2(2):22–35.
- Tegenthoff, M., Ragert, P., Pleger, B., Schwenkreis, P., Förster, A.-F., Nicolas, V., and Dinse, H. R. (2005). Improvement of tactile discrimination performance and enlargement of cortical somatosensory maps after 5 hz rTMS. *PLoS Biol.*, 3(11):e362.
- Teşileanu, T., Cocco, S., Monasson, R., and Balasubramanian, V. (2019). Adaptation of olfactory receptor abundances for efficient coding. *Elife*, 8.
- Tishby, N., Pereira, F. C., and Bialek, W. (2000). The information bottleneck method. *arXiv preprint physics/0004057*.
- Tishby, N. and Zaslavsky, N. (2015). Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pages 1–5.
- Tiwana, M. I., Redmond, S. J., and Lovell, N. H. (2012). A review of tactile sensing technologies with applications in biomedical engineering. *Sens. Actuators A Phys.*, 179:17–31.
- Turrigiano, G. (2012). Homeostatic synaptic plasticity: local and global mechanisms for stabilizing neuronal function. *Cold Spring Harb. Perspect. Biol.*, 4(1):a005736.
- Turrigiano, G. G. (2008). The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell*, 135(3):422–435.
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391(6670):892–896.
- Turrigiano, G. G. and Nelson, S. B. (2000). Hebb and homeostasis in neuronal plasticity. *Curr. Opin. Neurobiol.*, 10(3):358–364.
- Vallbo, A. B. and Johansson, R. S. (1984). Properties of cutaneous mechanoreceptors in the human hand related to touch sensation. *Hum. Neurobiol.*, 3(1):3–14.
- Van Boven, R. W., Hamilton, R. H., Kauffman, T., Keenan, J. P., and Pascual-Leone, A. (2000). Tactile spatial resolution in blind braille readers. *Neurology*, 54(12):2230–2236.

- van Westen, D., Fransson, P., Olsrud, J., Rosén, B., Lundborg, G., Larsson, E.-M., and Van Westen, D. (2004). Fingersomatotopy in area 3b: an fMRI-study. *BMC Neurosci.*, 5(28):28.
- Vanderhaeghen, P., Lu, Q., Prakash, N., Frisé, J., Walsh, C. A., Frostig, R. D., and Flanagan, J. G. (2000). A mapping label required for normal scale of body representation in the cortex. *Nat. Neurosci.*, 3(4):358–365.
- Vega-Bermudez, F. and Johnson, K. O. (2001). Differences in spatial acuity between digits. *Neurology*, 56(10):1389–1391.
- Verendeev, A., Thomas, C., McFarlin, S. C., Hopkins, W. D., Phillips, K. A., and Sherwood, C. C. (2015). Comparative analysis of meissner’s corpuscles in the fingertips of primates. *J. Anat.*, 227(1):72–80.
- Vidyasagar, R., Folger, S. E., and Parkes, L. M. (2014). Re-wiring the brain: Increased functional connectivity within primary somatosensory cortex following synchronous co-activation. *Neuroimage*, 92:19–26.
- Vitureira, N. and Goda, Y. (2013). Cell biology in neuroscience: the interplay between hebbian and homeostatic synaptic plasticity. *J. Cell Biol.*, 203(2):175–186.
- Wang, X., Merzenich, M. M., Sameshima, K., and Jenkins, W. M. (1995). Remodelling of hand representation in adult cortex determined by timing of tactile stimulation. *Nature*, 378(6552):71–75.
- Ward-Cherrier, B., Pestell, N., Cramphorn, L., Winstone, B., Giannaccini, M. E., Rossiter, J., and Lepora, N. F. (2018). The TacTip family: Soft optical tactile sensors with 3D-Printed biomimetic morphologies. *Soft Robot*, 5(2):216–227.
- Wässle, H., Grünert, U., Röhrenbeck, J., and Boycott, B. B. (1990). Retinal ganglion cell density and cortical magnification factor in the primate. *Vision Res.*, 30(11):1897–1911.
- Wei, X.-X. and Stocker, A. A. (2016). Mutual information, fisher information, and efficient coding. *Neural Comput.*, 28(2):305–326.
- Weinrich, M., Armentrout, S., and Reggia, J. (1995). A neural model of recovery from lesions in the somatosensory system. *J. Neurol. Rehabil.*, 9(1):25–32.
- Wells-Gray, E. M., Choi, S. S., Bries, A., and Doble, N. (2016). Variation in rod and cone density from the fovea to the mid-periphery in healthy human retinas using adaptive optics scanning laser ophthalmoscopy. *Eye*, 30(8):1135–1143.
- Wesselink, D. B., Sanders, Z.-B., Edmondson, L. R., Dempsey-Jones, H., Kieliba, P., Kikkert, S., Themistocleous, A. C., Emir, U., Diedrichsen, J., Saal, H. P., and Makin, T. R. (2020). Malleability of the cortical hand map following a finger nerve block.

- Wesselink, D. B., van den Heiligenberg, F. M., Ejaz, N., Dempsey-Jones, H., Cardinali, L., Tarall-Jozwiak, A., Diedrichsen, J., and Makin, T. R. (2019). Obtaining and maintaining cortical hand representation as evidenced from acquired and congenital handlessness. *Elife*, 8.
- Wiesel, T. N. and Hubel, D. H. (1963). Single-cell responses in striate cortex of kittens deprived of vision in one eye. *J. Neurophysiol.*, 26:1003–1017.
- Wiesel, T. N. and Hubel, D. H. (1965). Comparison of the effects of unilateral and bilateral eye closure on cortical unit responses in kittens. *J. Neurophysiol.*, 28(6):1029–1040.
- Willshaw, D. J. and von der Malsburg, C. (1976). How patterned neural connections can be set up by self-organization. *Proc. R. Soc. Lond. B Biol. Sci.*, 194(1117):431–445.
- Wilson, S. P. and Bednar, J. A. (2015). What, if anything, are topological maps for? *Dev. Neurobiol.*, 75(6):667–681.
- Witham, C. L. and Baker, S. N. (2011). Modulation and transmission of peripheral inputs in monkey cuneate and external cuneate nuclei. *J. Neurophysiol.*, 106(5):2764–2775.
- Witney, A. G., Wing, A., Thonnard, J.-L., and Smith, A. M. (2004). The cutaneous contribution to adaptive precision grip. *Trends Neurosci.*, 27(10):637–643.
- Wong, M., Peters, R. M., and Goldreich, D. (2013). A physical constraint on perceptual learning: tactile spatial acuity improves with training to a limit set by finger size. *J. Neurosci.*, 33(22):9345–9352.
- Woods, T. M., Cusick, C. G., Pons, T. P., Taub, E., and Jones, E. G. (2000). Progressive transneuronal changes in the brainstem and thalamus after long-term dorsal rhizotomies in adult macaque monkeys. *J. Neurosci.*, 20(10):3884–3899.
- Woolsey, T. A. and Van der Loos, H. (1970). The structural organization of layer IV in the somatosensory region (SI) of mouse cerebral cortex. the description of a cortical field composed of discrete cytoarchitectonic units. *Brain Res.*, 17(2):205–242.
- Xerri, C., Coq, J. O., Merzenich, M. M., and Jenkins, W. M. (1996). Experience-induced plasticity of cutaneous maps in the primary somatosensory cortex of adult monkeys and rats. *J. Physiol. Paris*, 90(3-4):277–287.
- Xerri, C., Merzenich, M. M., Jenkins, W., and Santucci, S. (1999). Representational plasticity in cortical area 3b paralleling tactual-motor skill acquisition in adult monkeys. *Cereb. Cortex*, 9(3):264–276.
- Xu, J. and Wall, J. T. (1999). Functional organization of tactile inputs from the hand in the cuneate nucleus and its relationship to organization in the somatosensory cortex. *J. Comp. Neurol.*, 411(3):369–389.

- Yamins, D. L. K., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., and DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proc. Natl. Acad. Sci. U.S.A.*, 111(23):8619–8624.
- Yang, T.-H., Kim, J. R., Jin, H., Gil, H., Koo, J.-H., and Kim, H. J. (2021). Recent advances and opportunities of active materials for haptic technologies in virtual and augmented reality. *Adv. Funct. Mater.*, 31(39):2008831.
- Yarrow, S., Challis, E., and Seriès, P. (2012). Fisher and shannon information in finite neural populations. *Neural Comput.*, 24(7):1740–1780.
- Yau, J. M., Kim, S. S., Thakur, P. H., and Bensmaia, S. J. (2016). Feeling form: the neural basis of haptic shape perception. *J. Neurophysiol.*, 115(2):631–642.
- Yerxa, T. E., Kee, E., DeWeese, M. R., and Cooper, E. A. (2020). Efficient sensory coding of multidimensional stimuli. *PLoS Comput. Biol.*, 16(9):e1008146.
- Yu, Y., Huber, L., Yang, J., Jangraw, D. C., Handwerker, D. A., Molfese, P. J., Chen, G., Ejima, Y., Wu, J., and Bandettini, P. A. (2019). Layer-specific activation of sensory input and predictive feedback in the human primary somatosensory cortex. *Sci Adv*, 5(5):eaav9053.
- Zeharia, N., Hertz, U., Flash, T., and Amedi, A. (2015). New whole-body sensory-motor gradients revealed using phase-locked analysis and verified using multivoxel pattern analysis and functional connectivity. *J. Neurosci.*, 35(7):2845–2859.
- Zenke, F., Gerstner, W., and Ganguli, S. (2017). The temporal paradox of hebbian learning and homeostatic plasticity. *Curr. Opin. Neurobiol.*, 43:166–176.
- Zhang, J. (1991). Dynamics and formation of Self-Organizing maps. *Neural Comput.*, 3(1):54–66.
- Zheng, J. Z., Rosa, S. D. L., Dollar, A. M., and De La Rosa, S. (2011). An investigation of grasp type and frequency in daily household and machine shop tasks. *proceedings of the 2011 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4169–4175.

Appendices

A Chapter 3: Non-linear scaling of resource allocation.

A.1 Differences between 1D and 2D regions

Here we demonstrate similarities and differences between 1D and 2D. The one-dimensional problems are easier to solve analytically and generally have less complex covariance patterning and subsequent eigenvalues. In 1D, there is a natural ordering of the eigenvalues (see section A.6, that does not need any approximation). Here we present the results for 1D, showing heterogenous receptor densities and activation alone, their interaction, and finally a comparison to 2D.

Heterogeneous density and activation alone in 1D

A key difference between 1D and 2D for density is that the density is considered over the whole region for 2D, which has two axes of tiled points. For this case, the relative density with another region means that in higher density regions, points are spaced further apart along each axis, with the number of points being \sqrt{d} to give $(\sqrt{d})^2 = d$, whereas, in 1D, the number of points on the single axis is scaled by d . Therefore, the points are less independent for the same correlation extent over the 2D sheet. In 1D the form of the convergence limits are the same for both the activation and density ratio only cases, $\frac{1}{1+\sqrt{a}}$ and $\frac{1}{1+\sqrt{d}}$ respectively.

Figure S.1A,D gives an example of the effects of heterogeneous density and statistics in 1D respectively. The key difference between 1D and 2D is the wider expansion of higher density and activation regions at narrow bottlenecks. However aside from this, the effects remain similar for varying the ratios (Fig. S.1B,C) and changes to σ (Fig. S.1C,F).

Combined density and activation for 1D vs 2D

The overall pattern for 1D and 2D for combined factors is similar; however, variations in spatial correlations extending in two axes over the sensory sheet causes differences in the initial allocations and convergences. Figure S.2 demonstrates allocations for a fixed density ratio of 1:5 for region 1 and region 2, respectively. The activation ratio is varied from being higher on region 1 to higher on region 2. For the higher density region, only those ratios which are the same or higher for this region lead to expansion over almost the whole bottleneck. Activation ratios lower than the density lead to contraction and expansion, and where the ratio is the opposite or higher, e.g. 1:5 density, but 5:1 or greater activation in favour of the opposite region, the representation is entirely contracted, see Fig. S.2A. This effect is demonstrated over a range of activation and density ratio combinations in Fig. S.2B. For 2D, the effect of the spatial correlations and density spacing over two axis leads to a narrowing of the possible combinations where both expansion and contraction occurs over the bottleneck (see lighter blue shading). Overall, the activation ratio determines more strongly whether the bottleneck regime will have expansion only, contraction only or both over the bottleneck.

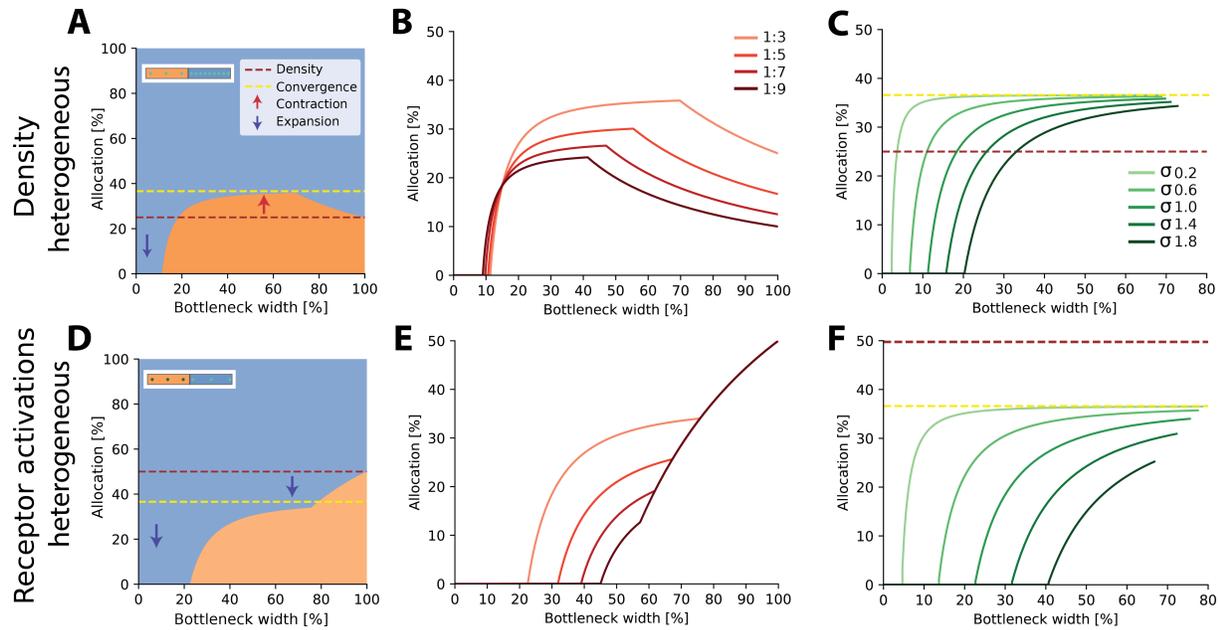


Figure S.1: Resource allocation for heterogeneous receptor densities and variations in input statistics in 1D. **A.** Examples of allocation for two input regions with differing receptor densities and activation (see insets) for different bottleneck widths, demonstrating complex trade-offs in resource allocation. Here ratio for low:high density = 1:3, $\sigma = 0.6$. Dotted lines show representation if there is no bottleneck effect (red, constant according to density) and convergence in the limit (yellow). Arrows indicate expansion (down) and contraction (up) of the higher density (blue) region. Both expansion and contraction of the higher density region are seen, similar to the 2D case. **B.** Allocation boundaries for different density ratios. Ratios are given as low:high density. **C** Effect of changing the decay σ parameter. The density ratio is set at 1:3 for all sigmas. Panels **D-F**, same as above but for density constant across the two regions and receptor activation manipulated. Activation ratio is set at 1:3 in panels **A,F**. Note that in panel **D**, the representation of the higher activation (blue) region is expanded for all bottlenecks.

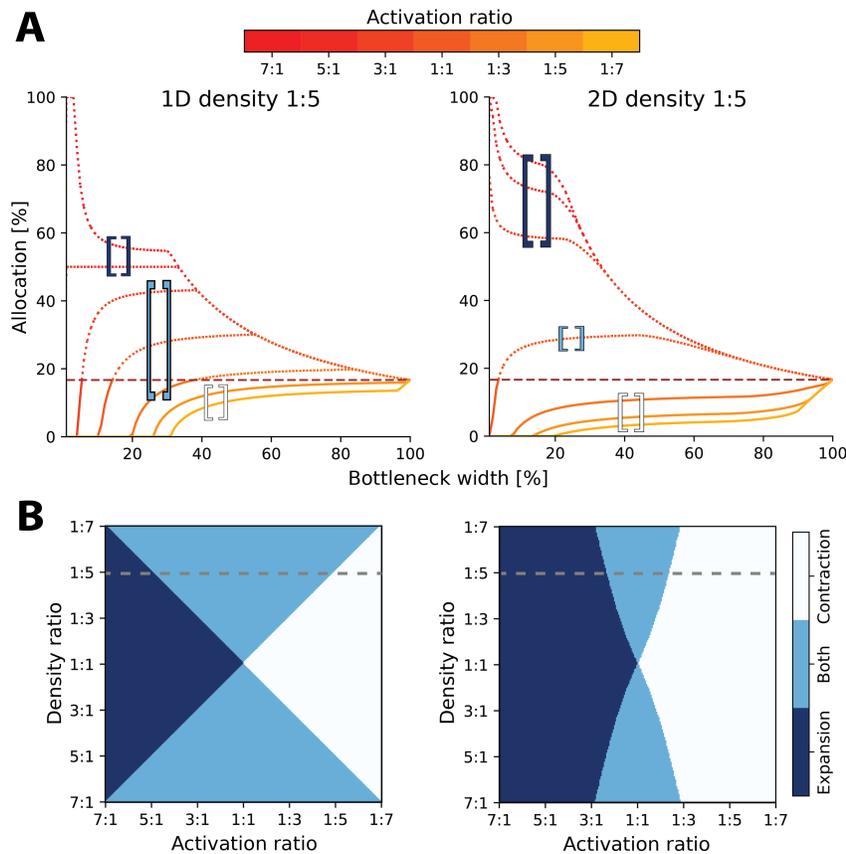


Figure S.2: Comparison between 1D and 2D results for heterogeneous statistics and density.

A. Effect of changing both the density and statistics and possible resource allocations for two regions. Plots show the same density ratio, 1:5, considering 1D (left) and 2D (right) regions. Allocation % is for the baseline region. Each line denotes the variance ratios between region 1:region 2 (see colourbar). Solid lines denote the contraction of region one's representation versus density, whilst dotted lines are expansion. The horizontal purple dashed line shows the density proportional representation. Across both 1D and 2D, different patterns of expansion and contraction occur. Region one is more likely to either be fully expanded or contracted for the whole bottleneck in 2D, whereas, in 1D, both expansion and contraction occur (denoted by the light blue square, which covers three activation ratios in 1D, but only 1:1 in 2D — see vertical box lines). In 1D, the same activation and density ratio lead to the convergence at the proportional density, whereas for 2D, the convergence is always lower.

B. Comparison of the possible allocation schemes between 1D (left) and 2D (right). The dotted grey line denotes the density ratio from panel A. In 2D, the representation's likelihood of being both expanded and contracted (blue region) over the bottleneck width decreases.

A.2 Expression as the total amount of information

Previously we described the bottleneck width in terms of the number of output neurons. Here instead, we demonstrate allocations in terms of the total amount of information (total variance captured). This leads to curves of similar shapes to those previously shown, with the initial expansion of narrower bottlenecks for both factors and their combination (Fig. S.3). In both the activation only (B) and combined cases (C), the higher region can be solely expanded for a much larger bottleneck width compared to the heterogeneous case (A).

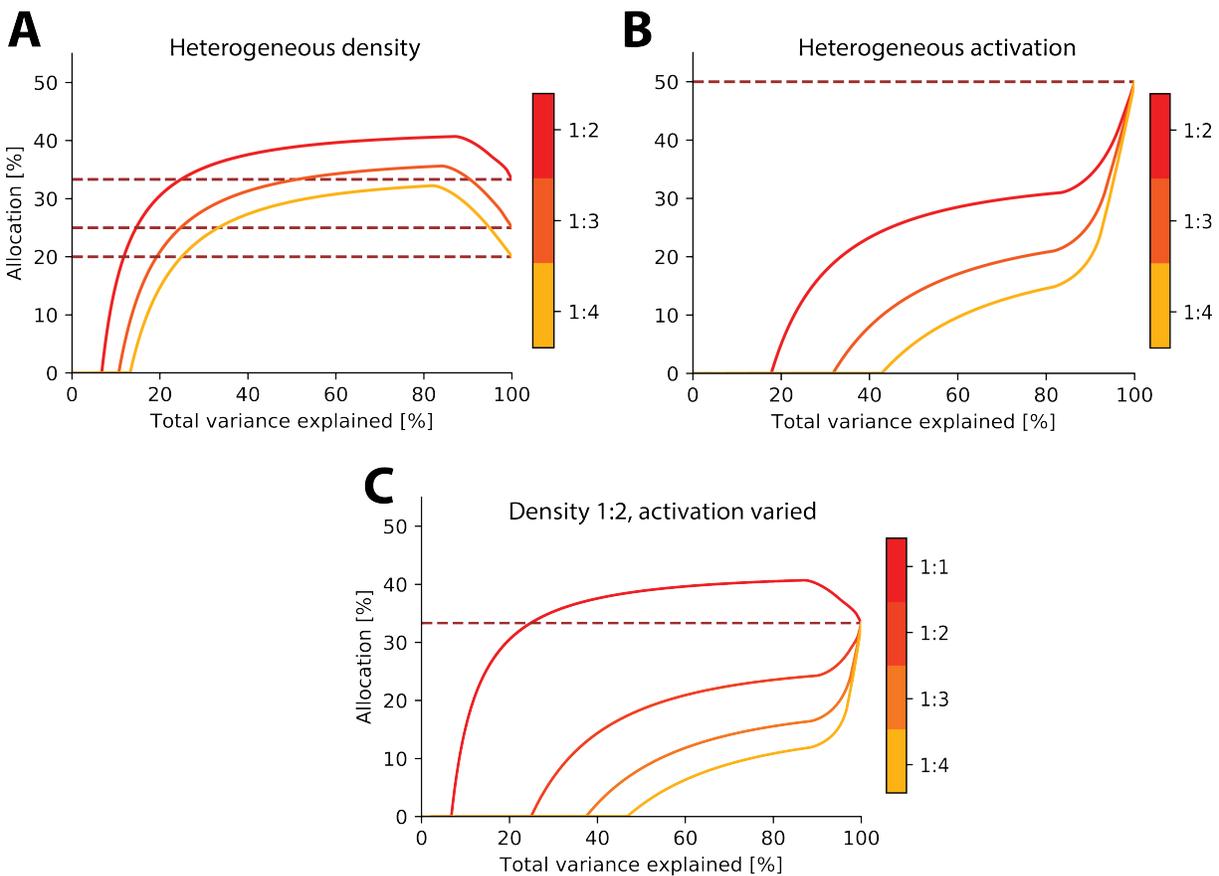


Figure S.3: Limit on information rather than the number of neurons. Allocations where the bottleneck is expressed as percentage variance explained. As eigenvalues decrease dramatically in size, this re-expression results in a 'squashing' of the allocation curve towards larger bottleneck sizes, as the initial eigenvalues in the sorted set explain a much larger amount of the variance. **A.** Examples with heterogeneous density. **B.** Examples with heterogeneous activation. **C.** Examples with heterogeneous density and activation.

A.3 Full derivation of problem

Problem set-up

Here we utilise the work of Jimenez Rodriguez (2018) who demonstrated the analytic solution for the eigenvalues of simple exponential covariance decays. We build on this by implementing the possible changes in both receptor density and region activation. From this we find an analytical solution to our bottleneck allocation problem.

Rationale

Let Ω be a one dimensional region, $\Omega = [0, L]$, with density of receptors denoted by ρ . The density can be written as $\rho = \frac{N}{L}$, denoting the number of receptors per unit length. We assume that correlation decay exponentially, such that the covariance between pairs of receptors i and j is:

$$C(i, j) = e^{-\sigma|i\Delta x - j\Delta x|}, \quad (\text{S.1})$$

where the distance between receptors is $\Delta x = 1/\rho$. Here, i and j are integer indexes, $1, 2, 3, \dots$. The space can be subsampled by taking a fraction N/d of the original receptors (where $d > 1$). The covariance for positions i, j becomes:

$$\tilde{C}(i, j) = e^{-\sigma|id\Delta x - jd\Delta x|} = e^{-d\sigma|i\Delta x - j\Delta x|}. \quad (\text{S.2})$$

The changes in receptor densities are therefore encoded by a scaling of the exponential decay rate. For a given distribution of receptors, there is an induced partition of the interval $[0, L]$. For a fixed $x = i\Delta x$, the covariance in the j th bin is approximately equal to the area of the exponential covered in that bin:

$$\int_{(j-1)\Delta x}^{j\Delta x} C(x, y) dy \approx C(i\Delta x, j\Delta x) \Delta x, \quad (\text{S.3})$$

By summing over all the bins, we arrive at the PCA problem:

$$\sum_{j=0}^n C(i\Delta x, j\Delta x) \phi(j\Delta x) \Delta x = \lambda \phi(i\Delta x). \quad (\text{S.4})$$

The continuum limit is found formally as $\Delta x \rightarrow 0$.

Derivation

To find the optimal assignment for a given receptor density, we are interested in solutions to the following equation which can be seen as a continuous version of the traditional PCA problem with an exponentially decaying covariance matrix:

$$\lambda_k \phi_k(x) = \int_0^L e^{-\sigma|x-y|} \phi_k(y) dy, \quad (\text{S.5})$$

where σ is the decay rate. We are interested in solutions $\phi \in C^2(\mathbb{R})$, that is, solutions that are twice differentiable and satisfy certain boundary conditions.

Theorem 1. If ϕ is a solution of (S.5), then it is an eigenfunction of the Laplacian operator with eigenvalues:

$$\mu_k = \frac{2\sigma}{\lambda_k} - \sigma^2. \quad (\text{S.6})$$

In one dimension, solutions to ϕ_k satisfy:

$$-\frac{d^2}{dx^2} \phi_k(x) = \mu_k \phi_k(x). \quad (\text{S.7})$$

Proof. Differentiating equation (S.5) twice using the Leibniz rule we obtain:

$$\begin{aligned} \frac{d}{dx} \phi_x &= \frac{\sigma}{\lambda} \left\{ - \int_0^x e^{-\sigma(x-y)} \phi_k(y) dy + \int_x^L e^{\sigma(x-y)} \phi_k(y) dy \right\}, \\ \frac{d^2}{dx^2} \phi_x &= \frac{\sigma}{\lambda} \left\{ -2\phi_k(x) + \sigma \int_0^L e^{-\sigma|x-y|} \phi_k(y) dy \right\}. \end{aligned}$$

The second term in the RHS can be replaced using (S.5) obtaining the desired result:

$$\frac{d^2}{dx^2} \phi_x = -\frac{2\sigma}{\lambda_k} \phi_k(x) + \sigma^2 \phi_k(x)$$

or

$$-\frac{d^2}{dx^2} \phi_x = \left(\frac{2\sigma}{\lambda_k} - \sigma^2 \right) \phi_k(x).$$

□

The previous is a sufficient conditions on the solutions to (S.5). A necessary condition is given in the following theorem:

Theorem 2. A solution to (S.7) is a also solution to (S.5) if it satisfies the following boundary

conditions:

$$\phi'(0) = \sigma\phi(0) \tag{S.8}$$

$$\phi'(L) = -\sigma\phi(L). \tag{S.9}$$

Proof. Assume ϕ is a solution to (S.7). We proceed by convolving (S.7) on both sides with the kernel $e^{-\sigma x}$:

$$\int_0^x e^{-\sigma(x-y)} \phi''(y) dy = \mu \int_0^x e^{-\sigma(x-y)} \phi(y) dy. \tag{S.10}$$

Integrating by parts twice we get:

$$\begin{aligned} \phi'(x) - e^{-\sigma(x-0)} \phi'(0) - \sigma\phi(x) + \sigma e^{-\sigma(x-0)} \phi(0) + \sigma^2 \int_0^x e^{-\sigma(x-y)} \phi(y) dy \\ = \mu \int_0^x e^{-\sigma(x-y)} \phi(y) dy. \end{aligned}$$

Using (S.6) and (S.9), we obtain:

$$-\phi'(x) + \sigma\phi(x) = \frac{2\sigma}{\lambda} \int_0^x e^{-\sigma(x-y)} \phi(y) dy. \tag{S.11}$$

Repeating the procedure with the kernel $e^{\sigma x}$ in the interval $[x, L]$, yields:

$$\phi'(x) + \sigma\phi(x) = \frac{2\sigma}{\lambda} \int_x^L e^{-\sigma(x-y)} \phi(y) dy. \tag{S.12}$$

Adding (S.11) and (S.12) we recover (S.5) which finalizes the proof. \square

Scaling

In some instances of our problem, the exponential covariance will be scaled by the activation ratio, a . In general, the same reasoning applies to any linear combination of solutions, therefore, our results extend to that case. In particular, we have the following result:

Theorem 3. The eigenvalues of the scaled covariance matrix, $C'(x, y) = aC(x, y)$ are:

$$\lambda_s = a\lambda, \tag{S.13}$$

where λ is an eigenvalue of the original problem.

Proof. Let $\phi(x)$ be a solution of (S.5). By linearity of the integral we have:

$$\int_0^L a e^{-\sigma|x-y|} \phi(y) dy = a \int_0^L e^{-\sigma|x-y|} \phi(y) dy = a \lambda \phi(x) \quad (\text{S.14})$$

□

Solutions

In the previous section we saw that solutions of the PCA problem (S.5) and Laplacian (S.7) coincide if boundary conditions (S.9) and (S.8) are met. In this section we show how these solutions relate to solutions of the boundary value problem of (S.7) with $\phi(0) = \phi(L) = 0$, which correspond to the eigenmodes of an idealized vibrating string fixed at the extremes. Such modes are considerably simpler than the exact ones and, as we show, are sufficient for our analysis.

Eigenmodes of the vibrating string

For $\kappa = \sqrt{\mu}$, $\mu > 0$, solutions can be found by assuming a general solution of the form:

$$\phi(x) = A \sin(\kappa x) + B \cos(\kappa x), \quad (\text{S.15})$$

and then using the boundary conditions to find the constants. The first boundary condition implies that $B = 0$. The second boundary condition gives the equation:

$$\sin(\kappa L) = 0$$

which is satisfied for $\kappa L = n\pi$ or

$$\mu = \frac{n^2 \pi^2}{L^2}. \quad (\text{S.16})$$

where $n = 1, 2, \dots$ is the index of eigenvalue.

Exact eigenvalues

In order to find the exact analytical eigenvalues of (S.5), we again assume $\mu > 0$ and a solution of the form (S.15). Using the first boundary condition (S.8, S.9) we get the following relationship:

$$A = \frac{\sigma}{\kappa} B, \quad (\text{S.17})$$

and with the second boundary condition, we obtain

$$\tan \kappa L = \frac{\kappa A + \sigma B}{\kappa B - \sigma A}; \quad (\text{S.18})$$

replacing (S.17), we find the transcendental equation

$$\tan \kappa L = \frac{2\sigma\kappa}{\kappa^2 - \sigma^2}, \quad (\text{S.19})$$

whose solutions lead to the exact eigenvalues of the continuous PCA problem.

Replacing (S.16) in (S.19) and reorganizing the right hand side, we find that the regimes for which our approximation is most accurate, are those for which

$$\left(\frac{2n\pi}{L\sigma}\right) \left(\frac{1}{\left(\frac{n\pi}{L\sigma}\right)^2 - 1}\right) \approx 0. \quad (\text{S.20})$$

We have made sure that this condition was met in the analyses presented. The criterion includes cases where the spatial extent of the correlations is relatively small compared to the size of the overall sensory sheet, which should be a realistic case in many biological scenarios. Finally, it should be pointed out that our procedure does not require the calculation of exact eigenvalues, but only relies on their relative ordering to be preserved and results should therefore be robust to small errors in the approximations.

A.4 Ordering in the 2D square case

For rectangle regions, the ordering can be solved by calculating the number of lattice points enclosed by a quarter ellipse (Strauss, 2007). Here we use square regions and therefore the solution is the number of points enclosed in a quarter circle.

The Gauss circle problem determines the number of integer lattice points which lie within a circle with radius $p \geq 0$, with its center at the origin:

$$N(p) = \#\{(l, m) \in \mathbb{R} \mid l^2 + m^2 \leq p^2\}. \quad (\text{S.21})$$

The number of lattice points within the circle is approximately equal to its area. The number of points within a square region can be approximated by calculating the area of the upper quarter of the circle (positive values only).

$$N(p) = \frac{\pi p^2}{4} \quad (\text{S.22})$$

The number of eigenvalues in each region is therefore the area of the intersection of the circle and region. This is demonstrated in Fig. S.4.

For each region we calculate the number of lattice points enclosed by a quarter circle with radius equal to $l^2 + m^2$ for R_1 , and the solution of equation (3.14)– $n^2 + o^2$ – for R_2 , where $l, m, n, o = 1, 2, \dots$. This number is approximately the area of the quarter circle. For values of $l^2 + m^2$ or $n^2 + o^2$ greater than the total number of eigenvalues in each dimension ($L\sqrt{d}$), the approximation diverges from the

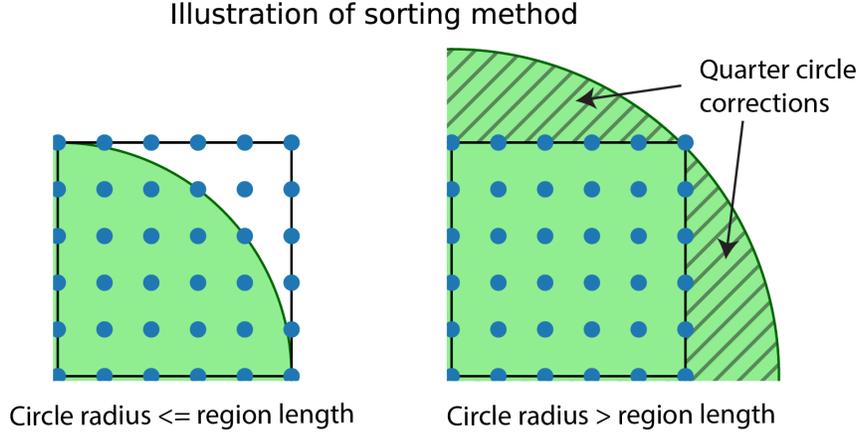


Figure S.4: Illustration of sorting method. The area of the grid enclosed by the circle is approximately equal to the number of allocated points. However, when the circle is large enough to encompass all points, there is an error as the circle area overflows the edges of the grid (right). In this case a correction term is added to the sorting calculation.

true ordering as the area of the quarter circle becomes larger than the area of the lattice (region). In this case a correction term is added:

$$N(p) = \begin{cases} \frac{\pi p}{4} - p \arccos\left(\frac{k}{\sqrt{p}}\right) - k\sqrt{p - k^2}, & \text{if } \frac{k}{\sqrt{p}} < 1. \\ \frac{\pi p}{4}, & \text{otherwise.} \end{cases} \quad (\text{S.23})$$

where p is either $l^2 + m^2$ or $n^2 + o^2$ for R_1 and R_2 respectively, k is the total number of eigenvalues in each region. Assuming region size L , where each receptor is spaced one unit apart, $k = L^2$ for R_1 , and $k = L^2 d$ for R_2 .

A.5 Allocation for multiple regions in 2D

For more than two regions, density and activation ratios for each additional region are calculated relative to a chosen baseline region. This leads to following general form for calculation of the eigenvalues of any Region x (R_x):

$$R_x: \lambda_{l,m} = \frac{2\sigma a \sqrt{d_x}}{l^2 \pi^2 L^{-2} \sqrt{d_b} + m^2 \pi^2 L^{-2} \sqrt{d_b} + \sigma^2 \sqrt{d_b}} \quad (\text{S.24})$$

where a is the region activation scaling ratio, d_b is the density of the baseline region, and d_x the density of region x . $l, m \in \mathbb{N}$ enumerate different eigenvalues for region x .

A.6 Allocation for the 1D case

The 1D case for changes in density has previously been addressed in Edmondson et al. (2019). Here we extend this to include changes in activation. For two regions R_1 and R_2 , we can calculate their eigenvalues as:

$$R_1 : \lambda_l^{(R1)} = \frac{2\sigma}{l^2\pi^2L^{-2} + \sigma^2} \quad (\text{S.25})$$

$$R_2 : \lambda_m^{(R2)} = \frac{2\sigma ad}{m^2\pi^2L^{-2} + \sigma^2} \quad (\text{S.26})$$

where d is the ratio of higher and lower densities, a is the ratio of receptor activation, L is the length of the region, and $l, m \in \mathbb{N}$ denote successive eigenvalues for regions R_1 and R_2 , respectively.

To calculate how many output neurons are allocated to region R_2 as a function of the number of neurons allocated to region R_1 , we set $\lambda_l^{(R1)} = \lambda_m^{(R2)}$ and solve for m . This yields

$$m = \frac{\sqrt{ad(l^2\pi^2 + L^2\sigma^2) - L^2\sigma^2}}{\pi}. \quad (\text{S.27})$$

It becomes apparent that for $l = 1$, i.e. the first neuron allocated to region R_1 , we have already assigned $m = \frac{\sqrt{ad(\pi^2 + L^2\sigma^2) - L^2\sigma^2}}{\pi}$ neurons to region R_2 . As we allocate more neurons to region R_1 , the ratio $\frac{m}{l}$ simplifies to: $\lim_{l \rightarrow \infty} \frac{m}{l} = \sqrt{ad}$. The fraction of neurons allocated to each region therefore depends on the size of the bottleneck and converges to $\frac{1}{1+\sqrt{ad}}$ and $\frac{\sqrt{ad}}{1+\sqrt{ad}}$ for R_1 and R_2 respectively.

B Chapter 4: Star-Nosed Mole

B.1 Star-nosed mole prey model

To calculate the sigmas for the covariance values of each of the rays, we created a prey model to simulate typical extents of contact across each ray, depending on stimuli sizes. Each ray does not contact stimuli of different sizes with the same frequency. Instead, smaller stimuli are more likely to contact the rays closest to the mouth, such as ray 11; more lateral and dorsal rays typically experience predominately larger stimuli. This is demonstrated in Fig. S.5.

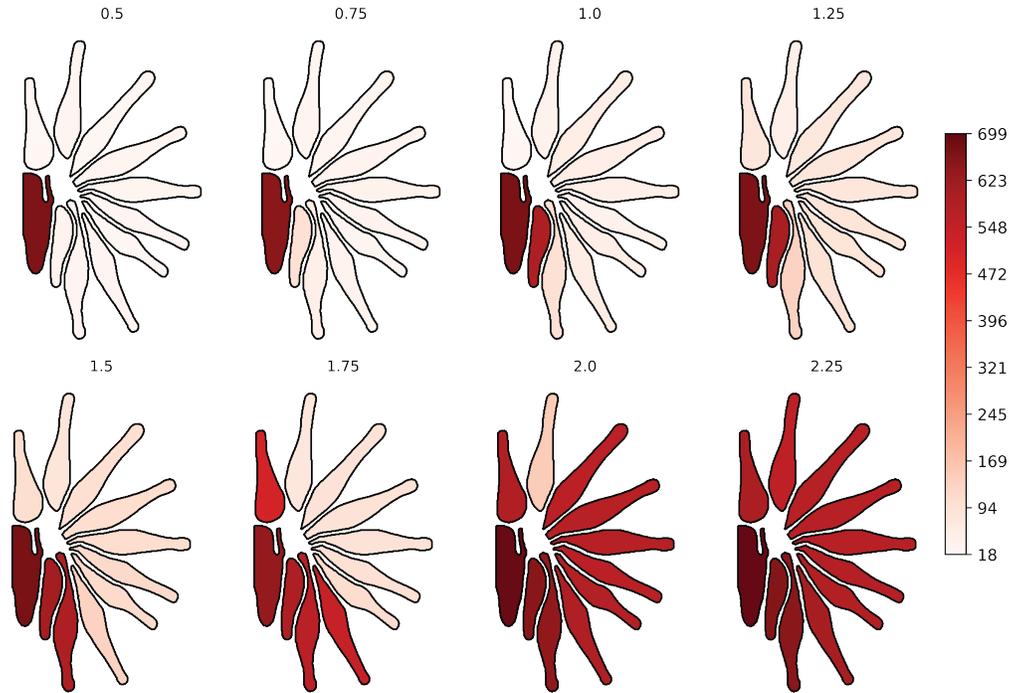


Figure S.5: Number of each model stimuli by size. Number of contacts on each ray by prey sizes. Ray 11 encounters many more prey than all other rays, particularly at small sized stimuli. Data taken from the prey model of Catania and Kaas (1997).

To simulate contacts, we used the 2D outline of flattened rays, shown in Fig. S.6A. Stimuli in their varying sizes and frequencies (according to Fig. S.5) were placed onto the ray. Stimuli were assumed to be circular in size as in previous work. As the rays are not flat when in use but a 3D shape, smaller stimuli areas are likely in contact with rays during interactions, such as edges of the stimuli. Therefore, we reduced the size of the stimuli, where the largest covers either most or all of ray. The relative sizes of the objects to the rays are shown in Fig. S.6A. Stimuli could be placed onto the ray, centred either within the ray or a boundary determined by the radius of the currently presented stimuli (see examples in Fig. S.6B). Rays were tiled with points as model receptors, set to '1' if covered by the stimulus and '0' otherwise. The density of these receptors does not need to correspond to the actual densities, as final computed covariances are based on distance. Model receptors were placed in the rays such that their density was much smaller than the size of

the smallest stimuli but large enough for speed of computation. From the responses, the sigma is estimated from a fitted exponential function to the response covariances between afferents based on their distances, Fig. S.6C. Given the distribution of stimuli sizes on each ray, fitted σ is similar for more distal and lateral rays and larger for rays closer to the mouth, where smaller stimuli are encountered.

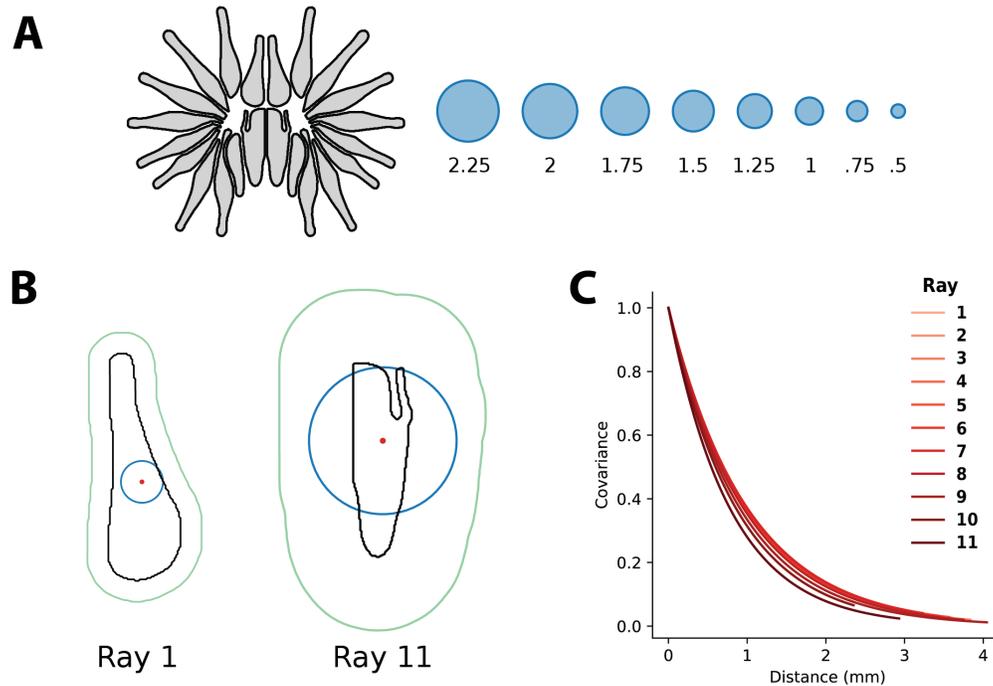


Figure S.6: A. Stimuli sizes. Differences between each star-nosed mole ray and the stimuli sizes. **B. Example model coverage.** Example coverage for two rays— one and eleven. Black outline shows the ray boundary. Each ray is tiled with uniformly spaced points—‘receptors’ (not shown). Blue circles are the stimuli outlines, and red dots their centroids. Any afferents covered by this blue stimuli are set to active; those outside are set to inactive. The green outline demonstrates the buffer where the stimuli centre can be placed. **C. Sigmas for each ray.** Exponential covariances for each ray. Darker colours denote higher-numbered rays. Ray 11 has the quickest drop-off and corresponding largest σ due to the smaller stimuli sizes typically encountered. Curves end at the maximum distance for that ray.

B.2 Comparison of bottleneck best fits

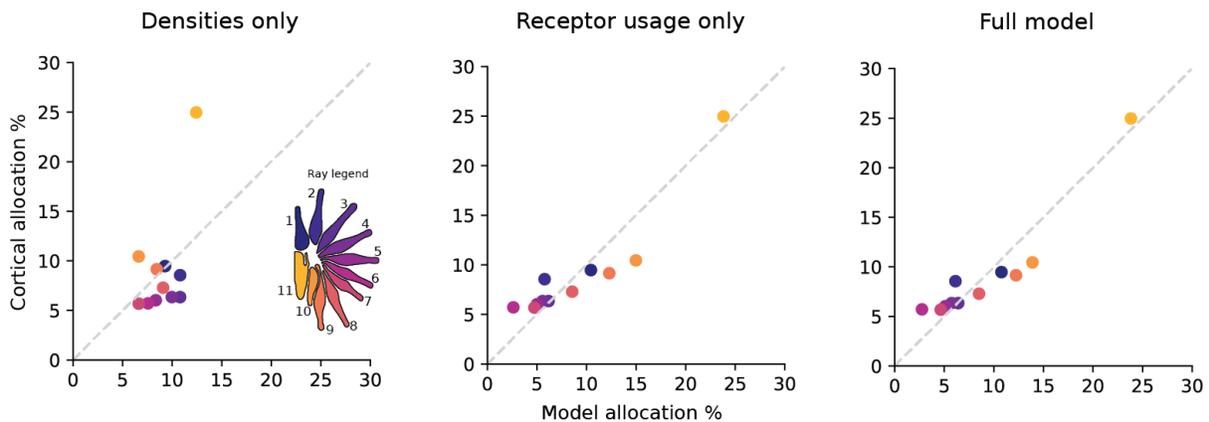


Figure S.7: Comparison fits at lowest RMSE. Comparisons between fits for ‘density’, ‘usage only’ models and ‘full’ model to the empirical cortical allocations for the lowest RMSE bottleneck. Colours correspond to ray legend. ‘Usage only’ and ‘full’ models provide the best fits versus cortical size.

B.3 Choice of error measure for calculating the bottleneck

The root mean square error (RMSE) was used in the fitting between the model and the cortical data. Ray 11 is considered the ‘fovea’ of the nasal rays (Catania and Kaas, 1997; Catania and Remple, 2004), and it is important to capture this accurately ray in the fitting. However, as ray 11’s cortical allocation is much larger than other rays, this could lead to an over-biasing of the fit towards this data point. Although it is unlikely that the use of a different error measure would change the main result (models considering usage have the best fits), the exact value of the best fitting bottleneck width could vary. As the RMSE can be sensitive to outliers, we tested this against two other error measures. As before, we calculate the bottleneck with the best fit between empirical cortical sizes and allocations from the ‘usage only’, ‘density only’ and ‘full’ models. First, we test the mean absolute error, which reduces the bias of larger data points. Second, we tested a ranked metric— Spearman’s rank, which does not consider the exact value of the ray fits, but the accuracy of their relative size order. The original RMSE fits over all bottleneck sizes, and values for the best fitting bottleneck are shown in Fig. S.6A and B, left panel, respectively. The RMSE was found to be similar for all bottlenecks in the ‘density only’ model, and between 38% and 46% for those that include accurate ray usages.

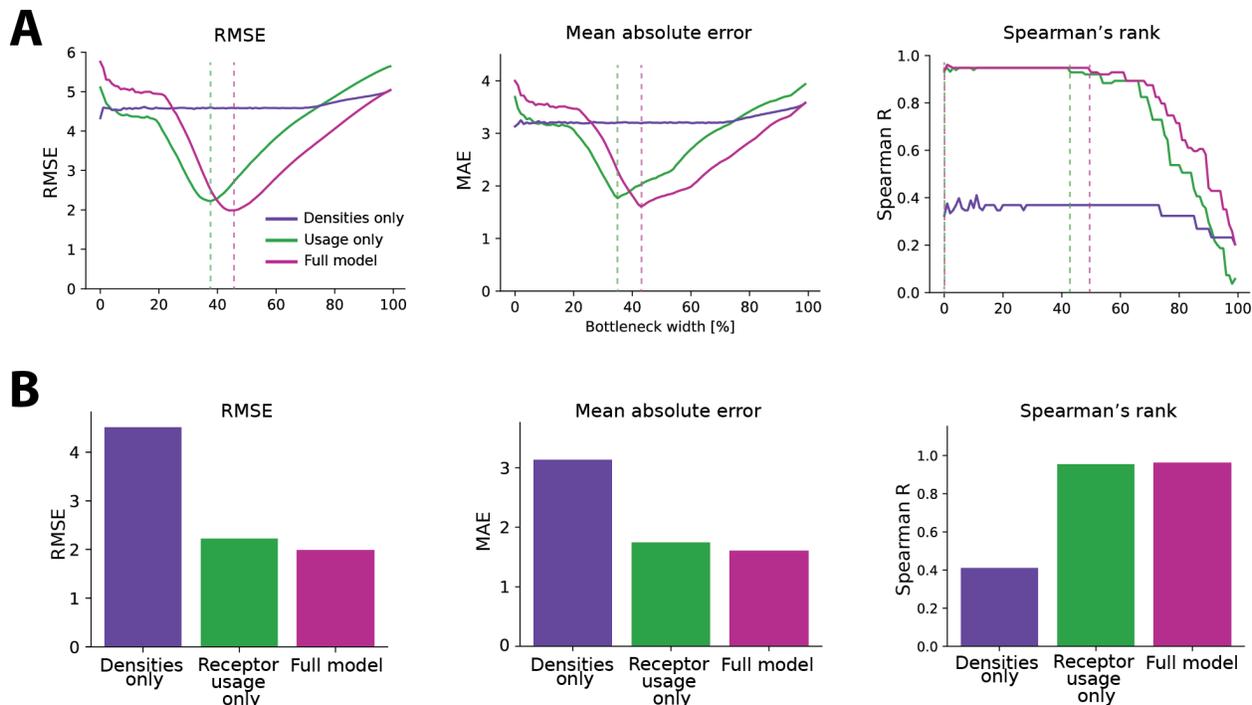


Figure S.8: Error measure comparison. **A.** Fits over all bottlenecks for three different error measures. Left panel shows RMSE used in main text, middle panel shows MAE, and right panel, Spearman's rank. Vertical dashed lines in the RMSE and MAE panels denote best fitting bottleneck for the 'usage only' model (green) and 'full' model (pink). For Spearman's rank, dashed lines denote where error increases. **B.** Values for each error measure at their relative best fitting bottlenecks. The main result that models with ray usage included have a better fit to the empirical data holds over all choices of error measure.

The mean absolute error (MAE) calculates the sum of the absolute difference between the model and cortical values. The error linearly scales, therefore data points that are off by a larger amount do not over-bias the error towards the fitting of that data point. The fits over all bottleneck widths are shown in Fig. S.6A, middle panel. Best fitting bottlenecks are similar to that of RMSE, at 36% for the 'usage only' model and 44% for the 'full' model, placing the bottleneck at an intermediate width.

Spearman's rank measures the ranked differences between the empirical cortex and model data. This metric does not compare the exact values, but whether their relative ordering (in terms of magnitude) is similar for both datasets. It is scored between 0 and 1, with 1 representing a perfect match between the ordering. The Spearman's rank values over all bottleneck sizes are shown in Fig. S.8A, right panel. The error is relatively consistent for all bottleneck widths up to 43% for the 'usage only' model and 50% for the 'full' model. As the width increases further, the error of the fit also increases. This suggests that the maximum width of the bottleneck is likely to be in the intermediate range, which is inclusive of the bottleneck size range calculated from the RMSE and MAE error measures. More detailed conclusions on the precise bottleneck fit below this value is not possible for this metric. At narrow bottlenecks, it is notable that ray 11 is more extremely

over-represented (see Fig. S.8A and B, main text), which is not captured by the Spearman's rank metric. However, a high score at these narrower bottlenecks suggests that relative ordering between ray sizes is consistently accurate.