Felt, Imagined, and Seen Touch Share a Substrate in Human Posterior Parietal Cortex

Thesis by
Srinivas Chivukula

In Partial Fulfillment of the Requirements for the degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2022
(Defended May 10, 2022)

Srinivas Chivukula
ORCID: 0000-0002-3570-162X

ACKNOWLEDGEMENTS

This work is a product of the time and effort that countless individuals invested in me over the years. I cannot adequately do justice to all in this brief section, but I would like to acknowledge a subset of those who have directly impacted my work and time at Caltech.

At the outset, I would like to thank my advisor, Richard Andersen for this truly remarkable and unique research experience. From the time I first approached him as a third-year resident in neurosurgery, interested in learning how to study brain function, he has been incredibly supportive and dedicated to my success, clearing countless obstacles for me to become a part of his lab. I would like to thank my co-advisor, Nader Pouratian. Over the many years that I have known him, he has been an unfailing mentor, guiding me, and opening doors for me, and creating opportunities for me. So much of who I am professionally, I owe to him, and I consider myself genuinely privileged to have people like Richard and Nader in my life. On the same level as Richard and Nader, I would like to thank my immediate mentor, Tyson Aflalo. It is hard to describe the role Tyson fulfilled for me – a friend, a colleague, a guide. If I have become a scientist, it is because of him. Tyson taught me how to listen to neurons, how to study them, not just one at a time but many at once. Tyson's ability to think critically in every situation, to think innovatively and inventively, to see the forest and not the trees, are just some of many aspects of his remarkably flexible intelligence that I have tried to imbibe, personally and professionally. Finally, I wish to thank the other members of my committee, Joel Burdick, and David Prober, for their time, their insightful feedback during all my meetings with them, their critical evaluation of my work and their thoughtful advice and guidance.

My heartfelt gratitude to all my many friends and colleagues within the Caltech family for productive discussions and for making this experience unforgettable: Spencer Kellis, Vasileios and Sofia Christopoulos, Luke Bashford, David Bjanes, HeyongChan Jo, Varun Wadia, Sarah Wandelt, Charles Guan, Kelly Kadlec, Isabelle Rosenthal, Sumner Norman, Kelsie Pejsa, and Viktor Shcherbatyuk. A special thanks to Matiar Jafari, a recent graduate from the Andersen lab and currently neurosurgery resident at UCLA. I remember when I walked into lab on my first day, in 2017, and sat hesitantly at the conference table, unsure how to get started. Matiar gave me a computer so I could pretend to be busy, and over the next year taught me so much of the basics.

I want to extend a heartfelt thanks to my UCLA family as well. Countless individuals worked to make this experience possible, and amongst them, foremost: Marvin Bergsneider and Linda Liau (my program director and department chair). They jumped through many administrative and logistical hoops to make it possible for me to dedicate time to focused research, leading to this work. I cannot also thank Ausaf Bari enough. He graciously obliged every time I requested time off to be in lab, or to work on research, even when I was on clinical service at the hospital. He is also an insightful and deep thinker and wonderful to have as a sounding board for ideas. I also want to thank Linda Demer, who heads the STAR

program, for always being an incredible resource, for her mediating between Caltech and UCLA, for her constant support and good cheer.

Most importantly of all, I would like to thank my family. I would not be where I am today if not for the steadfast love and support of my parents, Krishna, and Annapurna Chivukula. I have learnt, and have a platform to continue to learn, only because of their countless sacrifices, their unconditional commitment to me, their sweat, tears, and hard work. Both provided me with moral and emotional (and needless to mention, financial) support, endless encouragement, advice, and a comfortable home that have been the foundation and a support structure for anything that I have achieved so far, and I am forever grateful. I'd like to thank my brother, Ramanadh Chivukula; he was a voice of rationality. He was always available, uplifting me in moments of despair, and endlessly supportive, yet constantly challenging me to think more critically and consistently pressing me to perform at a higher level. My brother is a very quick thinker, and much like Tyson, able to bring seemingly isolated pieces of information into coherence and see meaning where many others see chaos. In trying to understand him, I find for myself benchmarks to work toward, higher standards of thought or insightfulness to work toward. I would like to thank my sister-in-law, Mythili Chivukula, for being so accommodating, caring, and loving. And my nephew, Ishan Chivukula, for always putting a smile on my face, a constant source of joy. Finally, during my time at Caltech, I had the fortune of adding to my family. I met my fiancé – Lalitha Voleti. She's truly a unique and wonderful person, an endless source of happiness – cheerful, lively, spirited, caring, and loving, while also insightful and sharp. She seems to float around spreading joy and warmth; I like to say, just like a happy bubble. She put up with me every time I refused to spend time with her – whether to go on a hike, or to the beach, or simply a walk with our dog – because I needed or wanted to work (on this thesis). She may deal with matters of the heart professionally, but she listens with interest every time I talk about neurons, and helped me find clarity in my own thinking about the work here. I'm very grateful to have found this special person, and her wonderful family, and excited to walk new roads together in the future.

ABSTRACT

One of the most remarkable aspects of human cognition is its flexibility. We can think new thoughts, infer meaning, plan actions, predict, extrapolate, and so much more. How do our brains enable this versatility? A growing ability to simultaneously record from large populations of single neurons in human cortex has begun to provide insight.[1] Recent studies have identified that shared populations of neurons in posterior parietal cortex (PPC) of a human subject (involved in a brain-machine interface (BMI) clinical trial) encode many aspects of motor cognition: ***attempted*** and ***imagined*** actions, ***observed*** actions and the ***semantic*** processing of action verbs.[2-4] Individual units are complex, but population representations manifest rich associations across neurons, supporting diverse behavioral contexts. Here, in novel work, we establish that the same PPC substrate also encodes aspects of sensory cognition, and unpack the functional organization of information that enables this versatility.[5] We record populations of neurons in PPC of the same human subject, a tetraplegic trial participant implanted with a 4x4 mm microelectrode array.[2-5] In a series of novel results, we first establish that neurons in this PPC substrate encode ***actual*** (or ***felt***) touch to oneself, at short latency, with bilateral receptive fields, organized by body-part.[5] We show that ***imagined*** touch to oneself and ***observed*** touch to others engage the same substrate.[5] To understand coding mechanisms further, we manipulated the touch location (cheek, shoulder), and the touch type (pinch, press, rub, tap). As in the motor domain, individual neurons exhibit highly variable responses. At the population-level, however, we find that the diverse touch conditions are explained by a small number of subspaces (meaningful groupings of neurons) that encode basic-level, elemental information such as touch location, and touch type. This suggests a *compositional* basis in PPC, such that various touch conditions are encoded through diverse combinations of common primitive elements. Moreover, these subspaces are *generalizable*, able to explain novel (held out) data. These principles of compositionality and generalizability suggest a basis by which PPC may support cognitive behaviors such as comprehension, in situations that extend beyond our experiences. In support of this interpretation, we show finally that this PPC substrate encodes ***seen*** touch universally – not only to insensate arm regions on the tetraplegic human subject, and to other human individuals, but also to a wide sampling of inanimate objects. As predicted, neural information combines and generalizes across conditions such that touch to objects with more similar features, is more similarly encoded. Taken together, our work is a novel, neuron-level characterization of how high-level cortex in humans may support diverse sensory, motor, and cognitive behaviors. We speculate that populations of neurons in PPC encode rich internal models of the world that can be flexibly repurposed for diverse (and novel) behavioral contexts.

PUBLISHED CONTENT AND CONTRIBUTIONS

Chivukula, S. et al. (2021). “Neural encoding of actual and imagined touch within human posterior parietal cortex”. eLife 10: e61646. doi: 10.7554/eLife.61646. url: https://elifesciences.org/articles/61646

S.C. participated in the conception, analysis, and interpretation of the data, as well as in the writing of the manuscript.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS AND/OR TABLES

Chapter 1

INTRODUCTION

From tying our shoelaces to solving complex equations or imagining life on Mars, every situation engages our brain's representations of our knowledge, yet in myriad ways. These representations are flexibly repurposed to meet diverse situational demands: to infer context, understand others' actions or words, think new thoughts, or plan and execute actions.[6] In every situation, knowledge is manipulated, information must be stitched together across domains into meaningful, coherent, context-relevant representations that inform our behavior. A fundamental question in neuroscience is: how do our brains represent knowledge in a manner that can support the depth and versatility in our repertoire?

This question has historically been difficult to address. On the one hand, human neuroimaging studies identify brain regions like posterior parietal cortex (PPC) that simultaneously support diverse sensory, motor, and cognitive behaviors, but lack the spatial resolution to clarify their neuron-level underpinnings.[7-11] On the other hand, neurophysiological studies have typically been conducted in animal models which offer high spatial resolution but limit the scope of testable behavior. For example, while it is possible to study how a monkey plans and ***executes*** movements, and how it responds to movements it ***sees*** performed by others, it is difficult to study how it may ***imagine*** the same movement (or to know if it is, in fact, ***imagining*** the movement).[12,13]

Moreover, until recently, electrophysiological studies relied on technologies that severely limited the number of neurons that could simultaneously be recorded and examined.[13] Recent technological advances offer the ability to simultaneously study large populations of single neurons, not only in animal models, but more recently, in humans.[1] This has begun to offer new insight into the computational basis by which our brains flexibly support diverse behaviors. Indeed, it has renewed support for a *neural population doctrine*, the idea that populations of neurons, and not individual neurons, are the fundamental units of computation in our brains.[1] While single neurons provide interesting snapshots of brain behavior, they do not adequately capture how populations of neurons collectively and coherently support behavior.

In recent population-level studies, shared neurons within a human PPC substrate have been found to encode many variables related to motor cognition: ***attempted*** and ***imagined*** actions, ***observed*** actions and the ***semantic*** processing of action verbs.[2-4] Neural information within this substrate is distributed across the population. Individual units are complex, but population representations manifest rich associations across neurons that form in a context-dependent manner, supporting diverse task conditions (i.e., behavioral contexts).[2] This organizational structure has been termed *partially mixed selectivity* and

allows relatively small number of highly variable and heterogenous neurons to represent many variables simultaneously.[4]

Here, we extend this work. We establish for the first time that this same human PPC substrate also encodes aspects of sensory cognition and unpack the functional organization of information that enables this versatility. To study ***actual*** touch to oneself and tactile cognition (***imagined*** touch and ***observed*** touch), we record populations of neurons within PPC through a microelectrode array implanted for an ongoing brain-machine interface (BMI) clinical trial.[5] This array (4x4 mm) is in the left hemispheric PPC, at the junction of the post-central sulcus (PCS) and intraparietal sulcus (IPS), a location we abbreviate PC-IP. All results presented in this thesis derive from single-unit recordings from this array.

In **Chapter 2**, we discuss theories for flexible human cognition that have been built from cognitive behavioral, and neuroscientific evidence. Specifically, we introduce the concepts of *internal models* and of *mirror neurons* and discuss how they have been proposed to support cognition.[6,13] We introduce the PPC, the brain region studied in this thesis, and discuss how prior PPC literature supports our hypothesis that it provides a rich, shared substrate within which to study how diverse forms of touch and tactile cognition engaged shared neural information. Lastly, we introduce the methods used in this thesis to investigate intended questions. In **Chapter 3,** we establish that populations of neurons in human PPC encode ***actual*** (or ***felt***) touch to oneself.[5] PPC neurons encode receptive fields to touch bilaterally, organized by body-part. We additionally show that tactile ***imagery*** shares this neural substrate, recruiting the same organizational structure. In **Chapter 4**, we show that neuronal populations from the same substrate encodes not only ***actual*** touch to oneself, but also ***observed*** touch to other individuals. We manipulated the touch location (cheek, shoulder), and the touch type (pinch, press, rub, tap). As in the motor domain, individual neurons exhibit highly variable responses. At the population-level, however, we find that the diverse touch conditions are explained by a small number of subspaces (meaningful groupings of neurons) that encode basic-level, elemental information such as touch location, and touch type. This suggests a *compositional* basis in PPC, such that various touch conditions are encoded through diverse combinations of common primitive elements. Moreover, these subspaces are *generalizable*, able to explain novel (held out) data. In **Chapter 5**, we show finally that this PPC substrate encodes ***seen*** touch universally – not only to insensate arm regions on the tetraplegic human subject, and to other human individuals, but also to a wide sampling of inanimate objects. As predicted, neural information combines and generalizes across conditions such that touch to objects with more similar features, is more similarly encoded. Taken together, our work is a novel, neuron-level characterization of how high-level cortex in humans may support diverse sensory, motor, and cognitive behaviors. We speculate that populations of neurons in PPC encode rich internal models of the world that can be flexibly repurposed for diverse (and novel) behavioral contexts.

Chapter 2

BACKGROUND

As humans, we have a rich understanding of the world that we flexibly implement in diverse situations. Even from just a still image of a person in action, for example, we can infer what he is doing, why is he doing it, imagine what might have just happened, or predict what he may do next. Such daily events reflect the remarkable versatility with which our brains can repurpose our knowledge stores, a core aspect of human intelligence. Yet, the way that neurons in our brains flexibly manipulate and repurpose our knowledge for different behavioral contexts and situational demands remains largely unknown.

2.1 A cognitive behavioral account – internal models

Much behavioral and observational evidence supports a view that at the heart of our cognition are flexible *internal models* of the world.[6] Internal models are our intuitive understanding of the world – our brains representations of the physical structure of the world and its dynamics, of how objects around us interact, of how they don't. This mental blueprint begins to develop early in life and becomes richer and faster as we grow, and as we learn. Infants, for example, can differentiate liquids from solids by five months of age, and by 11 months, can infer an object's weight from its compression of a soft material.[6,14] Much literature suggests that our foundational understanding develops not only regarding physics (objects and mechanics) but consistently in other domains as well including number (numerical and set operations), space (geometry and navigation), and psychology (agents and groups).[6] Each domain is organized by a set of entities and abstract principles relating the entities to each other.

By adulthood, we have a deep understanding of the world that can support complex cognition. We rely on this understanding in diverse real world-scenarios, such as to infer the meaning of a new word we hear or to plan a complex series of movements.[6] A simple example can illustrate the depth and versatility of these internal models. Read the following sentence: *imagine a white tiger standing on its hind paws, eating an ice cream cone with its right front paw, and holding a phone in its left front paw, and with a pink ribbon around its neck*. We may never have come across this type of tiger, but most of us can come up with a mental representation of such a tiger, flexibly manipulating our knowledge of a tiger, ice cream and a pink ribbon and combining them in a novel manner.

This example of imagining a tiger eating ice cream highlights at least two core features of how internal models are thought to support human cognition. First, internal models are *compositional*.[1,6] In this example, we have a mental representation of basic-level, primitive elements: tiger, ice cream cone, pink ribbon. We combine them as necessary to meet the task demands. Second, these elements *generalize* across contexts. Even if the only tigers

we have ever come across or heard of have been black and yellow, we can generalize that knowledge now to the new context of a white tiger. Such principles of compositionality and generalizability provide a way for internal models to enable novel constructions beyond our experiences, supporting complex cognitive behaviors such as prediction, extrapolation, imagination, and many other forms of cognition.[1,6]

Establishing the neural basis for a system of internal models has remained a challenge to date. Human neuroimaging studies show that multiple behaviors often activate shared cortex but cannot clarify the neuron-level mechanisms that underlie such shared cortical activation.[14] One brain region that may carry such a system is the posterior parietal cortex (PPC). The PPC is uniquely situated for such a role. It lies posterior to the primary somatosensory cortex (S1) and anterior to the visual cortex. Functionally, it has been described as an *association* region, integrating sensory information from multiple channels, and subserving coordination of movements of the body in space, specifically of the limbs, hands, and eyes.[15-19] It has been implicated in a variety of high-level brain functions including decision making, awareness, attention, and action planning.[20-23] Many of these functions have been studied in regions of PPC that lie along its intraparietal sulcus (IPS).

In support for a system of internal models in the human brain, magnetic resonance imaging (MRI) studies find that parietal and frontal regions are involved in intuitive physical inference.[14] PPC carries invariant representations of properties of external objects, like mass, underlying physical reasoning such as prediction and inference, consistent with an internal model of the physical world in PPC.[14] The PPC has also been implicated in thinking about physical concepts presented as words, consistent with a notion that it represents abstract, generalizable physical concepts. These studies suggest a critical role for PPC in implementing a system of internal models yet, these studies cannot elucidate the neuron-level underpinnings of this implementation.

2.2 Studies in neuroscience – mirror neurons

While the neuron-level basis for human cognition is unknown, *generalizability* in neuroscience has extensively been studied in the context of *mirror neurons.* Mirror neurons are a special class of neurons that were discovered in non-human primates (NHPs), and first reported in 1988.[13] They have typically been studied in the motor domain. They are neurons that respond to specific actions (e.g., grasping an object) both when the monkey ***executes*** the grasp as well as when the monkey ***observes*** the grasp being performed by another individual.[7,24,25] Different neurons were found that responded to grasps of different types: e.g., grasping-and-holding an apple versus grasping-and-eating an apple. Moreover, the neural response to both ***executed*** and ***observed*** action is matched, or congruent. The congruent activation of these neurons to specific actions was interpreted as a coding for the intentions of actions, for the purpose of action understanding. This became prominent as the *mirror hypothesis*: *we understand actions we see by simulating them within our own action production systems.*[7,24-26]

In the initial report, mirror neurons were identified within the frontal cortex in monkeys (specifically, in a region called F5).[13] Similar neurons have since been found in distributed cortical as well as subcortical regions.[24] Although examining human brain function at the neuron-level has historically been a challenge, similar neurons were reported in distributed brain regions in human subjects as well, recording from intracranial electrodes placed for mapping seizure foci in patients with epilepsy.[27] Moreover, in the years since their discovery, analogous neurons have been postulated to exist in other domains as well, such as for emotions (e.g., pain), or sensations (e.g., touch), together creating a general theory for how we understand the external world: the *mirror system.*[12,28]

This hypothesis, however, has significant shortcomings. Its basic units, mirror neurons, are single neurons characterized by specificity and congruency.[29] How might such neurons alone account for how we understand actions we cannot perform? Or sensations we cannot experience? We cannot fly, for instance, but we can certainly *understand* a bird flying. Or in the example above, even if we may never have seen a cow being milked, we can understand the activity occurring. We understand these activities even if they do not exist within our own action repertoire for our neurons to *simulate* ourselves performing these actions. Moreover, if mirror neurons respond to specific forms of actions (that are within our repertoire), a framework of mirror neurons alone cannot account for complex cognition.[29] How might they explain extrapolation? Or prediction? Or imagination?

Mirror neurons may certainly play a role in cognition, but they appear to be inadequate by themselves. As such, the mirror neuron theory is grounded in a special type of neurons, exemplifying the *single neuron doctrine*, or an idea that individual neurons are the fundamental units of computation in our brains.[1] The advent of newer technologies and the ability to simultaneously study large populations of single neurons in human cortex has begun to provide new insight. Growing evidence supports a notion that studying individual neurons cannot capture the complex ways that populations of neurons behave in, collectively.[2] Within this *population doctrine*, individual neurons provide intriguing samples but cannot capture the true mechanisms by which neural populations function.[1] The idea that ensembles of neurons may be more informative to studying cognition than individual neurons is not a new one, but it is seeing a new light given recent advances in recording technologies that allow sampling from hundreds, sometimes thousands, of neurons simultaneously.[1]

Indeed, many recent studies that have examined mirror neurons through recording from populations of neurons have found that they exist within a heterogeneous and diverse mix of neurons. In classical mirror type experimental paradigms in the motor domain, some code exclusively for ***executed*** actions, others for ***observed*** actions.[2-5] A small minority behave in the traditional mirror fashion, responding in a congruent manner to specific ***executed*** and ***observed*** actions. Moreover, recent evidence confirms that shared activation between ***executed*** and ***observed*** actions is not limited to congruent mirror neurons but instead involves subspaces of the population comprised of other types of neurons as well.[2-5] Such studies support a notion that examining cognitive behavior at the level of individual neurons

offers at best an impoverished understanding of the true mechanisms involved. Re-evaluating prior data (such as congruently activating mirror neurons to specific ***executed*** and ***observed*** actions) at the level of neural populations could offer remarkable new insight.

2.3 The posterior parietal cortex: a substrate for cognition

One brain region in which mirror neurons have been identified is PPC. PPC neurons in many regions around intraparietal sulcus (IPS) have been described to be multimodal: they integrate multisensory information (e.g., vision, proprioception) and transform these inputs into a representation useful for guiding our actions in the external world.[30-35] In NHPs, neurons in PPC regions respond congruently when the monkey ***observes*** specific actions, as well as when the monkey performs the corresponding action.[26,28] Even in the sensory domain, some PPC neurons have been identified that respond congruently not only when a monkey ***feels*** a touch but also when it ***sees*** touch to another monkey on the corresponding body part.[26,28] Such neurons have typically been interpreted as *mirror neurons* and taken as evidence of *mirroring*: using one's own body as a reference for perceiving actions by, or touch to, others.

The rich sensory responsiveness of PPC neurons makes it a good substrate for diverse human behaviors, including cognition. However, growing evidence raises suspicion that mirror neurons alone are adequate for this diverse encoding in PPC. For example, human neuroimaging studies indicate that not only ***seen*** touch to other individuals, but also to inanimate objects, activate overlapping regions of PPC.[36-39] It is difficult to explain how mirror neurons that activate similarly during ***felt*** touch and ***observed*** touch to similar body parts, might support touch to inanimate objects. Such studies, along with the previously discussed shortcomings of the mirror hypothesis, suggest that while mirror neurons do certainly exist, and may certainly play a role in diverse human behavior, broader mechanisms may be at play in substrates like PPC.[6]

Another line of evidence that PPC is a rich substrate for many cognitive variables, especially related to sensory-motor behavior comes from studies of its role in predictive (or forward) modeling.[18,21] This is a specialized type of internal model that has been studied in PPC. This model is concerned with utilizing a copy of the current motor command (known as efference copy), predicting upcoming states of movement, comparing it with all available, incoming sensory feedback, and enabling rapid, dynamic correction or updating of movement plans to increase the accuracy and dexterity of movements.[18,21] Such predictive modeling can compensate for delayed or noisy sensory information. It has been suggested that PPC is a prime location for such a predictive internal model because it receives (and associates) multiple sensory modalities.[18,21] Predictive or forward modeling within PPC during ***executed*** movements has been well studied in animal models. While such prediction itself is consistent with the broader idea of cognitive internal models, human neuroimaging evidence suggests that ***imagined*** movements also engage the same internal models, further supporting a view that these internal models (here, of the body) are flexible, and that PPC may provide a rich substrate for cognitive variables.[40]

It is likely that not individual neurons, but rather, populations of neurons in PPC, are critical to supporting these cognitive variables.[1] In direct support, many recent studies from the Andersen lab have identified that shared populations of neurons are mutually engaged during many aspects of motor cognition.[2-4,20,21,41] During an ongoing BMI clinical trial in the lab, a spinal cord injured (at cervical level three-four), tetraplegic human subject was implanted with a microelectrode array in the left hemispheric PPC. Neural populations recorded through this array were found to provide a shared substrate for not only ***executed*** movements, but also for ***attempted*** and ***imagined*** movements, ***observed*** actions and the ***semantic*** processing of action verbs.[2-4,20,21,41] In these studies, individual neurons are invariably complex, showing highly heterogeneous and variable encoding patterns. Yet, the population can support these various cognitive behaviors (or variables) in a seemingly systematic manner. For example, neural activity evoked by movements of a certain body part (e.g., shoulder) are very similar across behaviors (e.g., ***imagined***, or ***attempted*** movements).[2-4,20,21,41] In other words, the neural population appears to organize in a context-dependent manner, such that similar contexts (e.g., movement by a similar body part) are represented similarly across behaviors.

The encoding pattern within this population has been termed *partially mixed selectivity,* distinguishing it from *mixed selectivity.*[4,42,43] Traditionally, studies have looked at neural representations from a *pure electivity* perspective, where individual neurons are tuned to single task variables only, or a *sparse selectivity* perspective, where neurons are turned to specific combinations of task variables.[4,42,43] Mixed selectivity, on the other hand, is a framework where neurons in a population are tuned, often nonlinearly, to idiosyncratic combinations of variables.[4,42,43] In contrast, in partially mixed selectivity, neurons in a population are still tuned to idiosyncratic combinations of variables yet, across the population, organizes in a context-dependent manner allowing similar contexts across behaviors to evoke similar neural activity.[4,42,43]

The findings that populations of neurons in human PPC encode many aspects of motor cognition is consistent with the population doctrine, that populations of neurons (and not individual neurons alone) are at the heart of human behavior.[1] Moreover, such studies provide a way that mirror neurons might be involved yet allow for broader mechanisms to support cognition than a framework of mirror neurons alone. Such findings are also consistent with an idea that PPC may implement a framework of flexible, cognitive internal models.[18] If this is true, however, we can hypothesize that these models should support not only motor cognition but also sensory aspects. For example, this PPC population should be able to encode ***felt*** touch or ***seen*** touch. Moreover, it might encode ***seen*** touch more broadly, to inanimate objects as well. This thesis tests this hypothesis.

2.4 Methods for examining shared representations

Studying the brain basis for is challenged by the difficulty in recording neural activity at a high resolution in humans. Functional imaging techniques such as functional magnetic

resonance imaging (fMRI) cannot enable neuron-level precision.[44-46] Similarly, surface or intracortical electrode recordings often have limited neuron-level capabilities. For instance, electroencephalography (EEG) records the brain's electrical activity from the scalp, using an array of electrodes to measure the gross activity of a region.[47-49] Another technique is called electrocorticography (ECoG).[47-49] It is in invasive, intracortical recording directly from the brain. However, like EEG, it records aggregate voltage signals from the brain. The spatial resolution is much higher than with EEG, but nonetheless record local field potentials (LFPs) and are unable to resolve single unit neural activity.

One avenue for the exploration of such questions in humans is to record from high-fidelity, high-resolution, chronically implanted electrode arrays for research purposes (e.g., BMI studies).[48-50] BMI is a technology that allows for communication between the brain and an external device, such as a prosthetic limb or a computer. This may be unidirectional from brain to machine (or machine to brain), or bidirectional. BMIs measure brain activity either directly through neurons, or through a proxy or an aggregate measure (such as LFPs recorded through EEG).[48-50] However, in recent years, there have been an increasing number of human clinical trials in which chronic electrode arrays have been implanted directly on the surface of the brain.[3] These arrays enable the recording of activity from populations of single neurons within the brain.

Human BMI research has provided a unique opportunity to bridge between basic research and its clinical applications. For instance, several patients suffer strokes or other neurological conditions that leave them unable to communicate.[3,41] BMIs have demonstrated potential (albeit in research settings for the time being) in assisting these '*locked-in*' patients to communicate again, through machines.[3,41] Similarly, there are patients unable to move their limbs because of, for example, spinal cord injury. BMIs can help such paraplegic patients regain some lost functionality. Sometimes this is as simple as being able to drink from a bottle or from a cup. A proof of principle has been demonstrated through the control of a robotic limb, or through direct stimulation of their muscles to perform the movement.[3]

The Andersen lab has been involved in a BMI clinical trial as was described above. The work presented in this dissertation includes neuronal recordings from one human participant in this trial. The details of the trial are presented in later chapters, but briefly, the subject is a human tetraplegic subject that sustained a high cervical (between cervical levels three and four) injury. This spinal cord injury (SCI) is motor complete (American Spinal Injury Association [ASIA] score A) meaning that the participant has no sensation or motor function below approximately the level of her shoulders. This subject was implanted with a chronically implantable electrode array in her left PPC, approximately ten years prior to this report. Most recordings reported here were performed approximately 6 years after implantation. Although previous publications from the Andersen lab have occasionally labeled the site of the implant as the anterior intraparietal area (AIP), a region functionally defined in NHPs, we will refer to the implantation site as post-central-intraparietal (PC-IP) area, acknowledging that further work is necessary to characterize the functional

homologies between human and NHP anatomy. The implant is small (4mm x 4mm) with 96 recording channels (with 4 additional reference channels) and has offered a singular opportunity to record from human single units with high-fidelity.

.

Chapter 3

ACTUAL AND IMAGINED TOUCH WITHIN HUMAN POSTERIOR PARIETAL CORTEX

The following chapter's contents are taken and adapted from Chivukula et al. 2017, with modifications done to fit the dissertation format.

Chivukula, S. et al. (2021). "Neural encoding of actual and imagined touch within human posterior parietal cortex". eLife 10: e61646. doi: 10.7554/eLife.61646.

3.1 INTRODUCTION

Touch is a complex, multisensory perceptual process [33,51,52]. In non-human primates (NHPs), multisensory input (e.g., visual, tactile) converges upon neurons in higher-order brain regions such as the posterior parietal cortex (PPC) where they are integrated into coherent representations [30-34,53-56]. Recent human neuroimaging studies suggest that the PPC is also recruited during touch cognition in the absence of actual tactile input (e.g., seen touch, or imagined touch), supporting a notion that both higher-level touch processing and tactile cognition share a neural substrate [57,58]. To date, however, such a link has not been established at the single neuron level.

We recently reported an analogous relation in the parallel domain of motor function [2-4,21,59]. In these studies, we found that a shared PPC neuronal population coded for overt movements as well as cognitive motor variables including imagery, observed actions and action verbs [2-4,21,59]. This richness of representation is made possible through a *partially mixed* encoding in which single neurons represent multiple variables, allowing a relatively small neuronal population (recorded through a 4x4 mm implanted microelectrode array) to provide many movement related signals [4,43]. Here, we hypothesize that the same PPC neuronal population engaged by high-level motor cognition also encodes actual tactile sensations as well as tactile cognition within this partially mixed encoding structure.

The neural correlates of somatosensory perception are characterized by spatially structured receptive fields to touch that respond at short latency [38]. In NHPs, subregions of the PPC within and medial to the intraparietal sulcus (IPS) encode tactile receptive fields that respond to bilateral stimuli [30,33,53,55,56]. These are often large receptive fields, extending across multiple body parts [30-32,53-56,60]. In humans, functional magnetic resonance imaging (fMRI) studies support multisensory encoding of touch within and medial to the IPS in anterior portions of PPC [61-63]. Although these studies indicate that relatively small regions of PPC may encode touch to large portions of the body, the limited spatial resolution of fMRI precludes a characterization of tactile receptive fields. The inability to resolve single neurons in fMRI is especially problematic when attempting to understand the significance of the grossly overlapping representations of actual touch and cognitive representations of touch [57,58]. Spatial correspondence in fMRI cannot confirm whether representations share a

neuron-level substrate [64]. Taken together, it is unclear from the current literature whether individual neurons in human PPC discriminate touch to different segments of the body with spatially structured receptive fields, and, if so, whether cognitive processing of touch engages the same populations of cells.

In a unique opportunity, we investigated touch processing in a tetraplegic human subject, at the level of single neurons recorded from an electrode array implanted in the left PPC for an ongoing brain machine interface (BMI) clinical trial. We recorded single and multi-unit neural activity during the presentation of actual touch and during imagined touch to sensate dermatomes above the level of the participant's injury. We found that human PC-IP neurons encoded actual touch at short latency (~50 ms) with bilateral spatially structured receptive fields, covering all tested, sensate regions within the head, face, neck, and shoulders. The tactile imagery task evoked body part specific responses that shared a neural substrate with actual touch. Our results demonstrate that PPC neurons that discriminate touch are partially reactivated during a tactile imagery task in a body part specific manner. The latter represents a novel finding, thus far untestable in NHP models, and suggests PPC involvement in the cognitive processing of touch.

3.2 MATERIALS AND METHODS

3.2.1 Study participant

The study participant, NS, is a 60-year-old tetraplegic female with a motor complete spinal cord injury (SCI) at cervical level C3-4 that she sustained approximately ten years prior to this report. She has intact motor and sensory function to the level of her bilateral deltoids. NS was implanted with two 96-channel Neuroport Utah electrode arrays (Blackrock Microsystems model numbers 4382 and 4383) six years post-injury, for an ongoing BMI clinical study. Implants were made in the left hemisphere as human neuroimaging studies have typically reported stronger coding of intention related activity in left versus right PPC [65,66]. She consented to the surgical procedure as well as to the subsequent clinical studies after understanding their nature, objectives and potential risks. All procedures were approved by the California Institute of Technology (IRB #18-0401), University of California, Los Angeles (IRB #13-000576-AM-00027), and Casa Colina Hospital and Centers for Healthcare (IRB #00002372) Institutional Review Boards.

3.2.2 Implant methodology and physiological recordings

The electrode implant methodology in NS has been previously published [3,4,43]. One array was implanted at the junction of the left intraparietal sulcus with the left post-central sulcus in what we refer to as PC-IP. The other was implanted in the left superior parietal lobule (SPL). Implant locations were determined based on preoperative functional magnetic resonance imaging (fMRI). The participant performed imagined hand reaching and grasping movements during a functional MRI scan to localize limb and hand areas within this region. Following localization, a craniotomy was performed on August 26, 2014. The PC-IP

electrode array was implanted over the hand/limb region of the PPC within the dominant (left) hemisphere, at Talairach coordinates [-36 lateral, 48 posterior, 53 superior]. In the weeks following implantation, it was found that the SPL implant did not function. Although this electrode array was not explanted, only data recorded from the PC-IP implant were used in this study.

3.2.3 Experimental setup

All experimentation procedures were conducted at Casa Colina Hospital and Centers for Healthcare. Participant NS was seated in a motorized wheelchair in a well-lit room. Task procedures are presented in detail in the sections below. For most tasks, however, one experimenter stood directly behind the participant and was responsible for providing tactile stimuli to the participant. A 27-inch LCD monitor was positioned behind NS (visible to the experimenter but not to the subject) to cue the experimenter for the presentation of a stimulus. Cue presentation was controlled by the psychophysics toolbox (Brainard, 1997) for MATLAB (Mathworks) [67].

3.2.4 Data collection and unit selection

Data were collected over a period of approximately eight months in the fourth year after NS was implanted. Study sessions were conducted between two and three times per week, lasting approximately one hour each. Neural activity recorded from the array was amplified, digitized, and sampled at 30 kHz from the electrodes using a Neuroport neural signal processor (NSP). The system has received food and drug administration (FDA) clearance for less than thirty days of recording. We received an investigational device exemption (IDE) from the FDA (IDE #G120096, G120287) to extend the implant duration for the purposes of the BMI clinical study. Putative neuron action potentials were detected at threshold crossings of -3.5 times the root-mean-square of the high-pass filtered (250 Hz) full bandwidth signal. Each individual waveform was made of 48 samples (1.6 ms) with 10 samples prior to triggering and 38 samples after. Single and multiunit activity was sorted using Gaussian mixture modeling on the first 3 principal components of the detected waveforms. The details of our sorting algorithm have been previously published by our group [4]. To minimize noise and low-firing effects in our analyses, we used as a selection criterion for units, a mean firing rate greater than 0.5 Hz and a signal to noise ratio (SNR) greater than 0.5. We defined SNR for the waveform shapes as the difference between their mean peak amplitude and the baseline amplitude, divided by the variability in the baseline.

Each recording session was assumed to be independent in reporting the total number of units. However, it is likely that some fraction of units recorded on separate days are resamples of the same neuron, given that recordings for different days were made from the same array. Neural waveforms are largely a function of the geometry of the electrode with respect to the neuron, and not a unique signature that can be used to characterize a neuron. Thus, it is impossible to determine whether waveforms collected on separate days represent the same, or different neurons. However, the degree to which recordings change from one

day to the next can be taken as a general indicator of whether there may be some daily change in which neurons are recorded from. In other words, if the set of waveforms across the array are the exact same from one day to the next, it is likely that we are recording from the same neurons. Conversely, if the waveforms change substantially from one day to the next, it is likely that we are recording from, at least partially, distinct populations of neurons. We performed two analyses to quantify changes in neural recordings across days. In the first, more conservative of the two analyses, we compare the number of waveforms on each channel between days. If the number of waveforms on a single channel changes, then this is strong evidence that there has been some substantial change in the neural recordings. By this measure, an average of 29% ± 4.2% of channels change across days. In the second analysis, we used a permutation shuffle test to measure whether the recorded waveforms on the same channel were more similar than waveforms across different channels. By this assessment, 58% ± 8% of neurons were distinct from one day to another. These values indicate that there was some degree of neural turnover from one day to the next.

Well isolated single and multi-units were pooled across recording sessions. To ensure that such pooling did not bias the conclusions of this manuscript, we performed core analyses within this manuscript on single units alone, potential multi-units alone and all units together. The results of these analyses, shown as supplemental figures for key results, generally demonstrate that our results were robust to the pooling of all sorted units together. For the separation of spike sorted units into high quality single and multi-units, we used as a threshold the mean across all units of the base-10 logarithm of their cluster isolation distances, based on a previously described method [4,68]. Sorted units for which the cluster isolation distance was above this measure were considered single units, and those with a distance below this threshold were considered potential multi-units. Findings were robust to the exact choice of isolation distance threshold.

For measurements of neural latency to stimulus response (please refer to the task descriptions below for more information), a custom capacitive probe was used to record the exact time of tactile stimulation. This probe was built using a Raspberry Pi 2B and Adafruit Capacitive Touch Hat (Adafruit product ID 2340). The digital output (a binary output for touch or no touch) was transmitted through a BNC cable into the NSP at an analog signal sampling rate of 2 kHz.

3.2.5 Task procedures

We used several experimental paradigms to probe various features of actual and imagined touch representations in the PC-IP. In each paradigm, the participant was instructed to keep her eyes closed. The basic task structure comprised three phases. Each trial began with the presentation of a cue to the experimenter (or an auditory cue in the tactile imagery condition, see specific task description below), 1.5 seconds in duration, indicating the stimulus (for example, touch NS's left cheek). This was followed by a brief delay, 1 second in duration. Then written text appeared on the screen to signal the experimenter to present the instructed stimulus, for 3 seconds (in the tactile imagery paradigm, a beep indicated the "go" signal for

the participant). Exact time intervals varied depending on task. Trials were pseudorandomly interleaved; all conditions were necessarily required to be performed at least once before they were repeated. In tasks in which both left and right body sides (ipsilateral and contralateral to the implant, respectively) were tested, stimuli were delivered to one body side at a time.

3.2.5.1 Neural responsiveness to touch

This task variant explored neuronal responsiveness and selectivity to actual touch to body parts (receptive fields) with preserved somatosensory input (above the level of SCI). Body parts tested included the forehead, vertex of the head, left and right back of the head, left and right cheeks, left and right sides of the neck, and the dorsal surfaces of the left and right shoulders. As controls, the left and right hands (clinically insensate) and a null condition (no stimulus presentation) were also included. Actual touch stimuli were presented to each body part as finger rubs by the experimenter at approximately one per second. The experimenters' fingertips were used to provide touch stimuli over an approximately 3 cm x 5 cm patch of skin on each tested body part. Stimuli to the left and right body sides were delivered on separate trials to evaluate each side independently. To ensure that any neural activity observed arose from actual touch and not from observed touch or other stimuli, NS was instructed to close her eyes throughout the task. She additionally wore ear plugs to block auditory input. This task was performed on four separate days. On each day, ten trials per condition were conducted. In total, we recorded from 398 sorted units on four separate testing days.

3.2.5.2 Neural response latency

The purpose of this task was to determine the latency of neural response to actual touch for the left and right sides of the body. Tested regions included the left and right cheeks, the left and right shoulders, and as controls, the left and right hands (insensate). Actual touch stimuli were presented as in the task above. Instead of finger rubs, however, a capacitive touch probe was used to enable precise delineation of the actual time of contact (touch) before the onset of a neural response. This task was performed on eight separate days, with eight trials per condition in each run of the task. In total, we recorded from 838 sorted units.

3.2.5.3 Receptive field size

This task aimed to estimate the size of neuronal receptive fields to actual touch. Neural responses to nine equally spaced points were evaluated, two centimeters apart, along a straight line from NS's right cheek to her neck (**Figure 3.2-figure supplement 3**). Only the right side (contralateral to the implant) was tested in this task. The first of these nine points was on the malar eminence, and the ninth point was on the neck as shown. In addition to the nine points, a null condition (no stimulus presentation) was also included. Stimuli were presented via a paintbrush (three mm round tip) gently brushed against each location, at a frequency of one brush per second. Touch was restricted to a small region of approximately

0.5 cm (parallel to test sites) by 2.5 cm (perpendicular to the test sites) immediately above the marked points shown in **panel A of Figure 3.2-figure supplement 3**. The paintbrush was employed to deliver spatially localized sensations without accompanying skin distortion that could mechanically stimulate nearby sensory fields. Data were recorded on six separate days. On each day, ten trials of each condition were tested. In total, we recorded from 585 sorted units.

3.2.5.4 Engagement during tactile imagery

This task was intended to establish whether PC-IP neurons are engaged by tactile imagery, and whether neural patterns evoked by cognitive processing of imagined touch and actual touch share a common neural substrate (e.g., activate the same population of neurons in similar ways). In this variant, the participant was presented with either actual touch stimuli or instructed to imagine the sensation of being touched. NS was instructed to keep her eyes closed throughout. Actual stimuli were cued to the experimenter with written words that appeared on the monitor. Because the participant's eyes were closed, the participant did not receive any information about the body part that would be stimulated prior to experiencing the touch. The cue was followed by a one second delay and then at the sound of a beep (the "go" signal), rubs at 1Hz were presented with the capacitive touch probe to either the left or right cheeks, shoulders, or hands. During imagined touch trials, an auditory cue was presented to NS instructing her to imagine being touched on her right cheek, shoulder, or hand. The auditory cue consisted of a voice recording of the words "cheek", "hand", or "shoulder" with cue duration of approximately 0.5 seconds. After a one second delay, at the sound of the beep, NS imagined touch to the cued body part. We asked the participant to imagine the sensations as alternating 1Hz rubbing motions similar to what she actual during actual touch trials. A null condition (without actual or imagined touch), not preceded by an auditory cue was used to establish a baseline neural response. Data were recorded on eight separate days. Eight trials of each condition were performed on each testing day. In total, we recorded from 838 sorted units.

3.2.6 Quantification and statistical analysis

In the analysis of data from the various task paradigms used in this study, we utilized several statistical methods. Some were specific for certain tasks, but others were applicable to multiple sets of data from the different paradigms. For ease of reference, we have described all methods together in this section. Where necessary, we provide specific examples from tasks to help illustrate their use in our manuscript. Unless explicitly noted, all recorded units for a given task were used in the statistical analyses pertaining to that task.

3.2.6.1 Linear analysis (relevant for Figures 3.1 and 3.4, and for Figure 3.1-figure supplement 2, Figure 3.2-figure supplement 3 and Figure 3.4-figure supplement 1)

To determine whether a neuron was tuned (i.e., differentially modulated to touch locations), we fit a linear model that describes firing rate as a function of the neuron's response to each

touch location. Neuronal responses were summarized as the mean firing rate computed between 0.5 and 2.5 seconds after stimulus presentation onset. The starting time of 0.5 seconds was chosen to minimize the influence of variable experimenter delay in presenting the stimulus. The baseline response was summarized as the mean firing rate during the 1.5 second window before the stimulus presentation cue. The linear equation is written as:

$$FR = \sum_{c} \beta_c X_c + \beta_0$$

where FR is the firing rate, X_c is the vector indicator variable for touch location c, β_c is the estimated scalar weighting coefficient for touch location c, and β_0 is a constant offset term. X_c was constructed by assigning a value of 1 if the corresponding firing rate was collected when touch location c was being stimulated and with a 0 otherwise. All baseline samples were also assigned a 0, effectively pooling together baseline data independent of subsequent touch location. Here we used indicator variables as our predictors because our stimulus was applied in a binary manner, either touch was applied to a position on the skin or not. Note that in principle, the formalism of linear models allows multiple indicator variables to take on a value of 1 at the same time. In our experiment, this would amount to simultaneous touch of two or more body parts. However, in out experiments, simultaneous touch was not tested and thus only one indicator variable could take a value of 1 at a time. Neural responses to a particular body location were considered responsive if the t-statistic for the associated beta coefficient was significant (p<0.05, false discovery rate (FDR) corrected for multiple comparisons). A unit was considered tuned if the F-statistic comparing the beta coefficients was significant (p<0.05, false discovery rate (FDR) corrected for multiple comparisons).

The linear models for each task were computed using all test conditions within the task, except when comparing discriminative coding between the left and right body sides. For this analysis, the goal was to determine how informative information encoded for one body side was for the other. Each neuron was fit by two linear models, one for touch locations exclusive to sensate regions of the contralateral side (e.g., contralateral shoulder, neck, back, and cheek) and one for touch locations exclusive to sensate regions of the ipsilateral side (e.g., ipsilateral shoulder, neck, back, and cheek). More details regarding this analysis are in Methods: Tests for mirror symmetric neural coding of body locations: single unit analysis.

3.2.6.2 Population Correlation (relevant for Figures 3.2 and 3.4, and for Figure 3.2-figure supplement 4 and Figure 3.4-figure supplement 1)

We used correlation to compare the population neural representations of various tested conditions (stimulus presentations) against each other in a pairwise fashion. Correlation was chosen over alternative distance metrics (such as Mahalanobis or Euclidean distance) because it provides an intuitive metric of similarity that is robust to gross changes in baseline neural activity across the entire neural population. Alternative distance metrics were tested and gave comparable results (e.g., **Figure 3.2-figure supplement 5**).

To perform the population correlation analyses, we quantified the neural representations as a vector of firing rates, one vector for each condition (stimulus location) with each vector element summarizing the response of an individual unit. As before, neural activity was summarized as the mean firing rate during the stimulation phase window, defined as 0.5 to 2.5 seconds after the onset of stimulus presentation. Firing rate vectors were constructed by averaging the responses across 50-50 splits of trial repetitions. The mean responses across different splits were correlated within and across conditions (e.g., across stimulations of different sensory fields), then the splits were regenerated, and the correlation computed 250 times. Performing the splits 250 times was chosen based on an empirical analysis applied to preliminary data. For preliminary data, we performed the analysis with N splits, with N ranging from 5 to 200 in steps of 5. We found that the mean correlation across splits converged to a stable value by about 80 splits. We then roughly tripled that to ensure that the numerical sampling scheme would capture a stable value of our cross-validated correlation metric. The across condition correlations measured similarity between population responses for different sensory fields, answering the question - are the tactile sensations similar or dissimilar from the perspective of the recorded neural population? The within condition correlations assist in our interpretation of the across format correlations by allowing us to quantify the theoretical maxima of the similarity measure (e.g., if the within condition correlation is measured at 0.6, then an across condition of 0.6 suggests identical neural representations.)

To test whether the difference between any pair of conditions was statistically significant, we used a shuffle permutation test applied to the correlations computed over the 250 random splits. To illustrate, in **Figure 3.4E** we applied this analysis to test whether the correlation between actual and imagined cheek touch was significantly different from that of actual cheek touch and imagined shoulder touch. The true difference in the correlations was computed as the difference in the mean correlations between actual and imagined cheek touches (over the 250 splits) and the mean of the correlations between actual cheek touch and imagined shoulder touches. We then randomly shuffled the two distributions together (2000 times) and computed the difference in the mean correlations for each shuffle. The distribution of shuffled differences served as the null distribution, against which we compared the true difference to determine significance. As in the case above, our permutation shuffle test used 2000 shuffles to ensure that the numerical sampling scheme would capture a stable value of the percentile of our true difference as compared to the empirical null distribution.

3.2.6.3 Decode analysis (confusion matrix; relevant for Figures 3.1 and 3.4, and for Figure 3.1-figure supplement 3)

Classification was performed using linear discriminant analysis with the following parameter choices: one, only the mean firing rates differ for unit activity in response to each touch location (covariance of the normal distributions are the same for each condition); and two, firing rates for each unit are independent (covariance of the normal distribution is diagonal). These choices do not reflect assumptions about the behavior of neurons, but instead, were found to improve cross-validation prediction accuracy on preliminary data. In our

experiments, we acquired 10 repetitions per touch location, generally not enough data to robustly estimate the covariance matrix that describes the conditional dependence of the neural behavior on the stimulus. In choosing equal covariance, we are able to pool data across touch locations, achieving a more generalizable approximation of the neural response as verified by cross-validation.

The classifier took as input a matrix of firing rates for all sorted units. The analysis was not limited to significantly modulated units to avoid "peeking" effects [69]. Classification performance is reported as prediction accuracy of a stratified leave-one-out cross-validation analysis. The analysis was performed independently for each recording session and results were then averaged across days.

3.2.6.4 Tests for mirror symmetric neural coding of body locations: single unit analysis (relevant for Figure 3.2-figure supplement 6)

The purpose of this analysis was to assess whether neural responses to one body side were the same as neural responses to the alternate body side on a single unit basis. Heuristically, we used a cross-validation approach, similar in concept to the population correlation, to ask whether the neural responses to one body are similar to the other body side. The transition to single units required one major modification from the population approach: Instead of comparing the pattern of response across neurons (as in the population case), we compared the pattern of response across the set of lateralized body locations (shoulder, neck, cheek, and back). We first selected the set of neurons that demonstrated discriminative encoding to at least one of the body locations that was tested to ensure that there was a meaningful discriminative pattern across sites to form a basis of comparison. Then we used a cross-validation procedure to compare within and across body-side encoding. A schematic representation of how the two sides were compared is shown in panels B-F of **Figure 3.2-figure supplement 6**.

For each neuron, we created a linear model that explained firing rate as a function of the response to each touch location on the right side. The linear model was constructed using indicator variables as described above, however, the set of body locations was restricted to shoulder, neck, cheek, and back. In this way, the response of a neuron is quantified by the continuous set of beta values for the 4 locations. This model was then used to predict the responses for the same 4 locations on the left side. The ability to predict the responses was quantified as the $R^2_{\text{Right to Left}}$. This metric is hard to interpret on its own; a low $R^2_{\text{Right to Left}}$ could indicate that responses are very distinct between the right and left side or it could indicate that the neuron is not very discriminative (e.g., there is high trial-to-trial variability relative to the differences in response to the different touch locations). Therefor we also computed a cross-validated $R^2_{\text{Left to Left}}$ measure. This disambiguates the $R^2_{\text{Right to Left}}$ measure. If $R^2_{\text{Right to Left}}$ is low but $R^2_{\text{Left to Left}}$ is high then we know that the unit is discriminative, but that the patterns of response between the right and left side are distinct. To compare apples-to-apples both the $R^2_{\text{Left to Left}}$ and $R^2_{\text{Right to Left}}$ were computed using leave-one-out cross-validation. This is necessary to ensure that the two measures are computed based on the

same amount of training data. To directly compare these values, we plotted them against each other as a scatter plot. If the patterns of response are similar, this would lead to data points falling along the identity line. If the patterns are distinct, the data points should fall below the identity line.

3.2.6.5 Response latency (relevant for Figure 3.3)

We quantified the neural response latency to touch stimuli at the level of the neural population. Prior to the analysis, trials were aligned by touch onset as detected by the capacitive touch sensor (ground truth). Principal component analysis (PCA) was used to summarize the population-level temporal response of recorded neurons [70,71]. We constructed a matrix of neural data D that was (n) by (t * c * r) in size, with n being the number of neurons, t being the number of time points, c being the number of conditions, and r being the number of trial repetitions. For each neuron, activity was sampled every 2 ms and no additional smoothing was applied. 2 ms windows was chosen to allow high temporal resolution to precisely localize the timing of the neural response with respect to touch contact. We used t=201 time bins starting from -150 ms and stopping at 250 ms with respect to the time of touch sensor contact. Different ranges from time-of-contact (up to -500 ms before and 500 ms after probe contact) were tested and the basic average latency was robust to the exact window choice. c=2, including data for touch to the cheek and touch to the shoulder. r=10, as we acquired 10 repetitions per condition. Principal components were calculated based on the singular value decomposition algorithm.
The first principal component (1PC) was retained, and responses were averaged across conditions and repetitions. Single trial results were visually inspected, and basic temporal profiles were consistent across conditions and repetitions. This process was performed separately for data acquired for touch to the left side and right side of the body. The 1PC was then fit with a piece-wise linear function with two transition points. The choice of two transition points was set based on visual inspection of the data and allow for an initial baseline, a subsequent rise, and a final plateau. The time at which transitions occurred was not constrained, being purely a product of the fitting process. Latency was reported as the time the piece-wise linear fit crossed the 95^{th} percentile of the baseline data, as measured by the distribution of activity in the window between -150 and 0 ms. To compute bootstrapped quartile bounds of the latency estimates, the above process was repeated 1000 times while resampling with replacement from the 1PC single trial results. To verify that 1000 resamples was sufficient to estimate a stable estimate of the quartile range, we repeated the process with 1500 resamples and found that the quartile estimate changed less than 1%.

To determine whether the mean difference of latency estimates was significant between the right and the left side, we performed a permutation shuffle test. We used a rank test to compare the true difference in latency estimates against an empirical null distribution of differences in latency estimates generated by shuffling labels and repeating the comparison 2000 times.

3.2.6.6 Quantifying macroscale receptive field structure (relevant for Figure 3.2-figure supplement 2)

We found that many neurons responded to touch to multiple body locations. We wished to further characterize the receptive field structure to determine whether neurons were characterized by single-peaked broad receptive fields or discontinuous receptive fields with multiple peaks. To adjudicate between these possibilities, we selected touch locations to the contralateral (right) forehead, cheek, neck, and shoulder for further analysis because these locations are approximately collinear. We reasoned that if neurons are characterized by single-peak type responses, then responses across a collinear set of testing sites will result in a single local maximum (either with a single peak and fall off on either side, or as a monotonic increase to the edge locations). On the other hand, if receptive fields are characterized by multiple peaks, then responses should have multiple local maxima.
Neurons were first restricted to those demonstrating significant differential responses between the four sites (ANOVA, $p<0.05$, FDR corrected). Each neuron was then grouped according to its location of preferred (peak) response. This resulted in four groups of neurons: neurons that responded maximally to the forehead, the cheek, the neck or the shoulder. For each neuron, the goal was to identify if the firing rate monotonically decreases with increasing distance from the preferred location or rises again, allowing for a second maxima. For example, for a unit preferring the forehead, this would manifest as firing rate at forehead larger than at cheek, at cheek larger than at neck, and at neck larger than at shoulder. Tests of firing rate between adjacent locations were performed by one-tailed t-tests between the pair of locations, evaluating whether the firing rate at the location nearer the preferred response was greater than at the location more distant. In the example of the forehead preferring units, the t-tests evaluated whether cheek>neck and neck>shoulder. If it was found that any of those comparisons was not true (e.g., firing rate at neck greater than at cheek) after correcting for multiple comparisons, this implied a second local maxima. The unit was then classified as multi-peak. If no second local maxima was found the unit was classified as single-peak.

3.2.6.7 Receptive field size estimation (relevant for Figure 3.2-figure supplement 3)

In our first experiment, we tested touch responses across major body parts at a course resolution. Patterns of neuronal responses suggest that multiple body parts can be represented in individual neurons, although the response field around each body part is locally narrow (not expansive, covering all body parts). To evaluate this further, as a complimentary dataset, we tested tactile representations at a finer spatial precision to begin to characterize the size of their receptive fields. We characterized the response patterns of individual neurons to tactile stimuli delivered to each of nine points along the subject's face and neck. All units demonstrating a differential spatial response to touch to each of the nine fields were included in this analysis. For each of these units, we first identified the preferred site of stimulus delivery as the point associated with the largest firing rate. Next, we examined its response to delivering stimuli to the other points. To estimate the average size of a neuronal receptive field as a function of its preferred point of stimulus delivery, we fit a

Gaussian model to the average responses grouped by the preference of the neuron. The Gaussian model had three free parameters, and was defined as:

$$G(x) = Ae^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} + c$$

Here, A is the amplitude of the Gaussian, σ is the standard deviation, and c is the constant offset term. μ is the mean/center of the Gaussian and was fixed at the preferred point. A separate model (with the appropriate value of μ) was fit to each of the response groups. The field size was described as the full width at half maximum (FWHM).

3.2.6.8 Temporal dynamics of population activity during tactile imagery task: within category (relevant for Figure 3.5)

We performed a sliding-window classification analysis to quantify the strength and temporal dynamics of population coding in the tactile imagery task. In this task, the participant heard an auditory cue specifying a body part (“cheek”, “hand”, or “shoulder”) that lasted approximately 0.5 seconds, followed by an approximately two second delay, and finally a beep instructing the participant to initiate imagining a touch sensation at the cued body part. This task could engage at least four cognitive processes: 1) semantic processing of the cue; 2) preparation/anticipation for imagery; 3) attentional modulation; 4) imagined touch per se. We used a dynamic classification analysis to understand how the neural population evolved during the course of the trial to determine whether the population was best described as mediating a single cognitive processes or multiple cognitive processes. In brief, the analysis consisted of creating a dataset that consisted of the population response measured in small temporal windows throughout the course of the trial. We trained a classifier separately on each temporal window and applied each classifier to both temporal windows. In this way we can measure how information about the cued stimulus evolves in time (e.g., does there exist neural coding during the delay portion of the trial, and, if so, does the neural coding during the delay match neural coding during active imagery). Classification was performed using linear discriminant analysis as described above. We used cross-validation to ensure that training and predicting on the same time window was directly comparable to training on one window and testing on an alternate time window; in other words, we were careful to ensure that accuracy across all comparisons reflects generalization accuracy using the same amount of training and test data. Classifiers were trained and tested on neural responses to the three imagery conditions: cheek, hand, and shoulder. Population response activity for each time window was computed as the average neural response within a 500 ms window, stepped at 100 ms intervals. Window onsets started at -700ms seconds relative to auditory cue onset (cue-delay epoch) with the final window chosen 3.5 seconds after the beep (onset of the imagery epoch). Classification was performed on all sorted units acquired within a single session. Mean and bootstrapped 95% confidence intervals were computed for each time bin from the cross-validated accuracy values computed across sessions.

We used a fixed window size for averaging time series data for analysis (box-car smoothing) as it provides straight-forward bounds for the temporal range of data that are included in the

analysis for a particular time window. 500 ms was chosen as a good balance between temporal resolution and noise mitigation. We note that although the window size can influence various metrics (e.g., larger smoothing windows can increase coefficients of determination, R^2) the choice of smoothing size is largely inconsequential as long as the kernel size is kept consistent when making comparisons across conditions. The choice of a 100 ms step size was anchored to the choice of smoothing window. A small step, such as 1 ms would not be justified with a 500 ms time window. We chose 100 ms, representing a change in 20% of the data, to allow us the ability to temporally localize changes in neural response without unnecessarily oversampling a smoothed signal and thus not unnecessarily increasing computation time for analysis.

We believe that this technique, by helping us to understand when information appears and how information compares across task phases, provides a valuable approach to understanding how population activity relates to the underlying cognitive processes. For example, if neural decoding reaches significance only after the go cue, neural activity would be inconsistent with semantic or anticipatory processing. Alternatively, if neural processing begins with the cue, and the same pattern of neural activity is maintained throughout the trial, with no changes during the active imagery phase, then the data would be inconsistent with processing imagined touch per se.

The classification analysis described above was used to measure general similarity of the population response to the tested conditions across time. However, to explicitly test whether population activity was changing, we used Mahalanobis distance as our measure. This is necessary as classification involves a discretization step that makes the technique relatively insensitive to changes in neural population activity that do not cross decision thresholds. Mahalanobis distance, being a proper distance measure, is a more sensitive measure of change. To illustrate, imagine that a classifier is trained on time point A and tested on time point B. At time point A, the means of the two classes are 0 and 1 respectively and at time point 2 the means are 0 and 4 respectively. All classes are assumed to have equal but negligible variance (e.g., 0.01) in this example. When trained on time point A, the classifier finds a decision boundary at 0.5. with 100% classification accuracy. When tested on time point B, with the same 0.5 decision boundary, the classifier again is 100%. Naively, this could be interpreted as signifying that no change in the underlying data has occurred, even though the mean of the second distribution has shifted.

Separation in neural activity between the cue-delay epoch and the imagery epoch was quantified using a cross-validated Mahalanobis distance computed between the observed neural activity at a time point and a reference (baseline) defined as the neural activity immediately following the presentation of the auditory cue, from .25 to .75 seconds. Distances were measured separately for each of the three conditions and then averaged. The mean and standard error on the mean (SEM) were computed across sessions for the cross-validated distance measures and plotted in **Figure *3*.5C**. Activity during the cue-delay epoch and the go epoch were compared using a rank-sum test of the averaged activity during the phase averaged responses across sessions.

3.2.6.9 Temporal dynamics of single unit activity during tactile imagery task: within category (relevant for Figure 3.5)

We wished to understand the behavior of single neurons that led to the temporal dynamics of the population. The temporal dynamics of single unit activity during the imagery task (for the imagined touch conditions only) were quantified a principal component analysis (PCA, **Figure 3.5D**). A sliding-window classification analysis was first performed on each sorted unit from all testing days in the same manner as described above for the population activity, with the exception that classifier took as input a vector of the firing rates for a single unit as opposed to a matrix of the firing rates for all units recorded in a single session. This allowed a quantitative description of the temporal dynamics for each sorted unit. Next, a principal component analysis (PCA) was applied to the dynamic classification matrices with individual neurons counting as the independent observations. PCA has become a standard method for describing the behavior of neural populations [70]. Typically, PCA is applied to firing rate measurements of neurons. However, in our case, we were less interested in capturing the main modes of variability with respect to individual conditions, but instead wanted to capture the main modes of variability with respect to the temporal dynamics of information encoding.

3.2.6.10 Temporal dynamics of population activity during tactile imagery task: across category (relevant for Figure 3.6)

We wished to evaluate the similarity in neural representations of actual and imagined touch in a time resolved manner, as well as to compare the similarities in activity from one epoch (cue-delay) to another (stimulus: actual or imagined touch). We performed a sliding-window (dynamic) correlation analysis in a cross-validated manner to compute within-format correlation in addition to across-format correlations. For this analysis, we restricted the tested body sites in the actual touch format to the right cheek and right shoulder only. Neural activity from the left side was not used, to try and match the conditions for the imagined touch format, in which only touch to the right side was tested. Similarly, the hand was not included in this analysis to match conditions that evoked responses in both formats.

For cross-validation, the analysis began with splitting trial repetitions into training and testing sets (5 trial repetitions each). A sliding time-window was used for the analysis with window size of 500ms and step size of 100ms. Correlations were computed between training and testing sets for all combinations of windows, starting from 500 ms before the cue-onset to 1000 ms after the end of the stimulus phase. Within each window, we organized the neural response data into two matrices (one each for the training and the test trial splits) with two columns each. Each column contained trial averaged firing rates during the corresponding time window for each of the two tested stimulation sites (cheek and shoulder), with one value per unit. The columns represented the two formats. Thus, for N units recorded, two tested stimulation sites and two formats (actual and imagined), each matrix was of size (2*N) x 2. The mean response across each matrix (computed separately for training and test sets) was subtracted from each value to ensure that a positive correlation across formats reflected

a similarity in the pattern of responses to the two body sites and not general offsets in the mean response of the different neurons. Finally, correlations were computed between training and test sets for all combinations of time windows. This was done across 50 random 50-50 trial splits and the results averaged across these repetitions. The analysis was performed for each recording session independently and the depicted results averaged over days.

3.3 RESULTS

We recorded from a total of 101.6 ± 7.2 neurons (**Figure 3.1-figure supplement 1**) over 14 sessions in the PPC (left-hemisphere) of a tetraplegic human participant (spinal injury at level three to four; C3/4). In previous work, we referred to the implant area as the anterior intraparietal cortex, a region functionally defined in NHPs [2-4,20,43,59,72]. Here we refer to the recording site as the postcentral-intraparietal area (PC-IP), acknowledging that further work is necessary to definitively characterize homologies between human and NHP anatomy. Recordings were split across four tasks, designed to probe basic properties of the neuronal population during both actual and imagined touch. Recordings were made from chronic implanted arrays and thus neuronal waveform sorting resulted in both well-isolated neuronal waveforms and multi-neuron groupings. The main figures aggregate across sorted channels while key analyses are performed separately for well-isolated and multi-unit activity in supplemental figures.

3.3.1 PC-IP neurons encode bilateral tactile receptive fields

We first examined the hypothesis that PC-IP neurons encode tactile receptive fields to dermatomes above the level of the participant's spinal cord injury (SCI). Tactile stimuli were delivered as rubbing motions at approximately 1Hz, for 3 seconds. The subject was asked to keep her eyes closed to eliminate neural responses arising from visual input. Tactile stimuli were presented to bilateral axial (forehead, vertex, cheek, neck, back) and truncal (shoulder) body parts to determine the extent of body coverage of any tactile representations among PC-IP neurons. As controls, touch was also presented to the bilateral hands (insensate regions below the level of SCI), and a null condition was included (with no stimulus delivered), to verify that touch related neural responses did not arise by chance.

For each neuron, we fit a linear model that explained firing rate as a function of responses to each touch location. Neural responses to a particular body location were considered significant if the t-statistic for the associated beta coefficient was significant (p<0.05, false discovery rate (FDR) corrected for multiple comparisons). A significant fraction of the neuronal population encoded touch to each of the tested body parts with preserved somatosensation ($\chi^2(1)$=3908.98, p<0.05; **Figure 3.1A, Figure 3.1-figure supplement 2**). These results are consistent with bilateral encoding as the tested body parts included both body sides. Neither touch to the hands nor the null condition elicited significant neuronal modulation. Single neurons discriminated the location of actual touch. Of the 263 responsive units shown in **Figure 3.1A**, we found that 257 discriminated touch location (analysis of

variance (ANOVA), false discovery rate (FDR) corrected for multiple comparisons). Representative examples of neurons showing clear discrimination between the different touch locations are shown in **Figure 3.1B**. As expected, a population of discriminative cells enabled accurate cross-validated classification of the touched body part (**Figure 3.1C;** see **Figure 3.1-figure supplement 3** for single session examples).

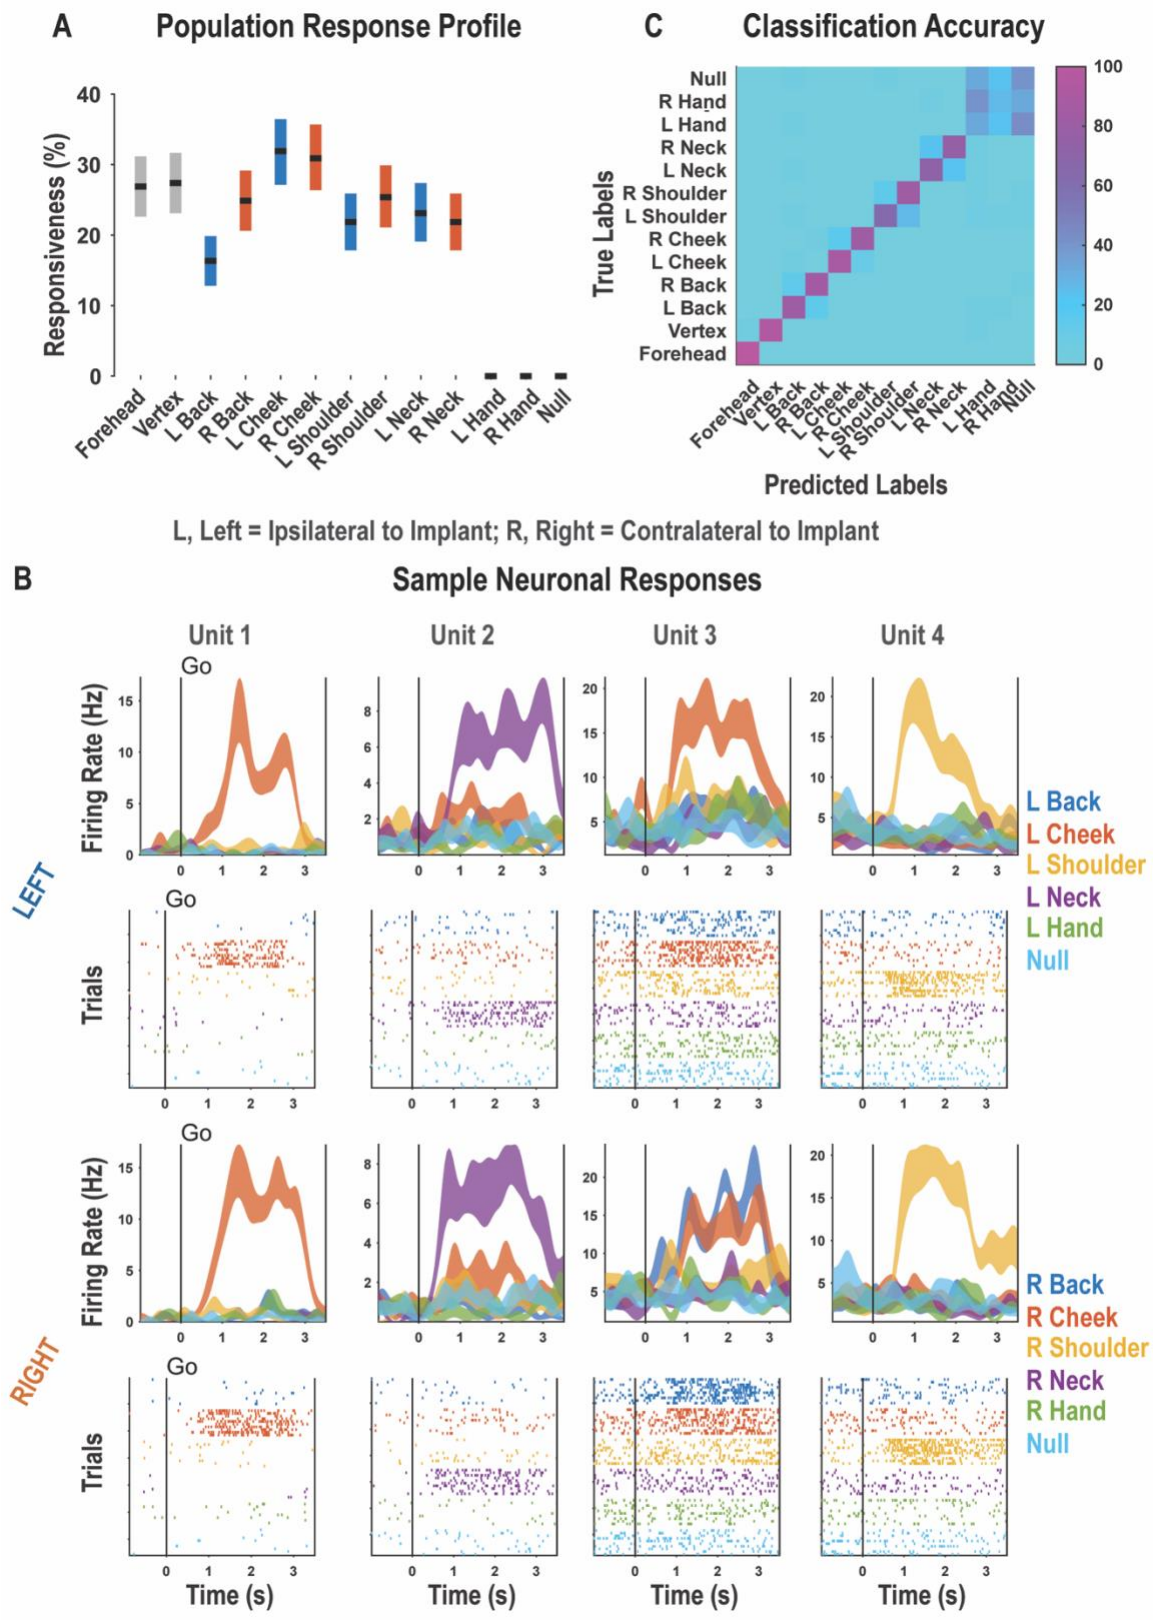
A
Population Response Profile
Responsiveness (%)
Forehead
Vertex
L Back
R Back
L Cheek
R Cheek
L Shoulder
R Shoulder
L Neck
R Neck
L Hand
R Hand
Null
C
Classification Accuracy
True Labels
Predicted Labels
L, Left = Ipsilateral to Implant; R, Right = Contralateral to Implant
B
Sample Neuronal Responses
Unit 1
Unit 2
Unit 3
Unit 4
Go
Firing Rate (Hz)
Trials
LEFT
RIGHT
L Back
L Cheek
L Shoulder
L Neck
L Hand
Null
R Back
R Cheek
R Shoulder
R Neck
R Hand
Null
Time (s)

Figure 3.1. PC-IP discriminably encodes bilateral tactile receptive fields.
A, Percent of the PC-IP neuronal population that demonstrated significant modulation relative to baseline for each tested stimulation site ($p<0.05$, FDR corrected, $n=398$ units). Results are shown as the mean percentage (horizontal black line) of the population ± bootstrapped 95% confidence interval (CI; bar height). Gray bars represent truncal (midline) body locations, blue bars represent left (ipsilateral) sided sites, and orange bars represent right (contralateral) sided sites. Population results were pooled across recording sessions (Figure 3.1-figure supplement 1) and were not qualitatively affected by pooling together single and potential multi-units (Figure 3.1-figure supplement 2). B, Representative neuronal responses illustrating body part discrimination. Each column of panels depicts the response for one neuron to body parts on the left (top two rows) and on the right (bottom two rows). The first and third row show the neural response (mean firing rate ± standard error on the mean; SEM, $n=10$ trials) as a function of time. The second and fourth rows show the spike rasters. Within these rows, each panel depicts the spike activity over each of the ten trials (rows) and time (x-axis), and are color coded by tested body site. The vertical line labeled "Go" indicates the start of the stimulus phase. C, Confusion matrix of the cross-validated classification accuracy (as percentage) for predicting body parts from population neural data. Colors represent the cross-validated accuracy, as in the scale. The matrix is an average of the confusion matrices computed for each recording day individually (Figure 3.1-figure supplement 3). (L, left body, ipsilateral to implant; R, right body, contralateral to implant; Hz, Hertz; s, seconds)

Single neurons were heterogenous, responding to variable numbers of touch sites (**Figure 3.2A, Figure 3.2-figure supplement 1**). Right and left sides tended to respond to the same number of fields (evidenced by the strong diagonal structure of **Figure 3.2A**). Tactile receptive fields of PC-IP neurons were diverse with evidence both for broad single peaked fields and multi-peaked fields characterized by spatially separated regions of enhanced response (**Figure 3.2-figure supplement 2, Figure 3.2-figure supplement 3**).

PC-IP neurons demonstrated mirror symmetric bilateral coding. We performed a cross-validated population correlation analysis to measure population level similarity in the responses to each touch location (**Figure 3.2B, Figure 3.2-figure supplement 4)**. In brief, the neural activation pattern elicited by touch to each body location was quantified as a vector, with each vector element capturing the mean response for a particular neuron during actual touch. These vectors were then pair-wise correlated in a cross-validated manner so that the strength of correlation between any two body-parts could be compared against the strength of correlation for repeated touches applied to the same body part. We found that responses to the same touch locations on the right and left sides are highly correlated, comparable to the correlation for repeated touches applied to the same body part. This result is consistent with a strong, mirror symmetric, bilateral encoding. As expected, correlations involving the hands and the null condition were distributed about zero, consistent with a lack of systematic neural population response to these conditions. The results from the correlation analysis were similar for alternative distance metrics (**Figure 3.2-figure supplement 5**). Further, analysis of single units revealed mirror symmetry in bilateral representation for the vast majority of the population, paralleling population level findings (**Figure 3.2-figure supplement 6**).

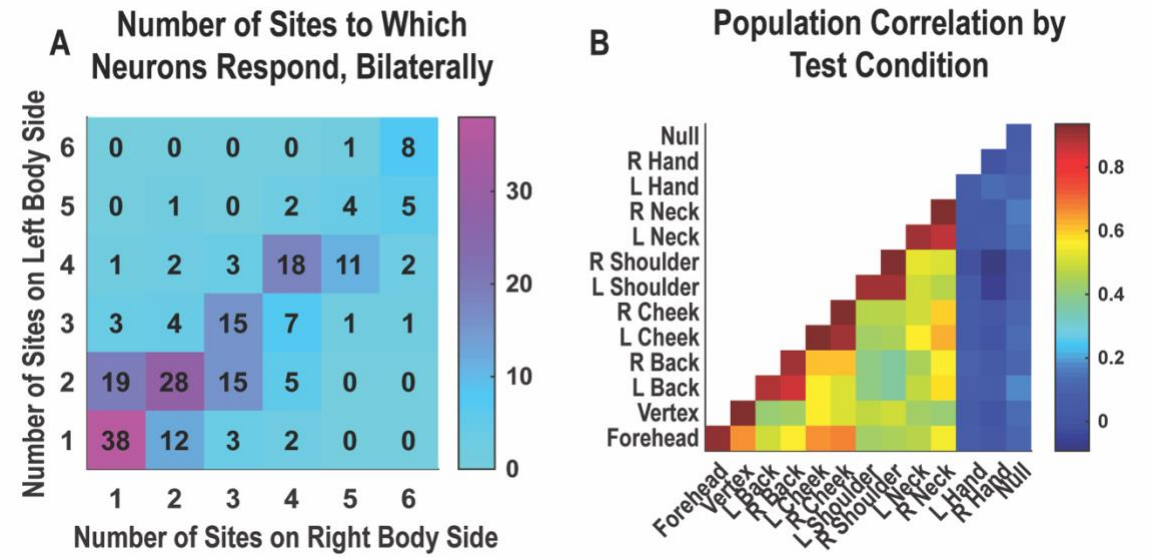

Figure 3.2. Neurons respond to variable numbers of bilateral receptive fields.
A, Matrix showing the number of neurons from within the PC-IP population that responded to the number of body parts shown, along each of the left and right body sides. Colors represent the number of neurons. Population results were not qualitatively affected by pooling together single and potential multi-units (Figure 3.2-figure supplement 1). Analysis of tactile recaeptive fields is shown in Figure 3.2-figure supplement 2 and Figure 3.2-figure supplement 3. B, Neuronal population correlation demonstrating the relation in encoding structure between body locations. Colors represent strength of correlation, as in the scale. Population results were not qualitatively affected by pooling together single and potential multi-units (Figure 3.2-figure supplement 4) or by choice of distance metric (Figure 3.2-figure supplement 5). For mirror-symmetry analysis at the single unit level, see Figure 3.2-figure supplement 6. (L, left body, ipsilateral to implant; R, right body, contralateral to implant)

3.3.2 Tactile responses occur at short latency to bilateral stimuli

We explored PC-IP population response latency to tactile stimulation on the contralateral and ipsilateral body sides. In a variation of the basic task paradigm, we used a capacitive touch sensing probe to acquire precise measurements of the time of contact with the skin surface in order to measure the latency in neuronal response from the time of tactile stimulation. We probed latency on the bilateral cheeks and shoulders. As a control, we included both hands in the task design.

We measured latency as the time at which the response of the neural population rose above the pre-stimulus baseline activity (**Figure 3.3**). The neural population response was quantified as the first principal component computed from principal component analysis (PCA) of the activity of all neurons [70,73]. The first principal component was then fit with a piece-wise linear function and latency was computed as the time the linear function crossed the baseline pre-stimulus response. Response latency was short for both body sides and was slightly shorter for contralateral (right) receptive fields (50 ms) than for ipsilateral (left)

receptive fields (54 ms) although this difference was not statistically significant (Permutation shuffle test, $p>0.05$). **Figure 3.3A** shows the time course of the first principal component relative to time of contact of the touch probe (stepped window; 2 ms window size, stepped at 2 ms, no smoothing) along with the piece-wise linear fit (dashed line). A bootstrap procedure was used to find the inter-quartile range of latency estimates (**Figure 3.3B**).

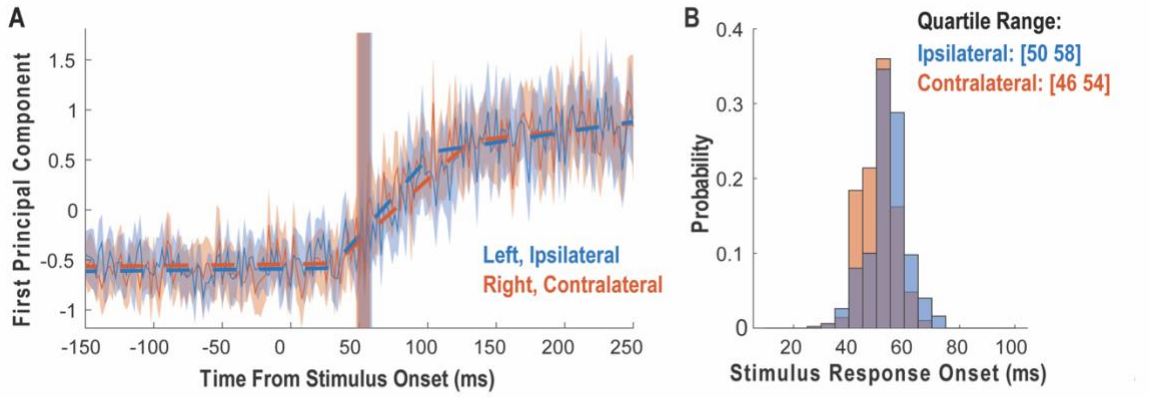

Figure 3.3. Tactile responses occur at short latency.
Population response was quantified as the first principal component (mean ± 95% CI). Population response was computed separately for the left (blue; ipsilateral to implant) and the right body sides (orange; contralateral to implant) and is shown as a function of time (2 ms window size, 2ms step size, no smoothing). Dashed lines show piecewise linear fit used to compute latency. Transparent vertical bar shows inter-quartile latency range based on a bootstrap procedure (see panel B). B) Distribution illustrating variability of latency estimates for the recorded data using a bootstrap procedure. Color code as in panel A.

3.3.3 Tactile imagery task evokes body part specific responses congruent with actual touch

The results thus far establish that PC-IP neurons have spatially structured tactile receptive fields that are activated at short-latency consistent with processing of tactile sensations. Are neurons that encode tactile sensations also recruited during tactile imagery? And if so, how might evoked neural responses compare to those arising from actual touch? To address these questions, we analyzed population activity elicited during a cue-delay-go tactile imagery task and compared the neuronal activity to that resulting from actual touch to matching body parts recorded during interleaved trials. During the imagery conditions, the participant was instructed to imagine touch to the right (contralateral) cheek, shoulder, or hand with the same qualities as the actual touch stimuli the participant actual during interleaved trials. A null condition was included as a baseline to measure neural activity when no stimulus was presented.

As with findings for actual touch, neuronal responses elicited during the tactile imagery task following the go cue (during the imagery phase) were discriminably encoded (**Figure 3.4A**, cross-validated accuracy 92%). High decode accuracy is consistent with the participant's compliance with task instructions and implies that the tactile imagery task elicited discriminative neural responses. A significant fraction of PC-IP neurons encoded actual touch to the cheeks and shoulders but not to the hands (**Figure 3.4B**; $\chi^2(1)=355.73$, $p<0.05$),

consistent with results presented in **Figure 3.1**. In comparison, a smaller fraction of the neuronal population was responsive to the cheek and shoulder during imagery of tactile stimuli (**Figure 3.4B)**. Of note, a significant number of neurons were active during imagined touch to the hand ($\chi^2(1)$**=**188.89, p<0.05), despite the hand being clinically insensate in the study participant (and despite actual touch to the hand not eliciting neuronal activation). The extent of overlap between the set of units active during actual and the tactile imagery condition is illustrated in **Figure 3.4C**. The degree of overlap, compared to what is expected by chance, was statistically significant (permutation shuffle test, p<0.05). Results were qualitatively similar for well-isolated single units (**Figure 3.4-figure supplement 1**).

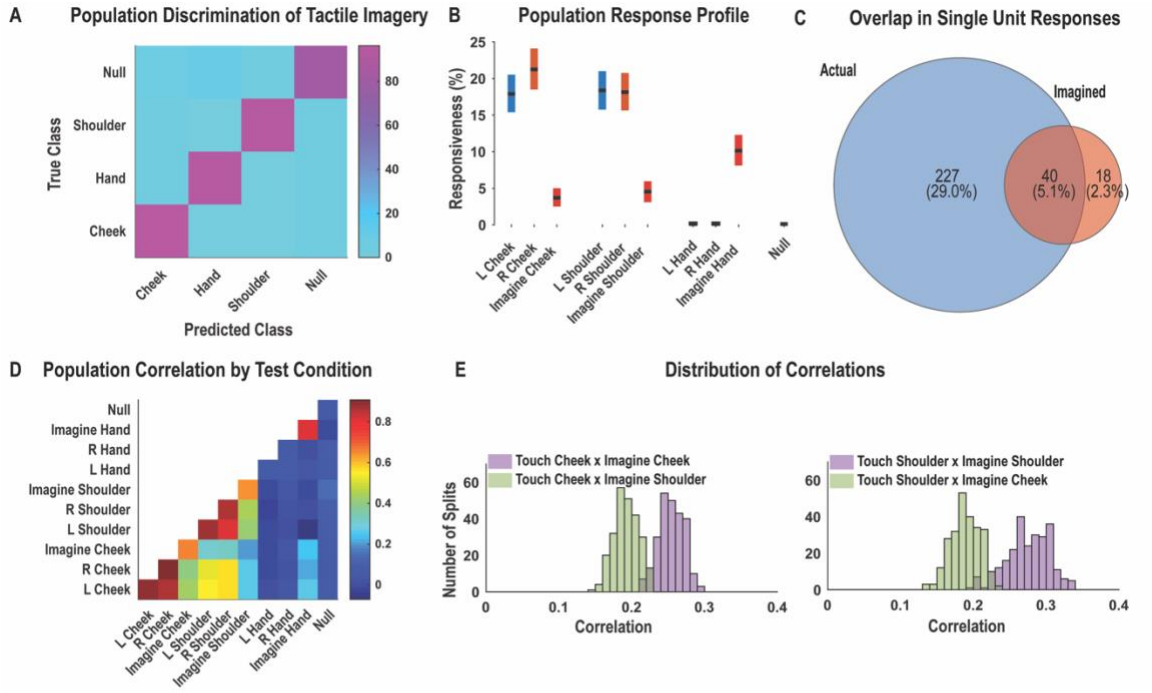

Figure 3.4. PC-IP neurons encode body part specific responses during imagery .
A, Average classification confusion matrix across recording sessions for body parts during tactile imagery and the baseline (null) condition. Colors represent prediction accuracy as a percentage, as in the scale. B, Percent of PC-IP neurons significantly modulated from baseline (mean ± 95% CI, p<0.05, FDR corrected, n=838 units) split by test condition. Population results were not qualitatively affected by pooling together single and potential multi-units (Figure 3.4-figure supplement 1). C, Venn diagram illustrating the numbers (percentages) of PC-IP neurons recorded that activated during actual and imagined touch, and their overlap. D, Population correlation matrix depicting similarity of the population response between all test conditions. Colors represent the correlation strength, as in the scale. E, Distribution of correlations between actual shoulder (left) and cheek (right) touch and imagined cheek/shoulder touches, with the distributions computed over different splits of the data (see Methods: Population correlation). (L, left body, ipsilateral to implant; R, right body, contralateral to implant)

We used the population correlation measure to compare population level neural activity across conditions (**Figure 3.4D**). Neural activity during the tactile imagery task shared a neural substrate with responses evoked by actual touch: representations evoked during the imagery task and during actual touch were more similar for matching body parts than for mismatched body parts (**Figure 3.4E,** permutation shuffle test p<0.05).

3.3.4 Dynamic evolution of population coding between task epochs suggests multiple cognitive processes

The analyses above were restricted to the mean neuronal activity following the go cue (e.g., during actual touch or during imagery) to allow a direct comparison with results reported for

the previous paradigms. We now expand this analysis. During the tactile imagery task, the participant heard a verbal cue specifying a body part (verbal cue = “cheek,” “hand,” or “shoulder”) followed approximately 1.5 seconds later by a beep instructing the participant to imagine the stimulus at the cued body part on the right side of the body. This cue-delay paradigm is standard in the motor physiology literature and is used to dissociate planning from motor execution related neural activity [3,74-76]. In our case, the cue-delay was unique to the tactile imagery condition. We utilized the cue-delay task to begin to dissociate in time neural activity related to different aspects of the task.

To leverage the benefits of the cue-delay paradigm, we performed a dynamic classification analysis (500ms windows, stepped at 100ms). Results are shown as a matrix (**Figure 3.5**). In brief, the diagonal elements represent the cross-validated prediction accuracy for a specific time window. The off-diagonal elements represent how well the classifier generalizes to alternate time windows. Each row can be interpreted as quantifying how well decision boundaries established for the diagonal time windows generalize to other time windows. This analysis allows us to measure when the neuronal population represents the different body parts (the diagonal) and whether population coding is similar or distinct during the task phases (the off-diagonal). We are interested in two main phases of the task, the early portion comprised of the cue and delay (cue-delay), and the later portion when the participant is actively imagining the stimulus (go/imagery). **Figure 3.5A** schematically illustrates examples of possible results. The examples are meant to be illustrative and are not an exhaustive list of possibilities. The population may be discriminative exclusively during the imagery phase, during the cue-delay and imagery phases but with distinct population coding, during the cue-delay and imagery phases with identical coding, or during the cue-delay and imagery phases with partially shared and partially distinct coding. Each pattern would suggest a different interpretation of various forms of cognitive processing that may be engaged in a tactile imagery task (see Discussion).

The results of our classification analysis (**Figure 3.5B**) are most consistent with body part selectivity during both the cue-delay and imagery phases, with partially shared and partially distinct population coding of the body parts between phases. The shared component is evident in the significant generalization accuracy in the off-diagonal elements, a representative row of which is shown in **Figure 3.5C** (blue portion) where cross-validated accuracy generalizes from approximately 70% within the cue-delay phase to approximately 60% during the imagery phase. The distinct population activity between phases is highlighted by a cross-validated Mahalanobis distance that provides a sensitive measure of change which is masked by the discretization process of classification (expanded rationale in Methods: Temporal dynamics of population activity). The findings demonstrate a significant change between the activity patterns in the cue-delay and the imagery epochs (**Figure 3.5C,** gray).

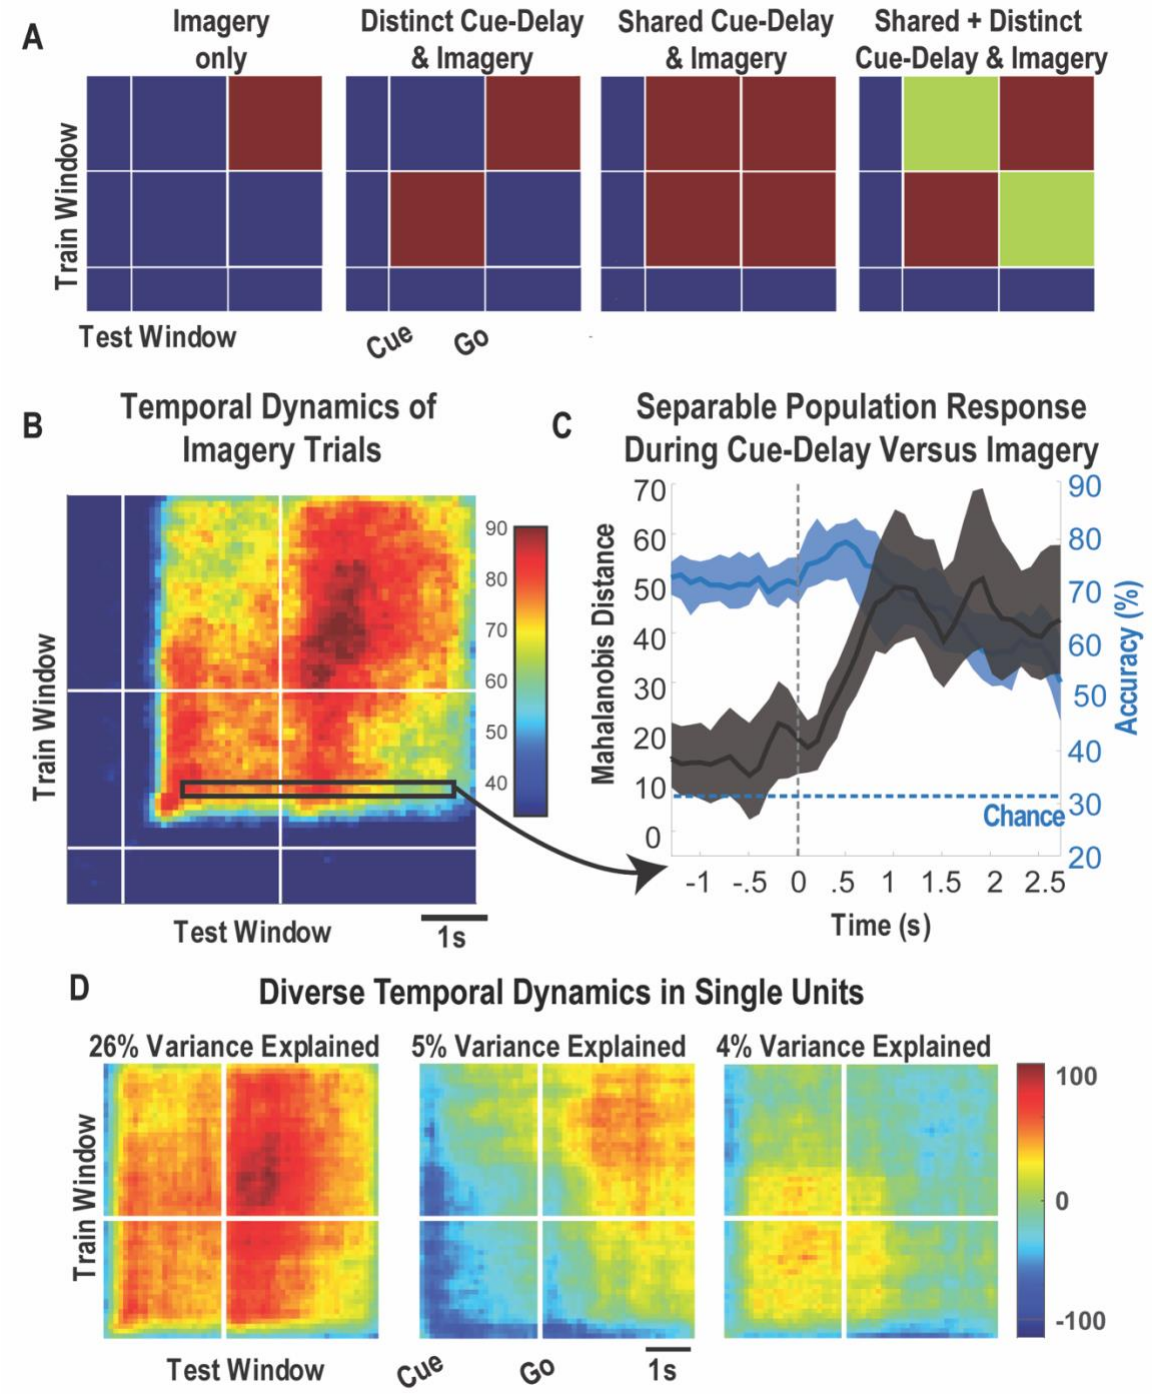
A
Imagery only
Distinct Cue-Delay & Imagery
Shared Cue-Delay & Imagery
Shared + Distinct Cue-Delay & Imagery
Train Window
Test Window
Cue
Go
B
Temporal Dynamics of Imagery Trials
Train Window
Test Window
1s
90
80
70
60
50
40
C
Separable Population Response During Cue-Delay Versus Imagery
Mahalanobis Distance
Accuracy (%)
Chance
Time (s)
D
Diverse Temporal Dynamics in Single Units
26% Variance Explained
5% Variance Explained
4% Variance Explained
Train Window
Test Window
Cue
Go
1s
100
0
-100

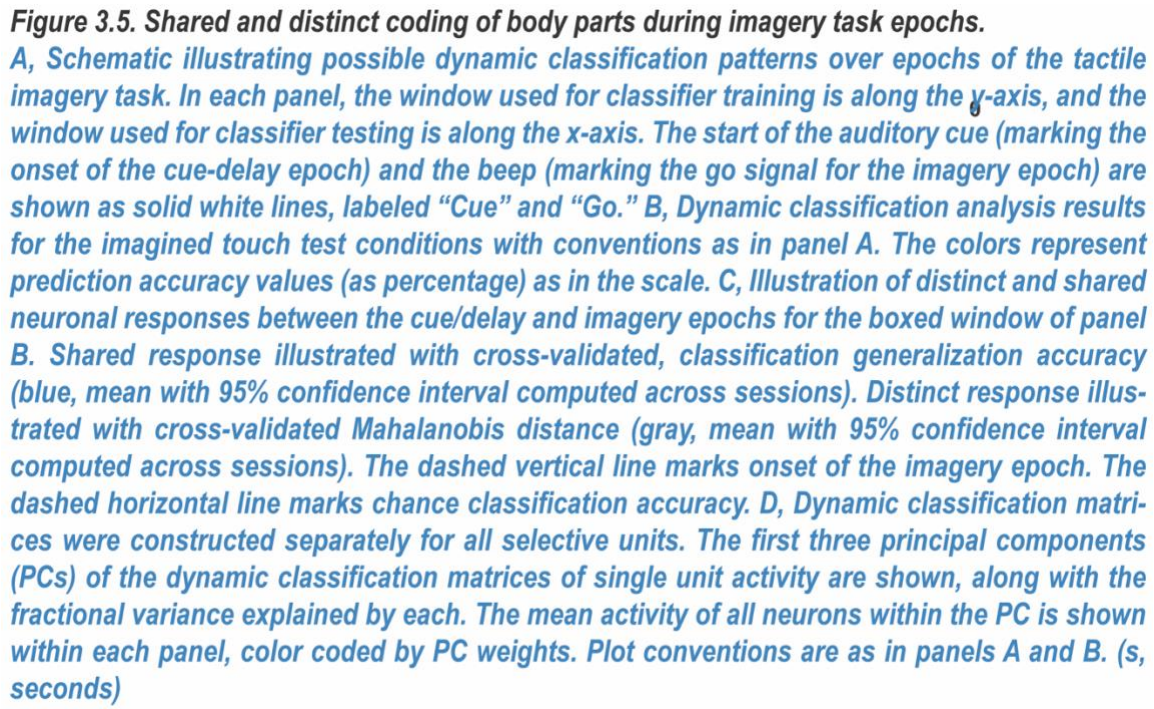

Figure 3.5. Shared and distinct coding of body parts during imagery task epochs.
A, Schematic illustrating possible dynamic classification patterns over epochs of the tactile imagery task. In each panel, the window used for classifier training is along the y-axis, and the window used for classifier testing is along the x-axis. The start of the auditory cue (marking the onset of the cue-delay epoch) and the beep (marking the go signal for the imagery epoch) are shown as solid white lines, labeled "Cue" and "Go." B, Dynamic classification analysis results for the imagined touch test conditions with conventions as in panel A. The colors represent prediction accuracy values (as percentage) as in the scale. C, Illustration of distinct and shared neuronal responses between the cue/delay and imagery epochs for the boxed window of panel B. Shared response illustrated with cross-validated, classification generalization accuracy (blue, mean with 95% confidence interval computed across sessions). Distinct response illustrated with cross-validated Mahalanobis distance (gray, mean with 95% confidence interval computed across sessions). The dashed vertical line marks onset of the imagery epoch. The dashed horizontal line marks chance classification accuracy. D, Dynamic classification matrices were constructed separately for all selective units. The first three principal components (PCs) of the dynamic classification matrices of single unit activity are shown, along with the fractional variance explained by each. The mean activity of all neurons within the PC is shown within each panel, color coded by PC weights. Plot conventions are as in panels A and B. (s, seconds)

To further clarify the properties of individual units, we conducted a dynamic classification analysis for each recorded unit. This resulted in the same matrices described above, but now each matrix represents how information coding evolves for a single unit. Time resolved classification data were then analyzed using principal component analysis, the first three principal components of which are shown in **Figure 3.5D**. A majority of variance (26%) is explained by units that are active during both epochs with similar coding. Coding during the imagery epoch exclusively or during the cue-delay epoch exclusively explained an additional 9% of variance.

3.3.5 Cognitive processing during the cue-delay and imagery epochs of the tactile imagery task shares a neural substrate with that for actual touch

Finally, we look at how encoding patterns through time generalize between the tactile imagery and actual touch conditions. A dynamic correlation analysis was applied both within and across the imagery and actual touch condition types (**Figure 3.6A**). In brief, the neural activation pattern elicited to each body location was quantified as a vector, and these vectors were concatenated to form a population response matrix for each condition type and for each point in time. These vectors were then pair-wise correlated in a cross-validated manner so that the strength of correlation between conditions could be assessed relative to the strength of correlation within condition, and across time. We found that the neural population pattern that defined responses to actual touch was similar to population responses both during the cue-delay or the imagery phases of the imagery task (**Figure 3.6A**). This implies that cognitive processing prior to active imagery as well as during imagery share a neural substrate with actual touch. Sample neuronal responses that help to understand single unit and population behavior are shown in **Figure 3.6B**.

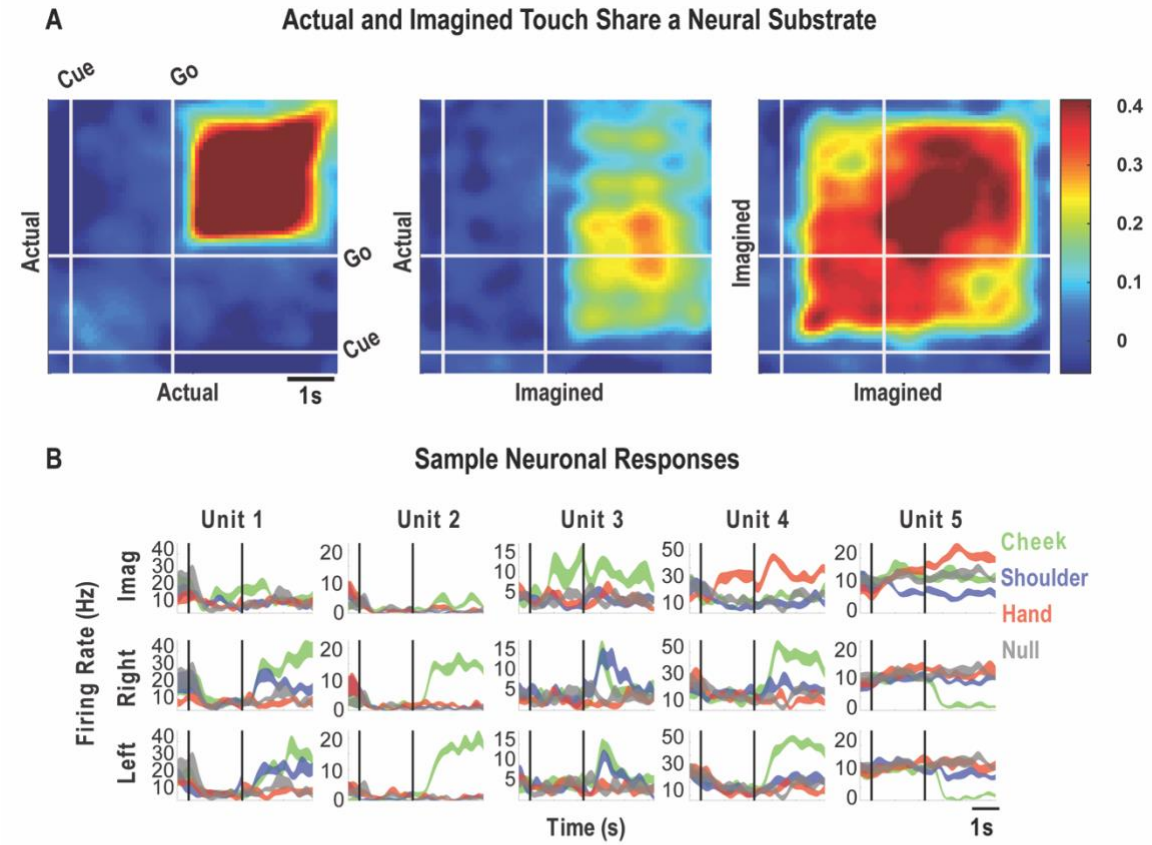

Figure 3.6. Cue-delay and imagery evoked neural activity shares a neural substrate with actual touch.
A, Within- and across-condition dynamic, cross-validated correlation analysis demonstrating a shared neural substrate between imagined and actual tactile sensations. Each panel shows how the neural population response at one slice of time compares to all other slices of time for the two formats being compared (x- and y- axis labels). Correlation magnitude is indicated by color as in the bar. The start of the auditory cue (marking the onset of the cue-delay epoch) and the beep (marking the go signal for the imagery epoch) are shown as solid white lines, labeled "Cue" and "Go." B, Representative neuronal responses illustrating selectivity during actual and imagined sensations. Each panel shows the firing rate (in Hertz, mean ± SEM) through time (vertical lines signal onset of cue/delay and go phases as labeled). Each column illustrates the responses of the same unit to tactile imagery of the right side (top), actual touch on the right side (middle), and actual touch on the left side (bottom) for matched body parts. (s, seconds; Hz, Hertz)

3.4 DISCUSSION

We have previously reported that human PPC encodes many action variables in a high-dimensional and *partially mixed* representation [2,4,43]. This architecture allows many parameters to be encoded by a small number of neurons, while still enabling meaningful relationships between variables to be preserved. Here we show that neurons recorded from the same electrode array in the same clinical trial participant are also selective for bilateral touch at short latency. Responses to actual touch are organized around body part, sharing population representations between the left and right side. Additionally, a tactile imagery task elicits body part specific responses that share a neural substrate with that for actual touch. Furthermore, we found neural selectivity during the active imagery epoch as well as during the cue and delay epochs that precede imagery. The distinguishable population

activity during these different phases indicates an encoding of multiple cognitive processes that may include semantic association, memory, attention, sensory anticipation, or imagery per se.

3.4.1 Human PC-IP encodes tactile stimuli with large and bilateral receptive fields.

Cortical processing of somatosensory information begins in the anterior portion of the parietal cortex (APC) within four cyto-architectonically defined areas termed BA 3a, 3b, 1 and 2 [77-79]. Each of these four sub-regions represents primarily contralateral somatosensory information [80-87]. Moving from the APC to superior regions of the PPC, spatially localized and segregated sensory representations become progressively more integrated, resulting in neuronal receptive fields that are larger, frequently encompassing multiple segments of the body [52,60,88-94] including bilateral encoding [85-87,95]. This process of integration is thought to play an integral role in sensory processing for the guidance of movement [18,32]. Our results, demonstrating that PPC neurons encode mirror symmetric spatially structured tactile receptive fields at short latency, are consistent with these prior reports. They further provide the first single neuron evidence supporting a role of human PPC in tactile processing. As hypothesized, when comparing the current results with our prior reports, we find that the same PPC neuronal population engaged by high-level motor cognition also encodes actual tactile sensations, providing a common neural substrate for sensory and motor processing [2,3,72].

3.4.2 Short latency tactile responses

In NHPs, reported latency to touch responses in primary somatosensory cortex (S1) from contralateral touch range between 19 and 23 ms [96,97]. PPC response latencies to touch are less clear, but neurons in the lateral intraparietal area within NHP PPC orient to visual stimuli at a mean latency of approximately 45 ms [98]. A recent human invasive electrocorticographic study reported mean latencies to visual response of approximately 60 ms in PPC [99] similar to the mean response latencies to visual stimuli within the occipital cortex [98,99]. Our own response latency to actual touch of ~50 ms compares well with these data and is consistent with rapid somatosensory processing within PPC for updating internal estimates of the body [18,32].

3.4.3 Tactile imagery dynamically invokes multiple cognitive processes in human PC-IP that share a neural substrate with actual touch

In motor neurophysiology, neural activity related to planning and execution are dissociated in time by introducing a delay between the cue instructing movement, and the movement in response to the cue [75,100]. We have previously found that such distinctions between planning and execution are preserved during motor imagery paradigms in tetraplegic individuals [3]. Here, a similar paradigm allowed temporal dissociation in cognitive processing during tactile imagery. Single units demonstrated three dominant response profiles (**Figure 3.5D**): 1) a shared selectivity pattern between the cue-delay and imagery epochs, consistent with

cognitive engagement during all phases of the imagery task, 2) selectivity exclusively during the cue-delay epoch but not the imagery epoch, and 3) selectivity exclusively during the imagery epoch but not the cue-delay epoch. In a previous study, we found similarly heterogeneous responses during the cue, delay and imagery epochs for imagined hand grasp shapes [101]. These single unit temporal selectivity profiles provide a basis for the population level findings of generalization in classification results between the cue-delay and the imagery epochs (**Figure 3.5B and Figure 3.5C**) but also a separation in neural state-space between these epochs (**Figure 3.5C**).

The tactile imagery task evoked body part specific cognitive activity that shared a neural substrate with actual touch within the PC-IP. Activity during imagined touch to the cheek, for example, was more similar in representation to actual touch to the cheek than to actual touch to the shoulder, and vice versa. Interestingly, the overlapping neural representations between actual touch and those elicited during imagery were not limited to the stimulus phase (actual touch and imagery) itself, but also extended to the cue-delay phase of the imagery task. This overlap echoes our recent findings for shared neural representations between imagined and attempted actions, as well as for shared neural representations between observed actions and action verbs [2,4]. These studies are consistent with views in which cognition recruits sensorimotor cortical regions [102-106]. We acknowledge that as with all passive neural recording studies, ours cannot establish a causal role for these neurons in tactile cognition. Understanding the unique contribution of PC-IP neurons within the larger network of brain regions engaged in cognitive touch processing remains to be explored. Nonetheless, our current results provide the first human single unit evidence of a shared neural substrate between tactile imagery and actual touch.

One concern with the use of all imagery experiments is that participant compliance cannot be externally validated. This raises the possibility that the participant is not performing the task or is performing the task in an unexpected manner. We think this is unlikely for three reasons. First, the subject by the time of this study was well versed in performing cue-delayed paradigms in the motor domain using both motor imagery and overt movements. In Zhang and Aflalo et al. 2017, the participant's performance of overt movements was perfect: the participant both performed the correct cued action and performed the action at the go cue (i.e., no movements began prior to the go cue as validated by measurements of electromyogram activity) [4]. Second, our current pattern of results that includes stable and accurate body part specific encoding within the cue-delay and imagery epochs, with a shift between epochs, is consistent with the participant performing the task as instructed. At a minimum, it is consistent with the participant's performing two distinct cognitive operations during the two primary phases of the task with remarkable trial to trial consistency. Third, evidence for a shared neural substrate between actual touch and the imagined touch conditions indicates that selective responses during the imagery task are related to tactile cognition.

3.4.4 What does neural processing within human PC-IP during tactile imagery represent?

While our task identifies dynamic engagement of multiple cognitive processes during tactile imagery, it is inadequate to precisely define the cognitive correlates of the observed neural activity. A number of cognitive processes may be engaged during the tactile imagery task including preparation for and/or execution of imagery, engagement of an internal model of the body, semantic processing of the auditory cue, allocation of attention to the cued body location or nature of the upcoming stimulus, and/or sensory memory for the corresponding actual sensation applied by the experimenter.

The precise neural correlates of tactile imagery are unknown, but evidence suggests that both imagined and actual touch may engage the same internal mental representations, or internal models, of the body [107,108]. Support for such a shared representation comes largely from the parallel domain of motor imagery [107]. Imagined and actual movements show similarity at the behavioral (e.g., similar duration), physiological (e.g., alteration of heart rate), and neural (e.g., activating the same neural substrates) levels [10,109-114]. These studies have been interpreted as evidence that imagined movements are the simulation of the internal models that track the state of our bodies during movement [40]. In powerful support of such a notion, we have shown that populations of neurons in human PPC code motor imagery and overt actions in highly similar ways [4]. The domain of tactile imagery has been less studied in comparison. However, relevant to the current paper, behavioral evidence has demonstrated that internal models of motor actions can influence sensory perception of touch [107]. Further, human neuroimaging studies suggest that overlapping brain regions are activated during both imagined and actual touch, including the PPC [108,115,116]. This points to not only a shared substrate for the representation of imagined and actual touch, but also to their likely engagement of similar internal models. Because an internal model may be involved in anticipatory or planning activity (and/or related to imagery), it could at least partly explain the pre-stimulus (post-cue, pre-imagination) neural activity we observed.

Another possibility is that the neural activity following the auditory cue in our study represents semantic processing of the cued word. Evidence suggests that a network of brain regions is activated in processing word meaning, including those involved in processing their higher-order sensory aspects, or motor intentions such as PPC [102,103,106,117-119]. Within this framework, semantic processing of the auditory cue (e.g., instructing imagined touch to the cheek) may engage the same population of neurons responsible for the higher-level processing of touch, consistent with our data. In support, we recently reported that action verbs and visually observed actions share a common neural substrate in the same PPC neural populations reported in the current study [2]. Results were consistent with automatic semantic processing as distinct from imagery. The current findings would extend possible semantic processing to the tactile domain and demonstrate neuronal selectivity for auditory cues (in addition to written text used in the previous study).

Hearing an auditory cue can direct the study participant's attention to the cued body part. Attention to a stimulated body part has been shown to enhance sensory processing in human neuroimaging [120-123]. In neurophysiological studies, this manifests as a gain in

stimulus responses [124,125]. However, during the imagery task, no stimulus was delivered to the participant. An attention account of our data would require that attention result in highly discriminable patterns of activity without a sensory stimulus (or pre-stimulus). Most studies of pre-stimulus attention report modest modulation of baseline neural activity [124-126]. However, the failure to find pre-stimulus effects may be the consequence of insensitive analysis techniques: Indeed, recent single neuron work in NHP visual cortical area 4 (V4) demonstrated discriminable coding for the locus of attention prior to stimulus presentation and, further, that the pre-stimulus activation patterns were systematically related to the post-stimulus response patterns [126]. These recent results suggest that attention may be decodable elsewhere, and they match the results presented in this manuscript. It is also consistent with what we have previously described as *partially-mixed* selectivity [4,43]. If our results are interpreted within the framework of attention, our current findings are inconsistent with a simple gain-like mechanism for attention, but instead suggest a richer mechanism by which information is selectively enhanced for further processing [126].

Our task was not designed to tease apart the different possible cognitive correlates of the observed neural activity engaged during imagery. We think the temporal dynamics of the signal indicate that multiple cognitive process may be engaged throughout the course of the task. The above cognitive phenomena may each independently engage the same neural population as distinct phenomena or may be distinct processes that nonetheless engage the same underlying neural substrate.

3.4.5 PC-IP and plasticity following spinal injury

The extent to which the human PPC reorganizes following SCI is unknown. Lesion studies in NHPs suggest that BA 3b and 3a, 1 and 2, show altered sensory maps following SCI, in a manner dependent on thalamic input from the afferent sensory pathways such as the dorsal column-medial lemniscus system [127]. With mid-cervical lesions, for instance, there is an initial loss of BA 3b hand representations, and a slight expansion in face representation at approximately two years [127,128]. Although significant axonal sprouting has been demonstrated to occur at the site of deafferentation in the spinal cord, with increased projections to brainstem nuclei, the changes observed in the somatosensory cortex are significantly smaller [127,128]. Moreover, the reorganization in higher order somatosensory centers such as the secondary somatosensory cortex is even more restricted than in BA 3b [128]. Similar stability in the topography of the somatosensory cortex has been identified in human subjects that have suffered limb amputations. In these amputees, there is a preserved digit map within the primary somatosensory cortex [129,130].

The results of our experiments suggest significant stability in tactile somatosensory architecture within the PPC. A substantial fraction of the neuronal population activated in response to imagined touch to the hand, where no response to actual touch was seen (insensate in the study participant), lending support to the idea that despite the lack of peripheral input from the hand due to the participant's spinal cord injury, the brain maintains an internal representation of tactile sensations [131]. The findings that intracortical

microstimulation produces discernable tactile perceptions from insensate body regions further reinforce a maintained representation of somatosensation after deafferentation [132,133]. These findings will prove useful for bidirectional neural prostheses. We acknowledge that while additional work probing cortical reorganization following SCI is necessary to fully understand its electrophysiological consequences, our results provide insight into the maintenance of basic tactile processing within the human PC-IP, after SCI.

3.5 CONCLUSION

Multiple lines of evidence indicate a critical role for the human PPC in the integration of convergent multimodal sensory information to enable complex cognitive processing and motor control. To date, however, its processing of somatosensory information at the single neuron level has remained fundamentally unexplored. In the unique opportunity of a BMI clinical trial, we examined the neural encoding of actual and cognitive touch within the human PC-IP. We found that local populations of PC-IP neurons within a 4x4 mm patch of cortex encode bilateral touch sensations to all tested body regions above the level of the participant's injury at short latency. A significant fraction of PC-IP neurons responded during the imagined touch condition with matching sensory fields to actual touch. The activity in the delay period of the task, between cueing and imagining touch, may reflect cognitive processes including tactile semantics, sensory anticipation, attention as well as active imagery. Together, our results provide the first single unit evidence of touch processing within the human PC-IP and identify a putative substrate for the encoding of cognitive representations of touch, thus far untested in animal models.

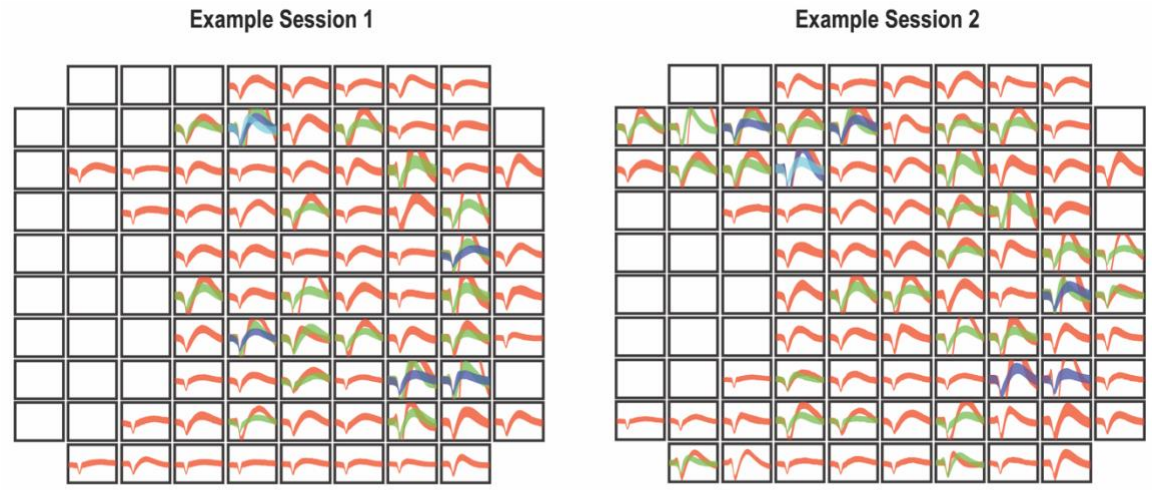

Figure 3.1-figure supplement 1. Nonstationary of waveforms across days indicate that recorded neurons are partially distinct.
We performed two analyses to quantify changes in neural recordings across days. In the first, we counted the number of waveforms on a channel and compared the number between days. If the number of waveforms changed, then this is strong evidence that there has been some substantial change in the neural recordings. By this measure, an average of 29% ± 4.2% of channels change across days. In the second analysis, we used a permutation shuffle test to measure whether the recorded waveforms on the same channel were more similar than waveforms across different channels. By this assessment, 58% ± 8% of channels change across days. These values indicate that there was some degree of neural turnover despite chronic recordings from the same implanted array. Representative changes are shown in the figure. In the two panels (left and right), the waveforms (mean ± SD) for all sorted neurons from each spatial position on the array are shown. An empty panel indicates no evidence of neural activity. Consistent with our analysis, both the presence and/or absence of waveforms as well as their shape can vary between sessions.

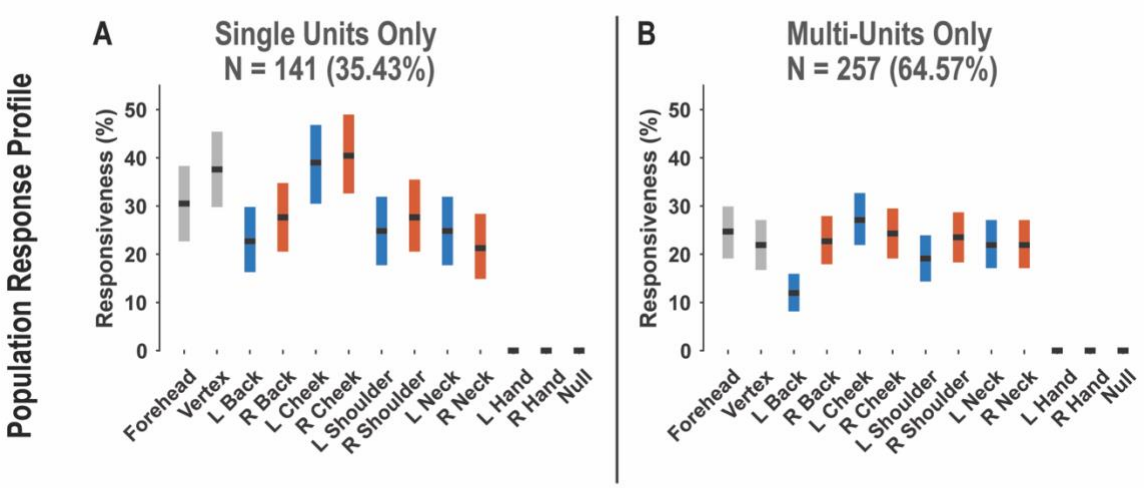

Figure 3.1-figure supplement 2. PC-IP population responsiveness is not qualitatively changed by pooling together single and multi-units.

A, Percent of the PC-IP neuronal population that demonstrated significant modulation (tuning) to each tested stimulation site ($p<0.05$, FDR corrected) when considering only high-quality single units. Within each bar, the black horizontal line represents the mean and the width of the bar represents the bootstrapped 95% confidence interval (CI). Gray bars represent truncal (midline) body locations, blue bars represent left sided sites, and orange bars represent right sided sites. B, Similar to panel A, but considering potential multi-units. (N, Sample size; L, Left; R, Right).

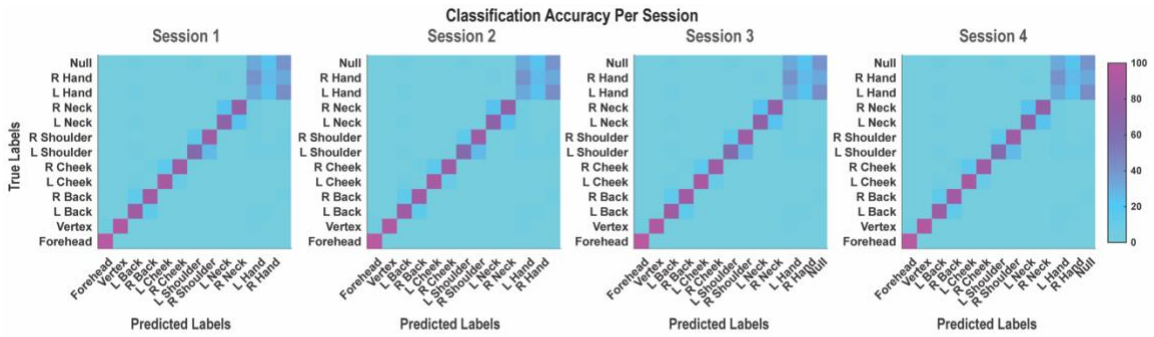

Figure 3.1-figure supplement 3. Population classification statistics were qualitatively unchanged between data recording sessions.
Each panel from left to right represents the confusion matrix of a labeled recording session. The cross-validated classification accuracy (as percentage) in predicting actual touch conditions from population neural data is shown. These data are averaged in the matrix shown in Figure 3.1C. (L, Left, R, Right).

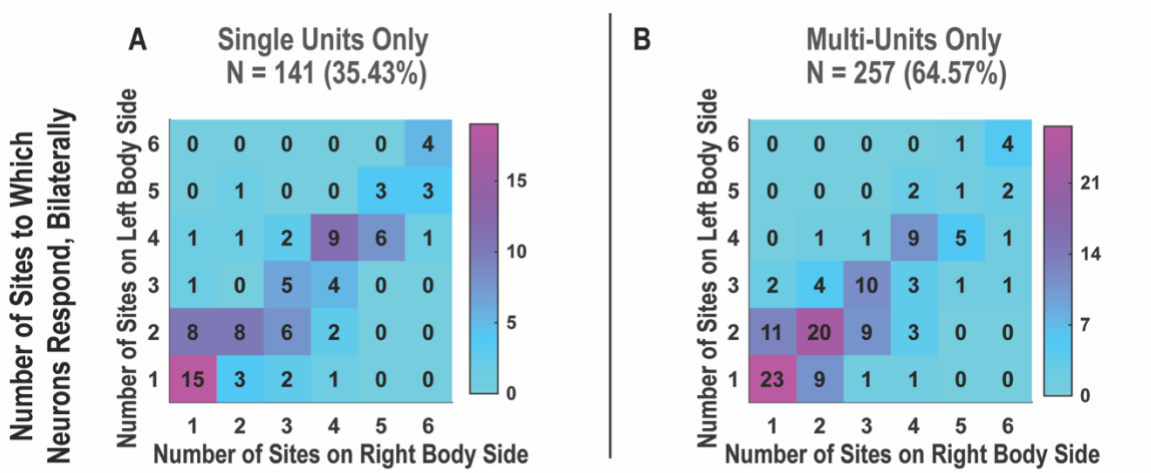

Figure 3.2-figure supplement 1. Bilateral responses to actual touch are not qualitatively changed by pooling together single and multi-units.

A, Matrix showing the number of neurons from within the population that responded to the number of body parts shown, along each of the left and right body sides. Only high-quality single units were considered in this matrix. The number of neurons is overlaid upon a heatmap in which the colors correspond to the number, as in the scale. B, Similar to panel A, but considering potential multi-units. (N, Sample size; L, Left; R, Right).

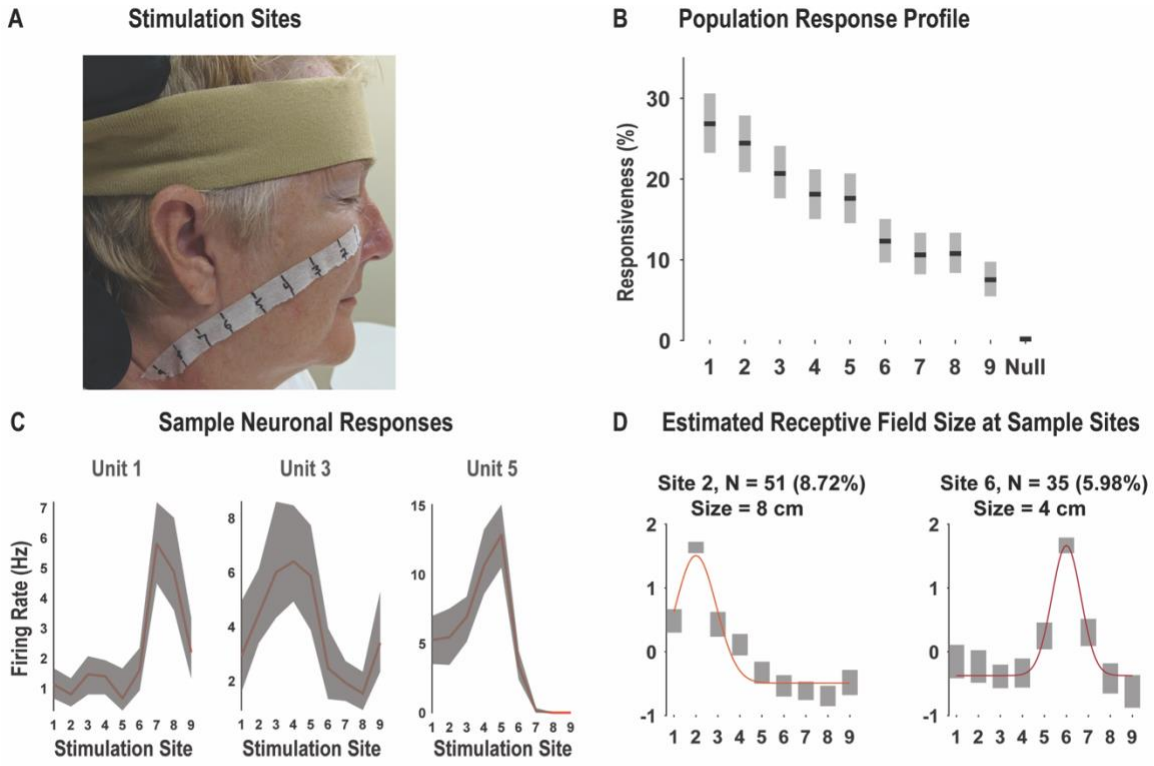

Figure 3.2-figure supplement 2. Diverse tactile receptive fields in human PC-IP neurons.

Tactile receptive fields of PC-IP neurons were diverse with evidence both for broad single peaked fields and multi-peaked fields characterized by spatially separated regions of enhanced response. A, Percentage of neurons characterized by a broad single peak response versus multiple peaks (mean ± 95% CI). The analysis is described in detail in Methods: Spatial extent of single neuron responses. In brief, for each neuron, we found the location of maximal response and asked whether we could find a second local maxima that rose significantly above the neighboring values. If no significant second local maxima was found, the neuron was categorized as single peak, otherwise, the neuron was categorized as multi-peak. B, Percentage of neurons characterized by a single peak versus multiple peaks split by their preferred response field (mean ± 95% CI). C, Similar to panel A, but considering potential multi-units. D, Similar to panel B, but considering potential multi-units. E, Sample single-peak neuronal responses. Four rows depict sample units grouped according to their preferred response field (from top to bottom: forehead, cheek, neck and shoulder). Within each panel, the neuron's firing rate (in Hertz) is depicted as a shaded blue region (the width indicating the interval from (mean - 95%CI) to (mean + 95% CI), to each of the four labeled locations. F, Sample multi-peak neuronal responses. Conventions as in panel C. (Hz, Hertz).

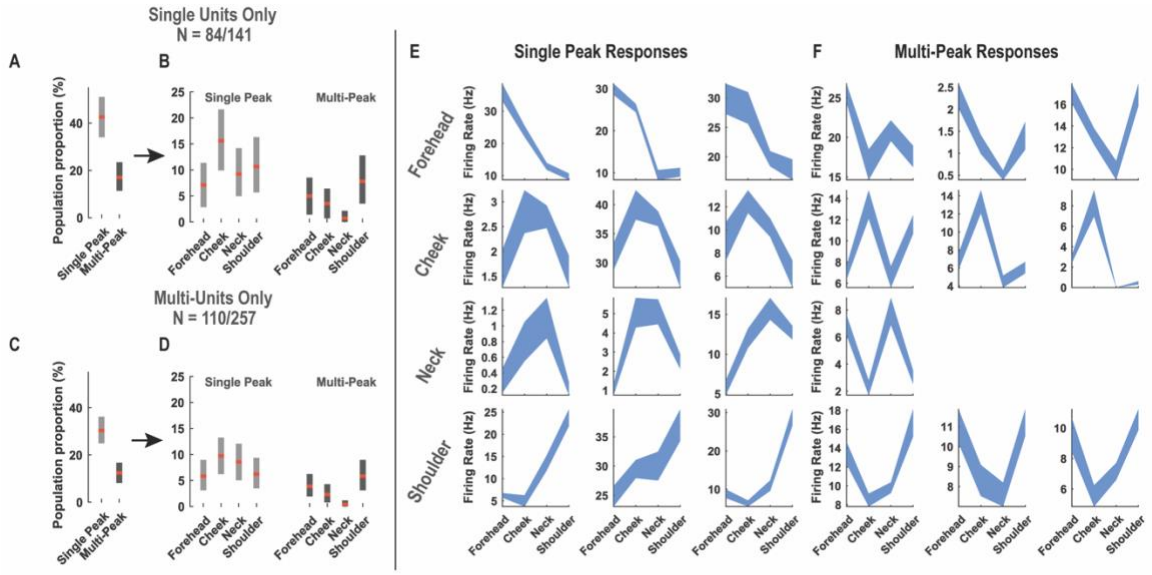

Figure 3.2-figure supplement 3. Receptive fields demonstrate local spatial structure.

In our first experiment (Figure 3.1) we found that neurons tended to demonstrate a region of peak response with gradual reduction of the response as a function of distance with the possibility for representations at larger distances. To better understand the local parametric structure of tactile receptive fields we performed a second experiment that tested collinear sites at finer spatial resolution. A, We performed a task variant aimed at estimating the size of neuronal receptive fields to actual touch (Methods: Receptive field size). In this variant, we tested actual touch to nine equally spaced points, two centimeters apart, along a line from the subject's cheek to her neck. This panel shows the location of stimulation points along the study subject's face and neck. Photo credit: Tyson Aflalo, California Institute of Technology. The face is masked to obscure identity per publisher's request. B, Percent of the neuronal population significantly modulated by touch to each stimulation location ($p<0.05$, FDR corrected, shown as mean±95% CI, n=585 units). C, Representative neuronal responses showing the firing rate at each stimulation site as a function of time (mean ± 95% CI; n=10 trials). D, Representative examples of neuronal responses and Gaussian fits used to estimate receptive field size (for full description, see Methods: Receptive field size estimation). Briefly, for each neuron demonstrating a differential response to touch to each of the nine fields, we identified the preferred site of stimulus delivery as the site associated with the largest firing rate. Next, we fit a Gaussian model to the average responses at the nine sites, with the mean/center of the model fixed at the preferred response site. The receptive field size was estimated as the full width at half maximum. In each of the examples depicted, the colored lines represent the Gaussian model fit to the average responses. The number of units that showed a preference for the labeled site (in the title of the subplot) is indicated above each subplot included in this analysis at each stimulation site is depicted above each subplot (as a percentage of the recorded neuronal population). The receptive field size was described as the full width at half maximum, shown above each subplot. In each subplot, the x-axis indicates the stimulation site. The y-axis is a standard (z) score. (Hz, Hertz; cm, centimeter)

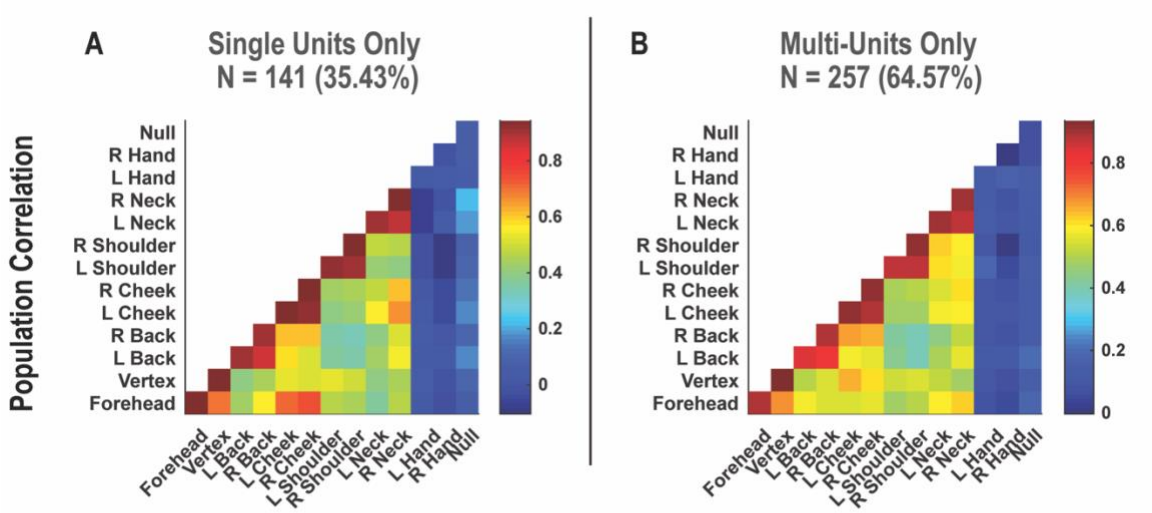

Figure 3.2-figure supplement 4. Symmetry in population level responses to bilateral touch are not qualitatively changed by pooling together single and multi-units.

A, Population correlation demonstrating the relation in encoding structure between test conditions, when considering only high-quality single units. Colors represent the strength of correlation, as in the scale. B, Similar to panel A, but considering potential multi-units. (N, Sample size; L, Left; R, Right).

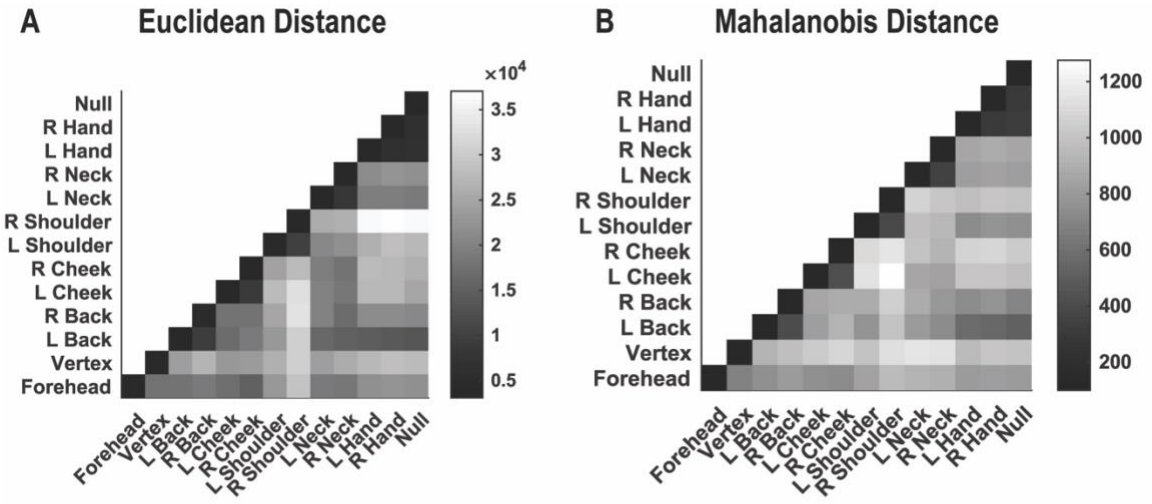

Figure 3.2-figure supplement 5. Relationships of population responses to touch are conserved across distance metrics.

A, Population similarity in encoding structure (for correlation data presented in Figure 2B) between test conditions when evaluated by Euclidean separation. Colors represent the distance, as in the scale. B, Population similarity in encoding structure (for the same data as in panel A) between test conditions when evaluated by Mahalanobis separation. Colors represent the distance, as in the scale. (L, Left; R,

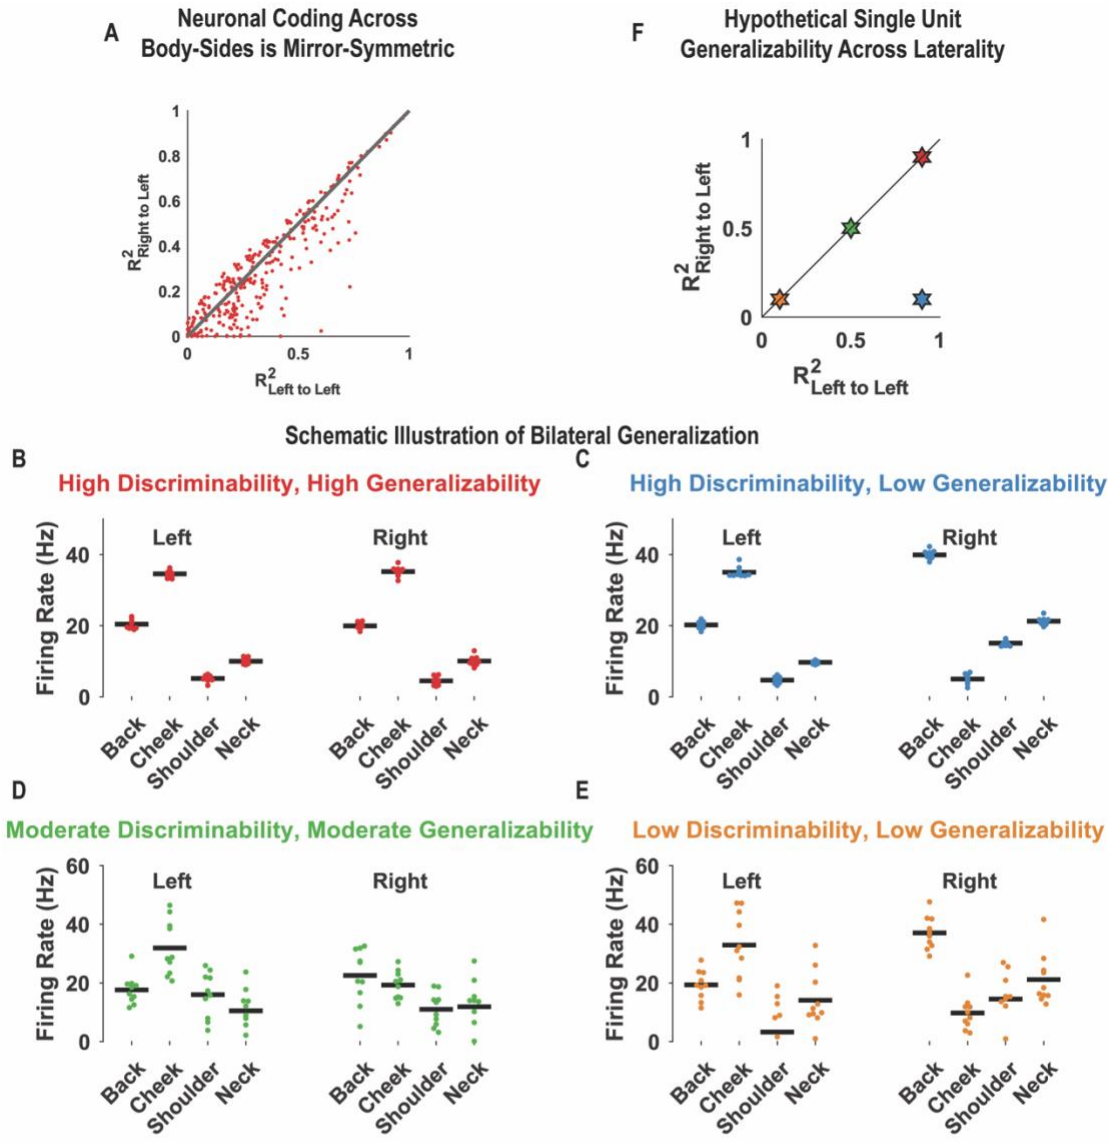

Figure 3.2-figure supplement 6. The majority of PC-IP neurons code tactile receptive fields bilaterally, in a mirror symmetric manner

We evaluated mirror symmetric bilateral coding at the level of single neurons. We used the same basic procedure as the population case (Figure 3.2B), although in the single unit case, cross-validated comparisons were done per-unit. All neurons demonstrating significant modulation from baseline for at least one location (263 of 398 units; 66.1%) were included in this analysis. A, Scatter plot comparing the consistency in touch responses to the same body side against responses across body-sides. Units close to the identity line code the left and right side with the same pattern of activity and thus in a mirror-symmetric fashion (see Methods: Tests for mirror symmetric neural coding of body locations: single unit analysis). The analysis used in this plot is depicted schematically in panels B-E. B, Firing rate for an example unit over ten trials (red dots) is shown, to touch to each indicated body part (both left and right). The black horizontal line indicates the mean response. The variance in trial-by-trial responses to each body part within a body side is low (i.e., high discriminability). Moreover, the pattern of responses on the left and right sides are closely matched (i.e., the encoding is highly generalizable across sides). C, Conventions are as in panel B. In this panel, the firing rate of a unit over ten trials (blue dots) is shown. As in panel B, the variance in trial-by-trial responses within each body side is low (i.e., high discriminability). However, the patterns of responses on the left and right sides are dissimilar (i.e., the encoding is poorly generalizable across sides). D, Conventions are as in panel B. Green dots represents the trial responses to touch. The inter-trial variance in responses is larger than in the top two panels (i.e., moderate discriminability). The pattern of responses between the left and right sides are moderately matched (i.e., moderate generalizability). D, Conventions as in panel B. In this panel, the yellow dots represent trial responses that have a high variance (i.e., low discriminability) and are also poorly matched in patterns between left and right sides (i.e., low generalizability). F, Scatter plot showing four illustrative neurons. These are color coded as in panel B-E and demonstrate how neurons with similar and dissimilar patterns of encoding will be arranged when their corresponding strength of encoding within a body-side (R2Left to Left) is compared to how well one side generalized to the other body-side (R2Right to Left).

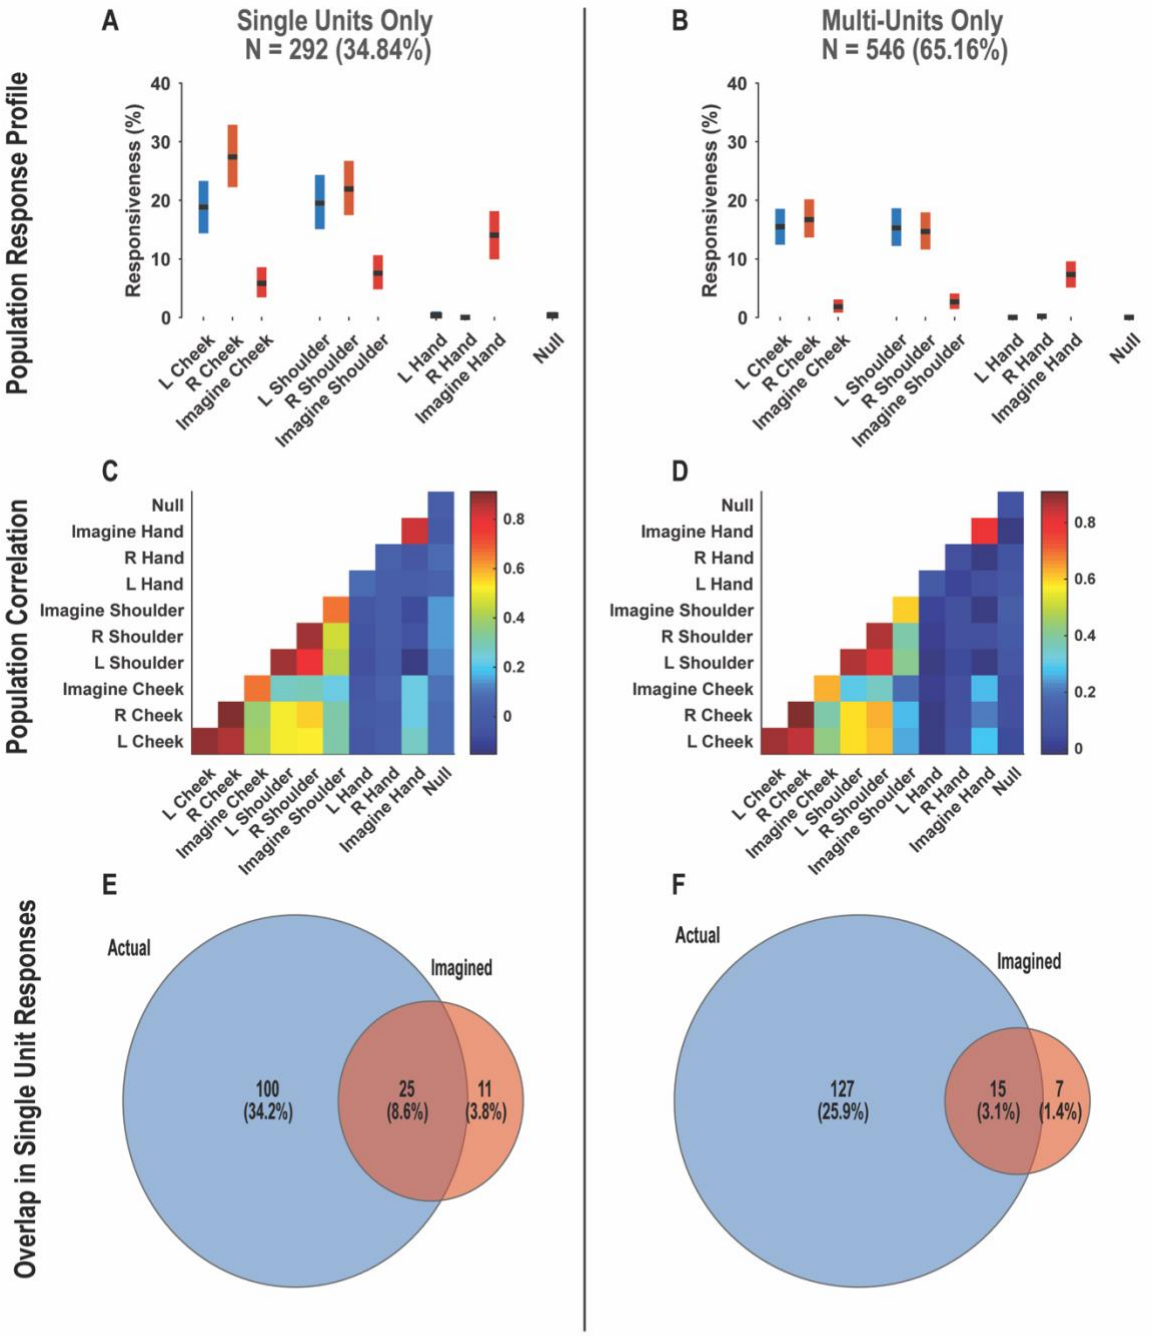

Figure 3.4-figure supplement 1. PC-IP responses during the tactile imagery task are not qualitatively changed by pooling together single and multi-units

A: Percent of the PC-IP neuronal population that demonstrated significant modulation (tuning) to each tested stimulation site ($p<0.05$, FDR corrected) when considering only high-quality single units. Within each bar, the black horizontal line represents the mean and the height of the bar represents the bootstrapped 95% confidence interval (CI). Blue bars represent actual touch to left sided sites, orange bars represent actual touch to right sided sites and red bars represent imagined touch (to right sided sites). B, Similar to panel A, but considering potential multi-units. C, Population correlation demonstrating the relation in encoding structure between test conditions, when considering only high-quality single units. Colors represent the strength of correlation, as in the scale. D, Similar to panel C, but considering potential multi-units. E, Venn diagram illustrating the numbers (percentages) of PC-IP neurons recorded that activated during actual and imagined touch, and their overlap, when considering only high-quality single units. F, Similar to panel E, but considering potential multi-units. (N, Sample size; L, Left; R, Right).

Chapter 4

MIRROR NEURONS ARE A MANIFESTATION OF A BROAD MECHANISM OF UNDERSTANDING

4.1 INTRODUCTION

We don't just see the world. We understand it.[7,13,24,25,27,106,134,135] From a brief video or even a still image of a person in action, we can infer what they are doing, why they are doing it, what they will do next, or what they might have done but didn't. A fundamental question in neuroscience is how neural populations transform sensory inputs into such deep and versatile understanding.

Mirror neurons have been proposed as the neural basis for such understanding, at least for how we understand what another person intends or feels.[24,25] Mirror neurons were discovered in high-level motor regions of the monkey brain and activated both when a monkey performs a specific movement, such as grasping an apple, and when the monkey observes someone else perform a corresponding movement.[25] These neurons inspired the *mirror hypothesis*: to understand the intentions of others, we map the visual representation of others' actions onto our own corresponding movement neurons. By simulating what we see within our own motor production circuits, we understand the intentions of others. Similar neurons have been postulated to enable understanding in other domains, such as emotions or sensations.[12,28,36-39,57] Collectively, these are termed the mirror system and are hypothesized to account for much of the way we understand the external world.

Yet, the mirror hypothesis has received numerous critiques.[29] For example, if understanding comes from activating our own high-level action representations, how can we understand actions we have never performed (e.g., jumping a skateboard)? Or could never perform (e.g., flying)? Human understanding is far more versatile than current data indicate mirror neurons can support. Limitations of the mirror hypothesis have led to speculation that mirror neurons may play other roles in behavior, such as guidance of movement based on the actions of others.[136] Alternatively, a potential role for these neurons in understanding may be clouded by an emphasis on interpreting the behavior of single neurons.[1] It is possible that analysis at the level of neural populations, and not individual neurons, can illuminate how these interesting neurons enable human-like understanding.[1]

We have recently recorded from populations of single neurons in the human posterior parietal cortex (PPC) during motor, sensory, and cognitive behaviors in human participants taking part in a brain-machine interface (BMI) clinical trial. We find that these neurons encode many diverse body-related variables such as action verbs, ***observed*** actions, motor and sensory ***imagery***, and motor plans.[2-5] Individual neurons are often complex, yet population level representations demonstrate shared encoding across these varied domains in a functional organization we have termed partially-mixed selectivity.[2-5] Based on these

past results we hypothesize here that “mirroring” is one manifestation of a more general mechanism by which we create shared internal representations of the world. To test this hypothesis, we recorded populations of neurons in human PPC while the participant experienced ***actual*** touch (to the subject) or ***observed*** touch (to another individual). We first validated the core mirror neuron phenomenon in the sensory domain of the participant. Then, to understand population-level coding mechanisms we manipulated the location of touch (cheek, shoulder) and touch type (pinch, press, rub, tap). We find that individual neural responses are highly variable. However, across the population, task conditions decompose along shared population-level neural subspaces that represent basic features of the stimuli such as body location and touch type. These shared, basic features generalize across task conditions, combining or being repurposed as necessary, like building blocks. Taken together, we speculate that populations of neurons within the PPC support versatile understanding not through a process of mirroring, but instead by encoding the representational building blocks of human cognition. Recent work in the cognitive neuroscience literature hypothesizes that human-like learning and thinking is largely enabled by how we build and utilize models of world to understand, explain, imagine, and plan. Human single neurons in PPC support this hypothesis by demonstrating that language, imagination, planning, and perception tap into the same underlying shared neural substrates.

4.2 METHODS

4.2.1 Subject details

All data were recorded from a 62-year-old tetraplegic female participant in a brain-machine interface (BMI) clinical trial. She has a high-cervical spinal cord injury (SCI) between cervical levels three and four, sustained approximately 10 years prior to the study, and with no preserved sensory or motor function below the shoulder-level. She was implanted with two 96-channel Neuroport Arrays (Blackrock microsystems model numbers 4382 and 4383) 6 years post-injury, in the left hemisphere. Informed consent was obtained, and she understood the nature, objectives, and potential risks, of the surgical procedure and the subsequent clinical studies. All procedures were approved by the Institutional Review Boards (IRBs) at the California Institute of Technology (IRB #18-0401), the University of California, Los Angeles (IRB #13-000576-AM-00027), and Casa Colina Hospital and Centers for Healthcare (IRB #00002372).

4.2.2. Experimental setup

All experiments were conducted at Casa Colina Hospital and Centers for Healthcare. NS was seated in a motorized wheelchair in a well-lit room. A 27-inch LCD monitor was positioned behind NS (visible to the experimenters but not to NS) too cue the experimenter for the presentation of stimulus. Cue presentation was controlled by the psychophysics toolbox (Brainard, 1997) for MATLAB (MathWorks).[67]

4.2.3 Physiological recordings

NS was implanted with one Neuroport array at the junction of the intraparietal sulcus (IPS) and postcentral sulcus (PCS), a region we refer to as PC-IP.[5] The other was in the left superior parietal lobule (SPL). Following surgery, the SPL implant did not function. Only data recorded from PC-IP were used in this study. Both arrays were explanted approximately one year after data in this study were collected.

Neural activity recorded from the array was amplified, digitized, and sampled at 30 kHz using a Neuroport neural signal processor. This system has received Food and Drug Administration (FDA) clearance for <30 days of recordings. We received an investigational device exemption (IDE) from the FDA (IDE #G120096, G120287) to extend the implant duration for the purposes of the BMI clinical study.

We have previously published our sorting algorithm.[4] Putative neuron action potentials were detected at threshold crossings of -3.5 times the root-mean-square of the high-pass filtered (250 Hz full bandwidth signal. Each waveform was made of 48 samples (1.6 ms), with 10 samples prior to triggering and 38 samples after. Single- and multi-unit activity was sorted using Guassian mixture modeling on the first three principal components of the detected waveforms. To minimize noise related effects, we used, as selection criteria, a mean firing rate greater than 0.5 Hz and signal to noise ratio (SNR) >0.5.

Well-isolated single and multi-units were pooled across recording sessions. To ensure that such pooling did not bias the conclusions of the paper, we performed core analyses on single-units alone, potential multi-units alone, and all units together. The results of these analyses, shown as supplemental figures for key results, and generally demonstrate that our results were robust to the pooling of all sorted units together.

4.2.4 Task procedures

4.2.4.1 Basic sensory mirroring task (BSMT; relevant for Figure 4.1). This task was performed to establish the shared responsiveness of PPC neurons to ***felt*** touch and to ***observed*** touch. NS sat facing an experimenter (actor). One experimenter stood behind the actor, and another behind NS. The task involved touch to one of two body parts (cheek, shoulder), to one of two persons (subject, or actor). Touch was provided as rubs performed bilaterally by the experimenter standing behind the person being stimulated, at approximately 2 rubs per second, for 2 seconds. Cheek touches were rubs parallel to the jawline (from cheek bone to chin and back again). Shoulder touches were rubs along the top of the shoulder, from near the neck to the outside of the shoulder and back. The task was performed on 6 individual recording sessions, with 10 trials per condition. In all, 805 units were recorded, of which 756 met selection criteria.

4.2.4.2 Multidimensional sensory mirroring task (MSMT; relevant for all Figures except 4.1). This task was performed to understand mechanisms by which neural information is shared

across populations of PPC neurons to support ***felt*** and ***observed*** touch. The basic setup was like the previous task. Here, however, we manipulated three dimensions: 2 body parts (cheek, shoulder), provided to 2 persons (NS, actor), in one of 4 types (pinch, press, rub, tap). As in the BSMT, touch stimuli were provided at approximately 2 per second, for 2 seconds. Rubs were as described. Pinches were performed in a non-painful manner with the thumb, index, and middle fingers. Presses were performed with the index and middle fingers and taps by the tips of the index and middle fingers. This task was performed on 7 recording sessions, with 10 trials per condition. In all 707 units were recorded, of which 699 met selection criteria.

4.2.5 Quantification and statistical analysis

4.2.5.1 Linear analysis (relevant for Figure 4.1B, Figure 4.2A). For each unit, we fit a linear model describing its firing rate as a function of response to each test condition. Response was defined as the mean firing rate between 0.5 after onset of the stimulus phase and ending 0.5 s after. These times were chosen to minimize the influence of experimenter delays in presenting the stimulus. The baseline was defined as firing rate during the 1 s prior to stimulus presentations. The linear model was written as:

$$FR = \sum_c \beta_c \, \mathrm{X}_c + \beta_0$$

Where FR is the firing rate, X_c is the vector indicator variable for test condition *c*, β_c is the estimated scalar weighting coefficient for each condition, and β_0 is a constant offset term. A neuron was considered responsive to a particular condition if the t-statistic for its associated beta coefficient was significant (p<0.05, false discovery rate (FDR) corrected for multiple comparisons).

4.2.5.2 Discriminability index (relevant for Figure 4.1C, Figure 4.2B). To quantify how well neural activity can be discriminated from baseline (pre-stimulus) activity, we developed a cross-validated measure: discriminability index (DI). As with the linear analysis described above, the stimulation phase window was defined as 0.5 after onset of the stimulus phase and ending 0.5 s after, and baseline was defined as the 1 s prior to stimulus presentation). The firing rate of all recorded neurons was concatenated into a vector, denoted by A. The firing rate of each neuron during the baseline phase was similarly concatenated to form a vector, denoted by B. Next, a non-dimensional discriminability index (DI) was computed as:

$$DI = \frac{\bar{A} - \bar{B}}{\sqrt{\frac{{\sigma_A}^2 + {\sigma_B}^2}{2}}}$$

Where $\bar{A}$ is the mean of the firing rate vector A, $\bar{B}$ is the mean of the firing rate vector B, σ_A is the standard deviation of the vector A, and σ_B is the standard deviation of the vector B.

4.2.5.3 Time-resolved classification (relevant for Figure 4.2C). Classification was performed using linear discriminant analysis (LDA) with the following parameter choices: (1) only the mean firing rates differ for unit activity in response to each test condition (covariance of the normal distributions are the same for each condition) and (2) firing rates for each unit are independent (covariance of the normal distribution is diagonal). The classifier took as input a matrix of firing rates for all sorted units. The analysis was not limited to significantly modulated units to avoid 'peeking' effects.[69] The analysis was performed independently for each recording session, and results were then averaged across days. In **Figure 2C**, this analysis was performed in a sliding-time window manner (300 ms each window, stepped at 10 ms intervals), beginning 0.5 s prior to the stimulation onset. Classification performance is reported as prediction accuracy of a stratified leave-one-out cross-validation analysis.

4.2.5.4 Correlation (relevant for Figure 4.1F). We performed cross-validated correlation to compare the neural representations of various test conditions (stimulus presentations) against each other in a pairwise manner. We quantified the neural representations as a vector of firing rates, one vector for each condition with each vector element summarizing the response of an individual unit. Neural activity was summarized as the mean firing rate during the stimulation phase window, defined as before (0.5 s after onset of the stimulus phase to 0.5 s after it ended). Firing rate vectors were constructed by averaging the responses across 50–50 splits of trial repetitions. The mean responses across different splits were correlated within and across conditions, then the splits were regenerated, and the correlation computed 250 times. The within-condition correlations assist in our interpretation of the across-format correlations by allowing us to quantify the theoretical maxima of the similarity measure (e.g., if the within-condition correlation is measured at 0.6, then an across condition of 0.6 suggests identical neural representations).

4.2.5.5. Event related averages (relevant for Figure 4.1E, Figure 4.2D, Figure 4.2E, Figure 4.2F, Figure 4.2-figure supplement 1). For each unit, neural activity was averaged within 750 ms intervals starting from 0.5 s prior stimulation onset, stepping to 2.5 s after, in 100 ms step intervals. Responses were grouped by condition, and a mean and standard error on the mean (SEM) were computed for each time window and for each condition.

4.2.5.6 Generalizability analysis (relevant for Figure 4.4 and Figure 4.4-figure supplement 1). This analysis was performed to understand how neural information belonging to one domain (e.g., body part) generalizes across another domain (e.g., touch type or person). For this analysis, we restricted the tested touch types to pinch and press only. Thus, the three dimensions we tested were: touch type (pinch, press), body part (cheek, shoulder), and person (***actual*** touch, or ***observed*** touch). In all, there are 8 test conditions.

We quantified the neural response to each condition as a vector of firing rates, one vector for each condition, with each element in each vector summarizing the response of an individual unit. All sorted units were used in this analysis. As in other analyses, neural activity was summarized as the mean firing rate beginning 0.5 s after onset of stimulation phase and ending 0.5 s after its conclusion. Next, we identified population-level neural subspaces

that optimally differentiate between pairs of conditions (training conditions) and testing how well these subspaces differentiate between pairs of test conditions. For example, in **Figure 4.4A**, we identified a subspace that optimally differentiates between body parts during **actual** pinches and asked: how well does this subspace also differentiate between body parts during ***observed*** pinches? To create the subspaces, we linearly regressed (using partial least squares regression) the vectors for the pair of training conditions, such that the variance (or separation) between each unit's response to the two conditions was maximized. We then used this model to compute the separability (or variance) in neural space between test data from the test pair of conditions. This separability in neural space for the test data was normalized to the cross-validated separability in space between held out data from the training pair of conditions (in a stratified, cross-validated manner), between -1 and 1. This basic computation was performed in reverse as well, such that in this example, a subspace was created that optimally differentiated between ***observed*** cheek pinch and ***observed*** shoulder pinch, and tested to identify separability between ***actual*** cheek pinch and ***actual*** shoulder pinch. The results in both directions were averaged and recorded as the normalized generalizability of information (in this example case, generalizability of information separating body parts across ***actual*** and ***observed*** touch). The generalizability was computed for each day independently and averaged across recording sessions. Confidence intervals were estimated using a bootstrap procedure.

The generalizability analysis was performed on both real neural data, as well as a population (same size as the real number of neurons) of synthetically generated mirror neurons. The neurons were generated by assigning identical firing rates to actual and observed touch to a randomly chosen body part, and a randomly chosen touch type (between pinch and press). A vector of firing rates (of length 10, for ten trials) was constructed, with a mean 1, and a standard deviation of 1.2, for that pair of conditions, and 0 for all other pairs of conditions. This population of synthetic neurons was then used as the dataset for the generalizability analysis (**Figure 4.4-figure supplement 1**).

4.2.5.7 Demixed principal components analysis (dPCA) (relevant for Figure 4.5). We used the most recent version of dPCA with the default parameters.[137] dPCA is an analysis technique that decomposes neural data along user-defined neural dimensions (marginalizations) that capture variance related to variables in the experiment. This decomposition allows understanding the structure in neural data as it relates to the experimentally manipulated variables. We used all sorted units in this analysis. dPCA takes as input a matrix that summarizes (for all units to be included), their firing rates to each of the test conditions (i.e., all combinations of experimental variables; in the MSMT this is the full factorial, or 16 conditions), along each trial repetition. Neural activity was averaged within 750 ms intervals starting from 0.5 s prior to the onset of the stimulation phase, stepping to 0.25 s after, in 50 ms step intervals. In our current study, we were interested in understanding how much of the population variance was explained by independent dimensions (i.e., body-part, touch type, and person being touched) as well as by interaction terms (touch type x body-part, body-part x person, touch type x person, and touch type x person x body-part). Thus, we used all 7 possible marginalizations within this analysis.

4.2.5.8 Modeling the single neuron response properties to various test conditions (BSMT) (relevant for Figure 4.1D). This analysis was performed to understand how individual neurons responded to four formats: ***actual*** cheek touch (Ac), ***actual*** shoulder touch (As), ***observed*** cheek touch (Oc), and ***observed*** shoulder touch (Os). Various possibilities exist. For example, the neuron might respond to ***actual*** touch to both body parts but not to any ***observed*** touch. Alternatively, it could respond to both ***actual*** and ***observed*** touch to the one body part but not to the other. We can model the firing rate for a given unit as:

$$fr = \alpha \cdot A_c + \beta \cdot A_s + \gamma \cdot O_c + \delta \cdot O_s$$

Where fr is the firing rate for the unit, A_c, A_s, O_c, O_s are the four formats, and α, β, γ, and δ are the weighting coefficients for each format, respectively. If the unit does not respond to a format, then the dot product of the unit's weighting coefficient and the format collapses to a scalar value. If a unit responds to two formats in a matched manner, then the weighting coefficient for these two formats will be identical. For the analysis, we allowed a weighting coefficient to be either 0 or 1, such that that across 4 formats, there are a total of 16 possible models for each neuron. We fit the parameters of each of the 16 models using standard linear regression techniques (see above), and the results were compared. As selection criteria to evaluate the "best" model from all candidate models, we used the Bayesian information criterion (BIC) and cross-validated coefficient of determination (cvR^2). The models were grouped according to four categories: **invariant** (in which the weighting coefficient was identical across all formats), **body part** specific (in which the weighting coefficient was invariant for matched body parts, but not for mismatched body parts), **person** specific (in which the weighting coefficient was invariant for touch to the same person) or **idiosyncratic** (all other combinations).

4.2.5.9 Modeling the single neuron response properties to various test conditions (MSMT) (relevant for Figure 4.3, Figure 4.3-figure supplement 1, Figure 4.3-figure supplement 2, Figure 4.3-figure supplement 3, Figure 4.3-figure supplement 4, Figure 4.3-figure supplement 5). This analysis is like the earlier modeling analysis for the BSMT, except it has been expanded to accommodate for more test conditions. To understand the breakdown of individual units that create the population response, we first defined four formats: ***actual*** cheek touch (Ac), ***actual*** shoulder touch (As), ***observed*** cheek touch (Oc), and ***observed*** shoulder touch (Os). An individual neuron could respond to one or more formats. If it responds to more than one format, it could respond with a matched selectivity pattern (SP; the precise pattern of responses) to each of the four touch types (pinch, press, rub, tap) within the format, or with a mismatched SP. Across the four formats, the firing rate for a given unit can be described mathematically as:

$$fr = \alpha \cdot A_c + \beta \cdot A_s + \gamma \cdot O_c + \delta \cdot O_s$$

Where fr is the firing rate for the unit, A_c, A_s, O_c, O_s are the four formats, and α, β, γ, and δ are the weighting coefficients for each format, respectively. If the unit does not respond to a

format, then the dot product of the unit's weighting coefficient and the format collapses to a scalar value. Within this type of a linear model, if a unit responds to formats with an identical SP, then the weighting coefficient for all those formats will have an identical weighting coefficient. In all, there are 51 unique models for all the ways in which SPs can be expressed across formats.

To determine how SPs compared across formats, we fit the parameters of each of the 51 models using standard linear regression techniques (see above), and the results were compared. As selection criteria to evaluate the "best" model from all candidate models, we used the Bayesian information criterion (BIC) and cross-validated coefficient of determination (cvR^2).

4.3 RESULTS

We recorded populations of single neurons from NS, a tetraplegic individual (spinal injury at levels 3–4; C3/4) participating in a BMI clinical trial. Recordings were made from a microelectrode array implanted at the junction of the post-central sulcus (PCS) and intraparietal sulcus (IPS). NS could not move her arms or hands, complicating studies that study mirror-like responses using ***executed*** and ***observed*** movements. However, NS had well-defined sensory receptive fields that responded at short latency thus opening the possibility studying mirror-like phenomena in the sensory domain.[5] We performed two primary experiments: The first experiment used a simple paradigm to validate whether PPC neurons have mirror-like properties, firing in similar ways to ***experienced*** and ***observed*** tactile sensations. After confirming mirror-like responses, we performed a second experiment, expanding the number of task dimensions to test whether PPC populations support a general mechanism.

4.3.1 Mirror-like responses in human PPC single neurons

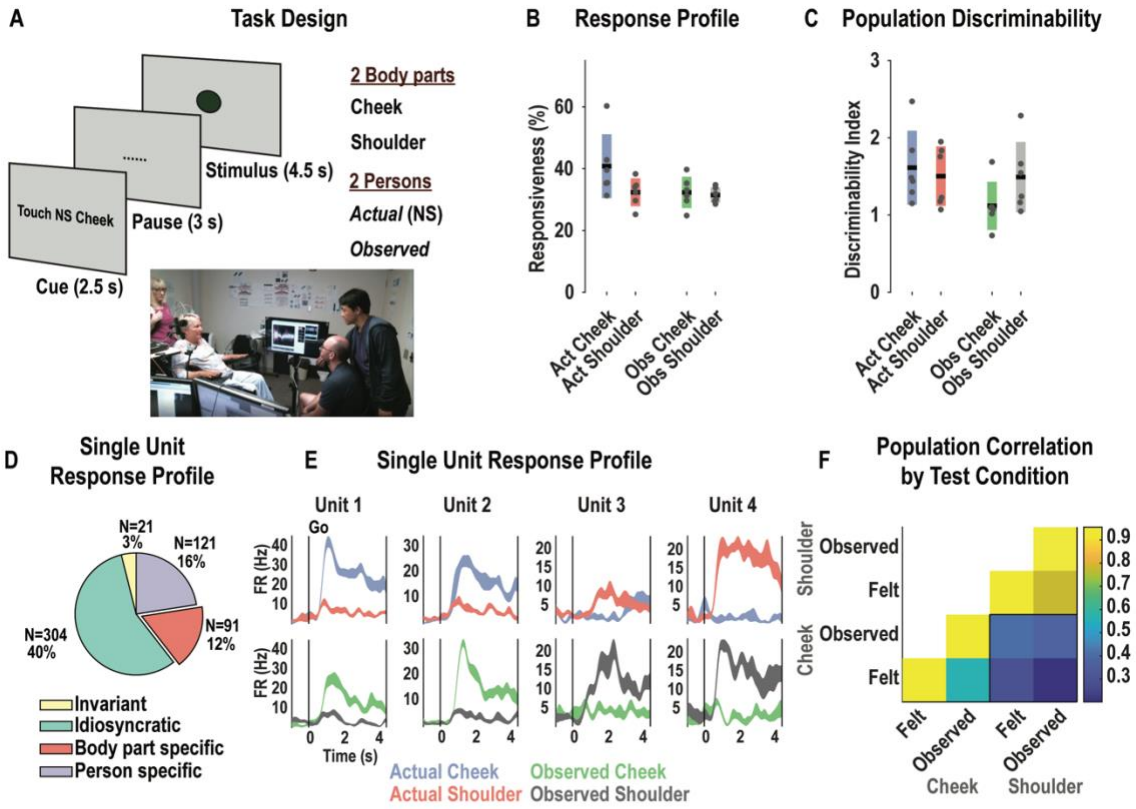

Figure 4.1. PPC neurons encode actual and observed touch in a mirror-like manner

A, Schematic illustrating the task design, in which PPC coding of touch responses was examined during manipulation of 2 dimensions: body part (cheek, shoulder) and person (actual touch to the subject, or observed touch to an experimenter). Thus, there were 4 test conditions in all: actual cheek touch, actual shoulder touch, observed cheek touch, and observed shoulder touch. Touch was provided as rubs by experimenters standing behind the subject and the seated experimenter as shown in the panel on the bottom. The phases of each trial (n=10) are indicated, along with the timings for each epoch. B, Percent of the PPC neuronal population that demonstrated significant modulation relative to baseline for each tested condition (p<0.05, FDR corrected, n=757 units). Results are shown as the mean percentage (horizontal black line) of the population ± bootstrapped 95% confidence interval (CI; bar height). Gray dots represent single session results. C, Discriminability of neural population activity (y-axis) during each of the test conditions (x-axis), from baseline neural activity. Results are shown as the mean value (averaged across sessions; horizontal black line) ± bootstrapped 95% CI. As in panel B, gray dots represent single session results. D, Pie chart illustrating the breakdown of single units (as a percentage of all significantly modulated units, n=537), according to their individual response properties: body part specific (invariant to whether touch was actual or observed), person specific (invariant to body part), invariant (responsive to all conditions), or idiosyncratic (other patterns). E, Example neurons showing mirror-like responses. Each column of panels shows the response for one unit to actual touch to both body parts (top row) and observed touch to both body parts (bottom row), indicated by the color code). The neural response to each condition is shown as mean firing rate ± standard error on the mean (SEM), n=10 trials. F, Neuronal population correlation demonstrating that at the population-level, actual and observed touch responses to matched body parts are more similar than to mismatched body parts. Colors represent the correlation strength, as in the scale. S, seconds; Act, actual; obs, observed

In preliminary investigations, we found many compelling examples of neurons demonstrating the specificity and congruency that are the defining feature of mirror neuron responses. Specificity is defined as the selective activation for a restricted set of tactile sensations, such as neural response when NS ***felt*** her outer shoulder being rubbed but not her inner shoulder. Congruency is defined as having the same selective activation for ***experienced*** and ***observed*** sensations. We performed a basic sensory mirroring task (Experiment 1 – BSMT, **Figure 4.1A**) to quantify whether sensory mirror-like responses could be robustly identified in our population of human single neurons. We recorded 757 ± 19.72 neurons over 6 sessions while the participant ***felt*** rubbing motions applied to her cheek or shoulder or ***observed*** rubbing motions applied to an experimenter's cheek or

shoulder. This two-factor (body part x person) design allowed us to test for specificity and congruency. We found robust coding of ***experienced*** and ***observed*** tactile sensations (population Mahalanobis distance from baseline with 95% CI, **Figure 4.1B**; percent of cells modulated from baseline, t-test $p<0.05$ corrected for multiple comparisons, **Figure 4.1C**). Many neurons demonstrated mirror-like responses, firing in a similar manner to touches to the cheek or shoulder, invariant to whether the touches were ***felt*** or ***observed*** (model analysis, body part specific, $p<0.05$ corrected, **Figure 4.1D**). Example neurons showing these mirror-like responses are shown in **Figure 4.1E** (firing rate thorough time). We found that the mirroring phenomena was robust at the neural population level. We summarized the response of the entire population during each condition as a vector of the mean firing rate while the participant ***experienced*** or ***observed*** touch. We found a significant correlation between the ***observation*** and ***experience*** conditions for matching body parts. Further, the correlation was higher for matching body parts than mismatched body parts (t-test, $p<0.05$, **Figure 4.1F**). These data provide powerful support for the basic phenomena of neurons responding in similar ways to both ***experienced*** sensations and ***observed*** sensations and is consistent with previous interpretations that we may understand the tactile sensations ***experienced*** by others through simulation within our own tactile processing system, either through single-unit or population mechanisms. However, from this simple paradigm, it is unclear whether these shared responses are direct evidence for a mirror mechanism, or instead, may be part of a more general mechanism in which shared representations support cognition.

4.3.2 Unpacking the population code: multidimensional sensory mirroring task

To better understand shared encoding between ***experienced*** and ***observed*** actions, we performed a second experiment (Experiment 2), that augmented the first experiment to include four different types of touch (pinch, press, rub, and tap). These touch-types were selected as they resulted in perceptually distinct stimuli both under ***observed*** and ***actual*** touch conditions, and not based on assumptions about the underlying selectivity of recorded neurons. Thus, in the updated task there are three manipulated dimensions (body part, touch type and person) combined in a full factorial design for a total of 16 conditions. Including the additional dimension allowed us to 1) test whether encoding similar variables in similar ways was a ubiquitous property of the neural population; and 2) test whether these shared responses are consistent with a compositional basis, encoding multidimensional sensations as a combination of basic sensory properties. We recorded from an average of 654 ± 22.40 neurons over 7 sessions using the second experiment. We briefly describe single unit response properties before jumping into our two main tests.

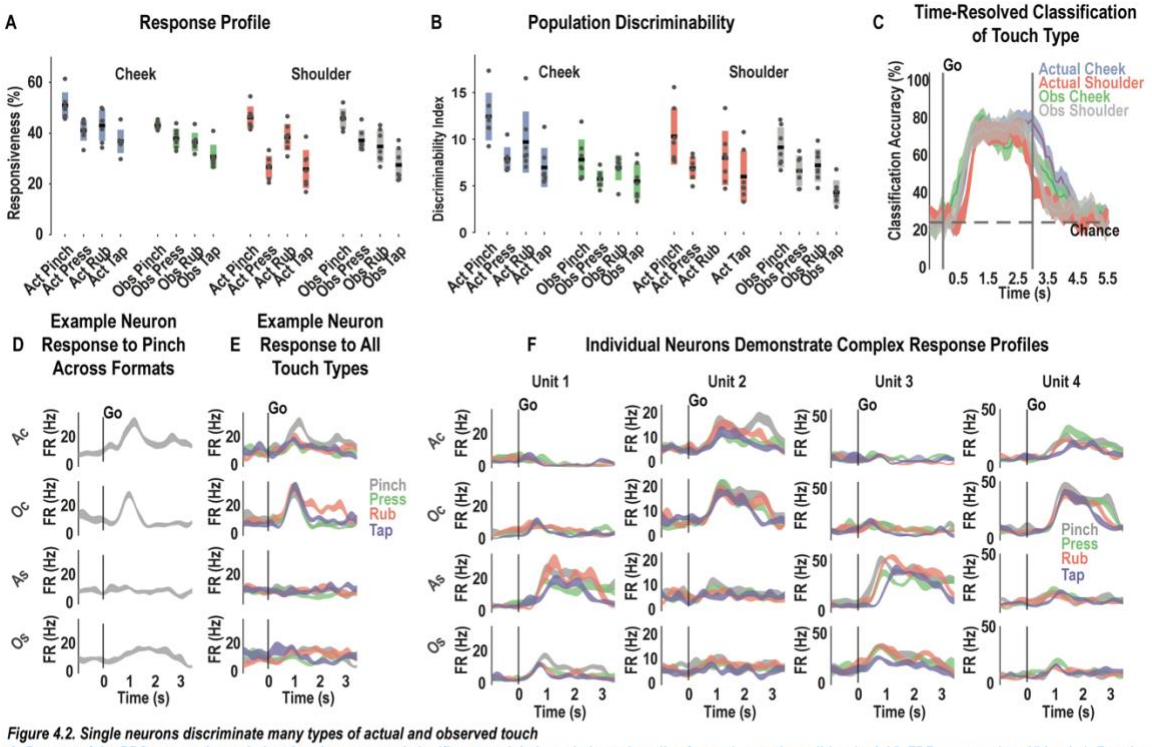

Figure 4.2. Single neurons discriminate many types of actual and observed touch

A, Percent of the PPC neuronal population that demonstrated significant modulation relative to baseline for each tested condition ($p<0.05$, FDR corrected, n=654 units). Results are shown as the mean percentage (horizontal black line) of the population ± bootstrapped 95% CI (bar height). Gray dots represent single session results. B, Discriminability of neural population activity (y-axis) during each of the test conditions (x-axis), from baseline neural activity. Results are shown as the mean value (averaged across sessions; horizontal black line) ± bootstrapped 95% CI. As in panel A, gray dots represent single session results. C, Time-resolved, cross-validated classification accuracy (y-axis) for discriminating the four touch types within each format (actual cheek, actual shoulder, observed cheek and observed shoulder touch). Results are shown as the mean across sessions (solid line, color coded) ± 95% CI (color coded) computed across sessions, plotted against time (x-axis). Chance accuracy is shown by the gray dashed line. The start of the stimulus phase is shown as the vertical line labeled 'Go', and the end of the stimulus phase is shown as an unlabeled vertical line at 3.0 seconds. D, Sample neuron response to pinch across all formats showing similar response to actual and observed cheek touches, with minimal response to shoulder touch. Each row depicts the response to one format as the mean firing rate ± SEM, n=10 trials. E, Response for the same example neuron from panel D across the four formats, now to all touch types. Touch types are color coded, as indicated. Other details are as in panel D. F, Example neurons illustrating often overlapping, yet discriminable responses to multiple touch types and formats (see also Figure 4.2-figure supplement 1). Each column depicts the response for one unit to each format (rows). Details are otherwise as in panels D and E. Act, actual; Obs, observed; s, seconds; Ac, actual cheek; Oc, observed cheek; As, actual shoulder; Os, observed shoulder; Hz, hertz

4.3.2.1 Single neuron responses are highly heterogenous. As in the first experiment, we found robust coding of ***experienced*** and ***observed*** tactile sensations (population Mahalanobis distance from baseline with 95% CI, **Figure 4.2A**; percent of cells modulated from baseline, t-test $p<0.05$ corrected for multiple comparisons, **Figure 4.2B**). The response to the different touch types could be discriminated for ***experienced*** or ***observed*** conditions (time-resolved classification, **Figure 4.2C**). However, the inclusion of additional touch types highlighted the near universal complexity of single unit responses that were difficult to reconcile with simple mirror neuron accounts. For example, **Figure 4.2D** shows the response properties of a single cell when NS ***felt*** a gentle pinch to her cheek or shoulder as well as when she ***observed*** another person being pinched on their cheek or their shoulder. This cell's response is consistent with the idea that NS knows what another person is experiencing because her own tactile cells are being recruited in a similar manner. However, testing the same neuron with additional touch-types reveals a more complicated pattern (**Figure 4.2E**), selective for pinch to her own cheek but responding to all touch-types during ***observation***. A straightforward interpretation of the mirror mechanism would predict that NS would understand all touch-types as a pinch, inconsistent with behavioral evidence that the touch types were easily discriminated and the finding that ***observed*** touch types are

discriminable (e.g. **Figure 4.2C**). Several additional example neurons illustrating the lack of specificity to touch-type are shown in **Figure 4.2E** and **Figure 4.2-figure supplement 1**.

We turned to a model selection analysis to summarize the behavior of single neurons across the entire population (**Figure 4.3**). The model selection analysis allows us to categorize neurons based on patterns of congruency in their responses to the four touch-types across all sensory fields, where sensory fields are defined as the location the touch was applied (i.e., cheek or shoulder, on NS or the experimenter). We began by constructing, for each neuron, a linear tuning model to describe its selectivity pattern (SP; the firing rate values for each of the four touch types) within each sensory field. We then use a model selection procedure to categorize whether these linear models demonstrated congruency (matched SP) across sensory fields. There are 51 possible models, each capturing one realization of SP to each sensory field and whether the responses are congruent (matched SP) or incongruent (or mismatched SP) across sensory fields. (**Figure 4.3-figure supplement 1**). For interpretative purposes, these 51 models were grouped into 7 categories (as labeled, **Figure 4.3-figure supplement 1)**. A few examples are illustrated schematically **Figure 4.3A-C** and in **Figure 4.3-figure supplement 2** (see figure legend). We computed all 51 possible models for every neuron. From amongst these possibilities, we identified the linear model that best described neural behavior using two metrics: Bayesian information criterion (BIC) and coefficient of determination (R^2). The percentage of PPC cells that behaved according to each model was summarized as the average percentage provided by these two measures. The breakdown of single unit activity according to all 51 models is shown in **Figure 4.3-figure supplement 3** (and split by BIC and R^2 in **Figure 4.3-figure supplement 4**). The breakdown according to the 7 categories is shown in **Figure 4.3D** (and split by BIC and R^2 in **Figure 4.3-supplement 5**). The PPC population was heterogeneous, composed of many complex patterns of congruency across sensory fields.

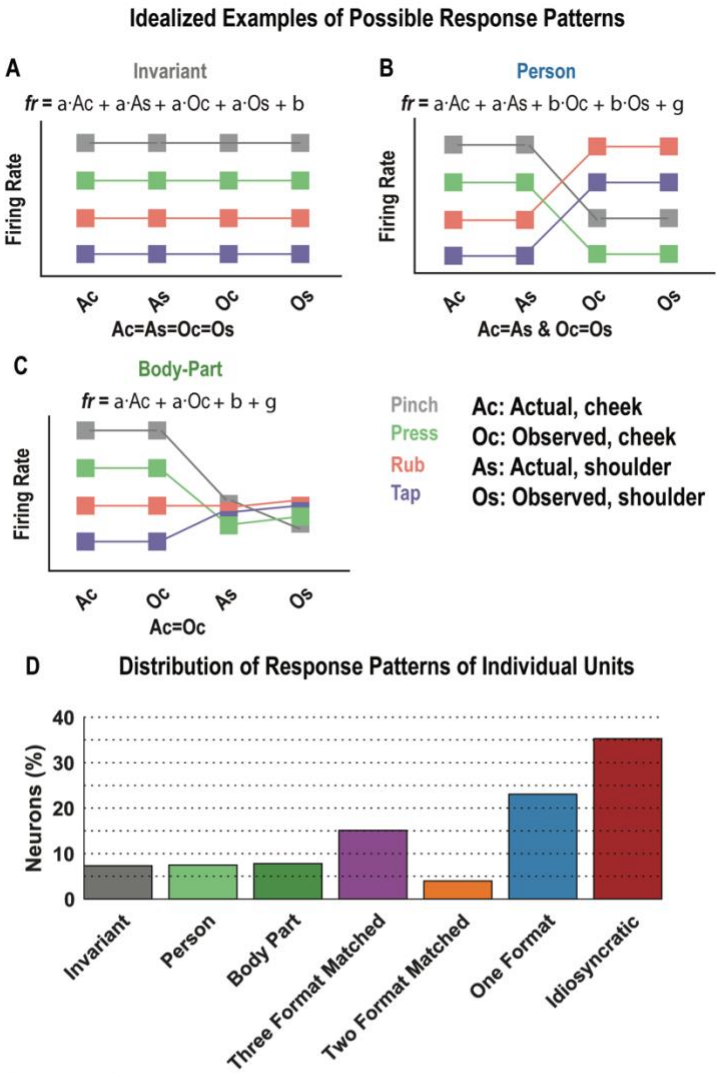

Figure 4.3. Single neurons are complex and heterogeneous

A-C, Schematic illustrations of three of the 51 possible linear models describing how a neuron's firing rate response (selectivity pattern, SP) to the four touch types (color coded), within each of the four formats (Ac, As, Oc, Os, see legend), can link together. The SP may be matched across all formats (panel A), matched within person (i.e., for actual cheek and actual shoulder touch) but mismatched during observed touch (panel B), matched during touch to cheek whether actual or observed but invariant to shoulder touch (panel C). These three example response patterns illustrate behavior that we categorized as invariance, person-specific, or body-part specific. The full list of categories of models is shown in Figure 4.3-figure supplement 1, and schematic examples of each category in Figure 4.3-figure supplement 2. The mathematical description of the model is shown above each illustration. Within each panel, a bold Greek letter (with a dot product with the format name) indicates the vector of coefficient values describing firing rate as a function of touch type within the corresponding format. A non-bolded Greek letter denotes a constant, or no selectivity to touch types for the corresponding format. Use of the same Greek letter across multiple formats indicates identical coefficient values, or matched SPs across formats. Use of different Greek letters indicates distinct coefficient values, and thus mismatched SPs across formats. An abbreviated name for the model is shown below each panel. The '=' symbol denotes congruency in SP between the format listed before and after the symbol, and the '&' symbol denotes mismatched SPs across the bookend formats of the symbol. D, Histogram showing the percentage of PPC neurons that behaved according to each category of models (see Figure 4.3-figure supplement 1). Ac, actual cheek touch; As, actual shoulder touch; Oc, observed cheek touch, Os, observed shoulder touch; fr, firing rate

4.3.2.2 Population-level neural subspaces mediate generalizability of information.
The mirror-mechanism is proposed to link what we see with what we intend or feel. Based on our previous data within this PPC substrate, we hypothesized that the mirror mechanism is a manifestation of a broader computational strategy by which shared, basic features within populations of PPC neurons generalize across task dimensions (behavioral contexts). To test this hypothesis, we train a linear model to discriminate the PPC population response for values along one dimension for fixed values of the other two dimensions. Mathematically, this model identifies a population-level neural *linear subspace* (or simple, subspace) that discriminates between values of the chosen dimension for fixed values of the other dimensions. Then, we test whether this subspace allows similar discrimination for alternate values of the fixed dimensions. For example, we train the model to discriminate touch types (dimension 1, pinch versus press), for fixed body part (dimension 2, e.g., cheek) and person (dimension 3, e.g., NS, ***experienced*** touch). Then we test the ability of this model to similarly discriminate touch type when switching body part and/or person. This analysis provides a simple way of quantifying whether information-rich (e.g., enabling discrimination) neural population codes are shared across task dimensions. If the mirror-mechanism is the dominant motif that determines population-level encoding, then we would predict preferential sharing or generalization when matching body locations between self and other (e.g., between cheek and cheek or between shoulder and shoulder). Otherwise, under a more general mechanism for shared encoding, we would expect generalizable information to be a ubiquitous phenomenon.

The analysis is schematically illustrated in **Figure 4.4A**. Here, a subspace is formed that optimally distinguishes ***actual*** pinches to the shoulder from ***actual*** pinches to the cheek. This example subspace thus answers the question: where on my body is being pinched? We can than ask how well this decision space generalizes e.g., when the person being touched is not NS but the experimenter, such that pinches to the cheek and shoulder are ***observed*** instead of ***experienced***, or to different touch types, such as ***actual*** presses to the cheek or shoulder instead of pinches. In the illustrative example of **Figure 4.4A**, the decision subspace also allows distinguishing the body part to which pinch is applied, if it were ***observed*** instead of ***experienced***.

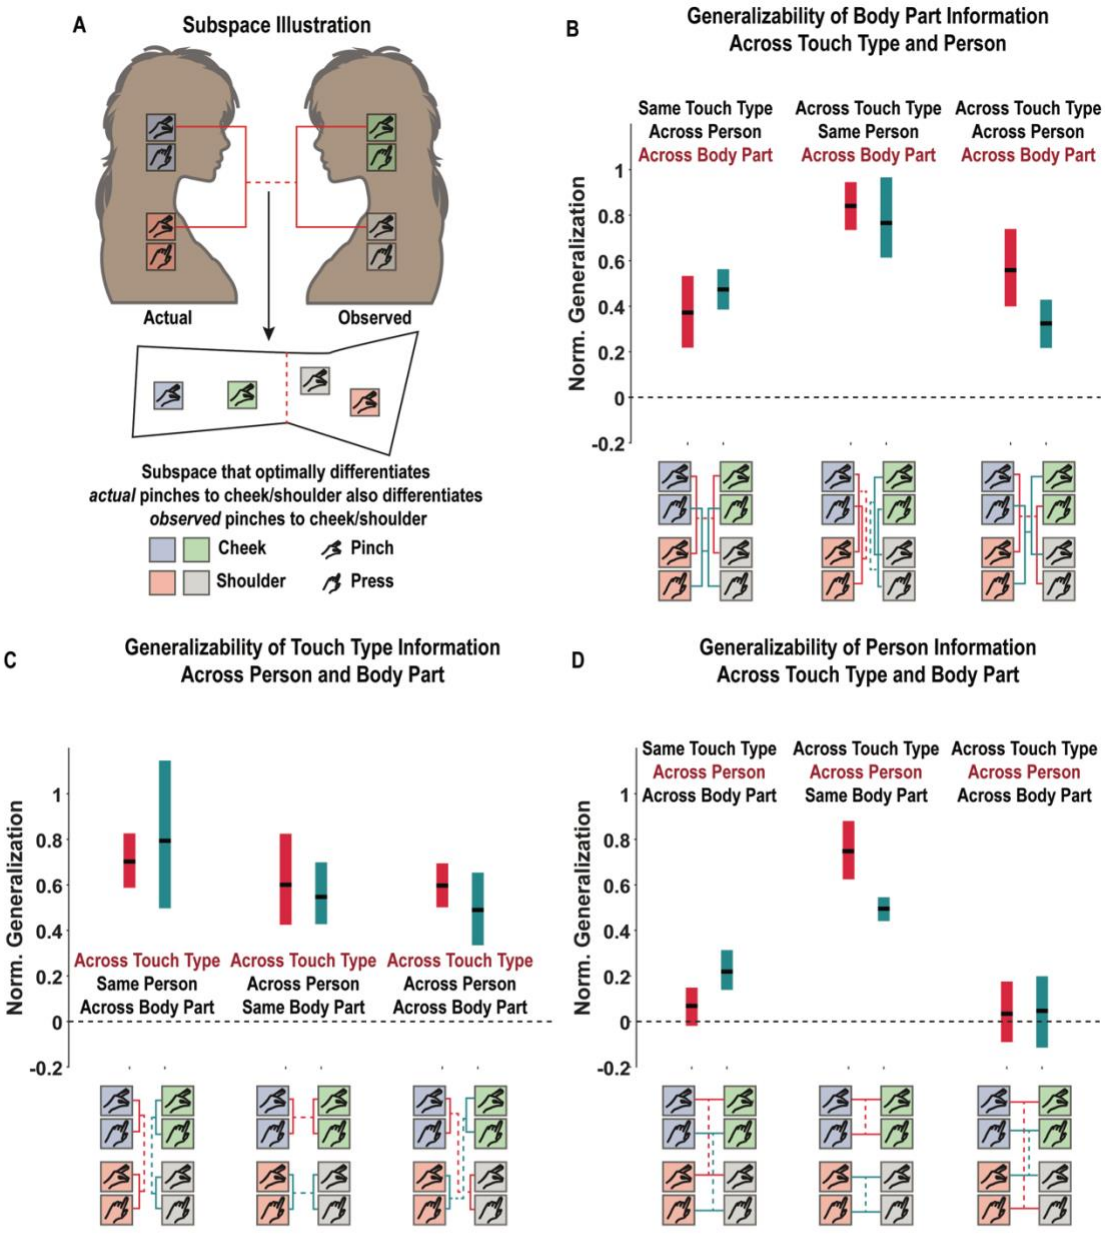

Figure 4.4. Population-level subspaces mediate generalizability of information.

A, Schematic illustration of subspace analysis. The two body parts, two touch types used in the analysis are schematically illustrated by the colored squares, and icons (see legend). The two torsos illustrate actual touch to the subject (NS) or observed touch to an experimenter (labeled). In this illustration, a subspace is constructed (the polygon) that optimally differentiates (shown by the red dashed line, the decision boundary) body parts (cheek, shoulder) during actual pinches. This subspace is shown by on the torsos by the red square bracket facing left. Next, this subspace is tested on whether it differentiates the same body parts during observed pinches. The test set is shown by the red square bracket oriented to the right. The result is shown schematically in the same polygon (subspace) below the torsos. Cheek and shoulder are on opposite sides of the boundary line during both actual pinches and observed pinches, indicating that this subspace (that carries body part information) generalizes across persons (i.e., actual or observed touch conditions). B, Results of the subspace analysis when testing how body part information generalizes across the other two dimensions: touch type and person. The normalized generalization (y-axis) is shown for each tested subspace (x-axis). On the x-axis below each group of bars is a condensed schematic (from panel A), showing the two body parts, the two touch types and two persons. The red and green lines illustrate two separate tests for generalizability. Note that within the red or green sets of lines, the test could be computed with the training pair for the subspace on the left or on the right, and the test pair similarly on the left or on the right. The bars show the mean generalization (horizontal black line) ± 95% CI computed across sessions. The mean is the average generalization not only across sessions but also in both directions (trained on pair on left, tested on right, and vice versa). Each group of bars tests the generalizability for information related to body part for the same or different values of the other two dimensions, as labeled above the group. The dimension whose generalizability is being tested (here, body part) is labeled in red, while the test dimensions are in black. A lack of generalization is shown, for reference, as the dashed line centered at 0. C, Similar to panel B, except here the generalizability of touch type information is being tested across the other two dimensions (body part, and person). D, Similar to panels B and C, except here the generalizability of person information is being tested across the other two dimensions (body part and touch type). Norm, normalized.

The results of this analysis are shown in **Figures 4.4B-D**. We constructed subspaces for all main pairs of variables and tested for generalization. In Figure **4B** we show the resulting patterns of generalization when discrimination subspaces are built around the two body parts: cheek and shoulder. The subspace generalizes to allow body parts to be discriminated, whether touch is applied to the same (or the other) person, and whether it is of the same type (e.g., pinch versus press) or the other touch type. Similarly in **Figure 4.4C**, we show that discriminable patterns built around coding of touch type (e.g., discriminating pinch versus press) generalize to allow discrimination whether for the same person or the other, and within and across body parts. Interestingly, **Figure 4.4D** shows a fundamentally different pattern, with generalization in some contexts but not all. In other words, touch type and body location are encoded in ways that generalized across contexts while the person did not. Such preferential encoding fits with the known functions of our cortical implant site (see discussion).

To assist interpretation, we repeat the analysis with a population of synthetically generated mirror neurons (**Figure 4.4-figure supplement 1**). As expected, the analysis shows that a population composed exclusively of mirror-like neurons only showing similar responses for matched conditions across self and other does not support ubiquitous generalization, unlike our PPC population.

4.3.2.3 Architecture of knowledge representation in human PPC: structured compositionality. The mirror-mechanism has been suggested to link visual inputs to single neurons encoding high-level action goals. We hypothesize that ***experienced*** and ***observed*** sensations are represented through a combination of simple primitive elemental features encoded by the population. We used a demixed-principal components analysis (dPCA) to visualize and quantify how much of the population-level response could be understood as being composed of a linear combination of these elemental features (such as information about body part or touch type) encoded by the population.[137] The dPCA analysis decomposed the pooled population response across all test conditions into independent factors (here, body part, touch type and person), interaction effects (body part-person, body part-touch type, and person-touch type), and a term accounting for a shared neural response to any form of touch.

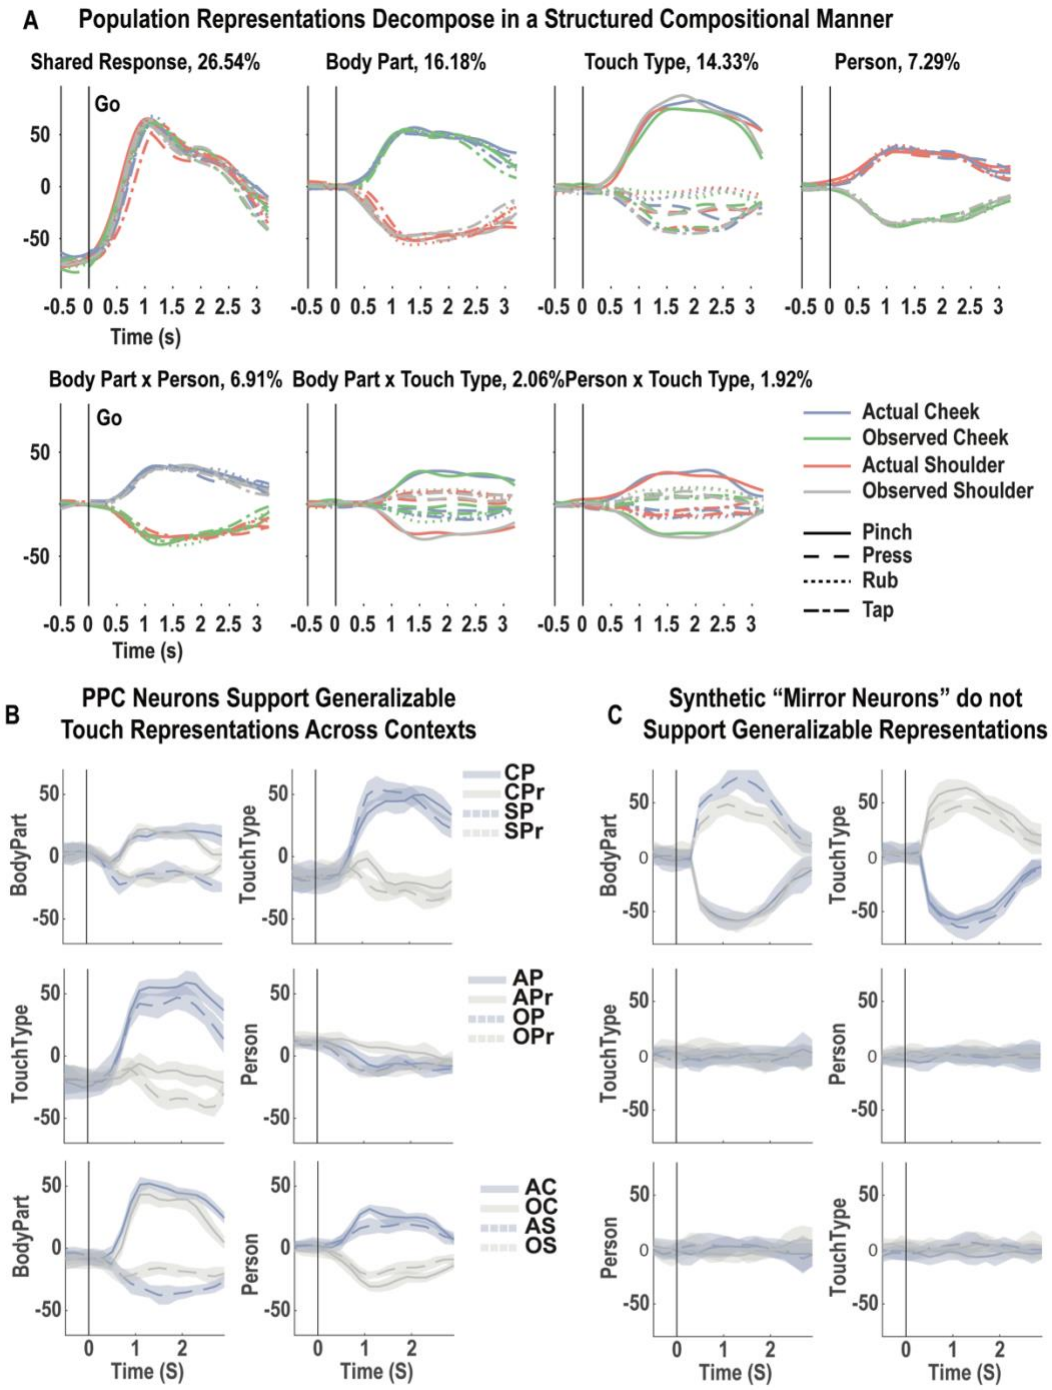

Figure 4.5. Structured compositional architecture of PPC responses to touch

A, Results of a dPCA performed on the pooled population data are shown. Each panel shows the first component along a marginalization (as labeled above the panel), along with the portion of population variance that it explains (as a percentage). In each panel, the projection of test condition is shown as a function of time (x-axis), along a dimensionless y-axis, coded by color and line type as shown on the right. B, Each row shows the results of a dPCA in which the population representations of four conditions (in the legend) were decomposed along two dimensions (labeled y-axis; each panel) for a fixed value of the third dimension. The learned mapping was then tested on the held-out value of the third dimension. The mean value across sessions (and both directions of the third dimension) ± 95% CI for each condition is shown, as a function of time (x-axis). C, Same analysis as in B, except performed on synthetic mirror neurons. CP, cheek pinch; CPr, cheek press; SP, shoulder pinch; SPr, shoulder press; AP, actual pinch; APr, actual press; OP, observed pinch; OPr, observed press; AC, actual cheek; OC, observed cheek; AS, actual shoulder; OS, observed shoulder; s, seconds.

The results are shown in **Figure 4.5A**. A significant fraction of population behavior, described here as the percentage of the population variance described by each component, relates to the shared neural response to any form of touch. Otherwise, the largest variance is explained by the three main effects with smaller amounts relating to the interactions. This breakdown suggests that most of the population's response to various forms of touch is comprised of information deriving from independent dimensions of body part, touch type, and person. However, amongst these three dimensions, body part and touch type contribute more to the population response than the third (person) dimension. This asymmetry suggests that within this population of PPC neurons, the compositionality may be asymmetric or structured.

To ensure that dPCA is capturing generalizable features of the neural population code, we performed a generalization test of the dPCA components. To this end, we performed a two-factor dPCA on two dimensions of the data for a single level of the third dimension. We then applied the learned mapping to the held-out level of the third dimension. The results are shown in **Figure 4.5B** for each possible combination of the two-factor dPCA. As expected, we find that main components of body part and touch type generalize, consistent with the hypothesis that PPC populations encodes generalizable basic components of sensations. Finally, we repeated this dPCA analysis for our synthetic mirror neurons. As shown in **Figure 4.5C**, canonical mirror neurons do not form a representational basis.

4.4 DISCUSSION

4.4.1 Mirror neurons reconsidered. Mirror neurons are an influential theory for how we understand the actions and experiences of others.[12,24,25,28] At a broad level, the mirror hypothesis makes the claim that neurons within high-level regions responsible for planning our own motor behavior or bodily sensations are also involved in understanding the intentions and experiences of others.[25] Our data are consistent with this view as we show that neurons active during ***experienced*** sensations also become active when ***observing*** somebody else being touched in similar ways. The mirror hypothesis also proposes a "mirror mechanism" whereby understanding is achieved through a process of simulation.[7,27] This has been primarily described in the motor domain where most mirror studies are carried out. In this view, individual neurons in higher-order motor cortices encode action goals. Different single neurons encode different goals, resulting in a full vocabulary.[7,27] These neurons are assumed to naturally be imbued with meaning under the assumption that we must understand our own actions. The mirror mechanism, by extending access to this vocabulary to what we ***observe***, can therefore directly impart understanding.[7,24,25,27] It is here that our interpretation of mirror-like responses diverges from the classical model.

Our data demonstrate that shared responses associated with mirroring are part of a more general mechanism in which shared representations support cognition. A shared representation of pinch, for example, informs not only ***felt*** cheek pinches, but also ***felt*** shoulder pinches, as well as ***observed*** pinches. A mirror mechanism implemented through

simulation loses intuitive appeal in clarifying such shared population activation, such as during rubbing motions whether ***experienced*** on one's cheek or ***observed*** on another's shoulder. Are we here understanding how a rub to the shoulder feels by simulating it in our cheek-sensitive neurons? Likewise, if the key operational feature is similar patterns of activation, the same logic would lead one to predict that we understand how a tapping motion feels by simulating it in our pinch neurons or vice versa. In our view, the shared representations reported in this, and in our recently published work, are the neural architecture that allows us to efficiently learn and generalize our experience more broadly. This basic view can connect to a recently developed theory of human-like cognition.

4.4.2 Potential basis for cognitive models of the world. A recent branch of cognitive neuroscience has put forth the idea that human-like learning and thinking is largely built on the internal models we construct of the world.[6,14] These models are thought to play a ubiquitous role, providing a common substrate to inform our perception, cognition, imagination, and planning. Yet, to date, the neural basis for this cognitive neuroscientific framework remains largely unexplored, though preliminary neuroimaging evidence points to a role for PPC [14]. Our results, demonstrating shared representations that can be decomposed into basic building blocks, support this computational architecture, and provide preliminary insight into its neural implementation within human PPC. This framework provides a unifying account of many of our recent results, suggesting that language, imagination, planning, and perception tap into the same underlying shared internal models.[2-4,23] Understanding our neural results in this framework also helps to address limitations associated with mirror neurons, as discussed below.

4.4.3 Generalization. By the mirror account, single neurons encode our own action goals, and the mirror mechanism extends access to these representations to ***observation***.[136] As stated above, such a mechanism does not allow for understanding action goals outside our own repertoire. In our account, high-level regions of the cortex encode representational building blocks that can be combined in novel ways to understand novel stimuli. Future studies can directly test this idea: For example, we would predict that if were to ask a participant to ***imagine*** what it would feel like if we pinched her tail (something clearly outside the participant's direct experience) we would find that the neural subspace associated with pinches would combine with cortical representations associated with tails, presumably built from observation of animals.

4.4.4 A selfless theory. The mirror account assumes a temporal dependence: individuals form motor representations through producing actions that can subsequently be accessed during ***observation***.[24] In our understanding, such temporal precedence is not necessary. Individuals can form representational building blocks using any possible information source. To use an example from Patricia Churchland, even if I have never myself milked a cow, I can readily build understanding by watching someone else milk a cow which in turn, can then help me quickly understand milking a goat or help inform my own attempts to milk a cow.[138] No doubt our own actions and experiences provide unique and unreproducible forms of knowledge. Just as the description of a sunset cannot replace witnessing a sunset, it is

likely that ***observing*** pinches cannot replace the ***experience*** of having been pinched. Thus, our own experiences can change the nature of our internal representations but are not a prerequisite to form these representations. Indeed, we predict that the population of neurons that have mirror-like sensory responses will also show encoding for observed contact between inanimate objects as they share an underlying basis in the physics of interacting objects.

4.4.5 Relationship to alternative accounts of mirror neurons. Considering arguments against the mirror hypothesis, a number of groups have hypothesized that visual information related to the actions of others arises in our motor system simply as a means to guide motor behavior e.g., by mediating motor imitation, observational learning, or planning in response to the actions of others.[136] These are compelling accounts given that animals clearly use observation of the actions of others to guide their own motor behavior. It is less clear how well such explanations can account for our data: There is no simple corollary of generating an endogenous sensory experience in response to the sensory experiences of others and thus these alternative accounts are less persuasive in the sensory domain. In our view, the compositional building blocks provide useful representations that can inform all relevant aspects of behavior. In the motor domain, this can include understanding the actions of others as well as guiding our own motor behavior based on the actions of others. Interestingly, the degree of population-level similarity between executed and observed actions appears to be smaller than we report here, in the sensory domain.[136] While there are many possibilities, one intriguing hypothesis is that the number of possible ways observed actions can inform our own cognition and behavior (e.g., understanding, imitation, learning, motor planning) is substantially larger than the sensory case and thus leads to more multi-faceted neural responses.

4.4.6 Relevance to BCI. Numerous clinical trials have shown that individuals with paralysis can use signals from motor regions of the brain to control external devices, such as robotic limbs or computer cursors.[3,45,50] The underlying brain signals are low-dimensional and roughly encode movement direction in a smooth continuous fashion enabling researchers to collect sufficient data to train a decoding algorithm in a few minutes time. Future implementations of BCIs could communicate the contents of the mind more directly, such as direct decoding of high-level concepts, visual imagery, or emotional state. The dimensionality of these signals is far larger than basic movements. However, if these high-dimensional datasets are encoded using generalizable relatively low-dimensional basis sets, then ability to read out these high-dimensional signals may be tractable. To this end, proof-of-concept studies have already demonstrated the ability to decode high-fidelity faces or the semantic content of visual scenes from rich low-dimensional basis sets.[139]

4.5 CONCLUSION

Understanding how systems of neurons transform data into understanding is central to understanding human cognition and how we can create smarter machines. Our data demonstrate that mirror neurons are one manifestation of a more general mechanism by

which neural populations encode generalizable representational building blocks. Recent cognitive science literature has demonstrated that such building blocks are a key feature that enable human-like learning and thinking.

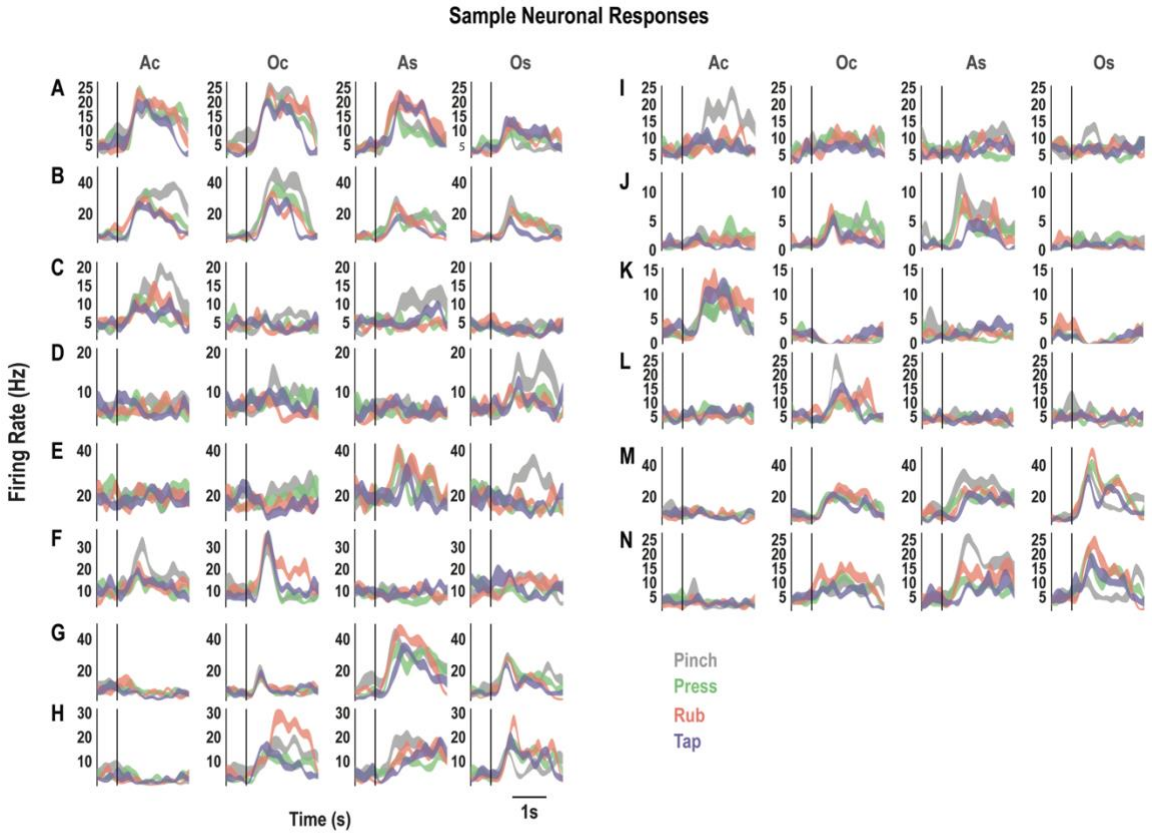

Figure 4.2-figure supplement 1. Individual neurons exhibit complex and variable response patterns.

A-N, Each row (labeled by an alphabet) depicts the response of one example neuron to the four different formats (columns; Ac, actual cheek touch; Oc, observed cheek touch; As, actual shoulder touch; Os, observed shoulder touch). Within each panel the neural response to each of the four touch types is shown, color coded, as the mean firing rate (y-axis) ± SEM, n=10 trials, as a function of time (x-axis). Ac, actual cheek; Oc, observed cheek; As, actual shoulder; Os, observed shoulder; Hz, hertz; s, seconds.

Invariant	Person	Body-Part	Three Format	Two Format	One Format	Idiosyncratic
Ac=As=Oc=Os	Ac=As&Oc=Os	Ac=Oc&As=Os	Ac=As=Oc	Ac=Os&As=Oc	Ac	As&Oc&Os
	Ac=As	Ac=Oc	Ac=Oc=Os	As=Oc	As	Ac&Oc&Os
	Oc=Os	As=Os	Ac=As=Os	Ac=Os	Os	Ac&As&Os
	Ac=As&Oc	Ac=Oc&As	As=Oc=Os	As=Oc&Os	Os	Ac&As&Oc
	Ac=As&Os	Ac=Oc&Os	Ac=As=Oc&Os	As=Oc&Ac		Ac&As&Oc&Os
	Oc=Os&Ac	As=Os&Ac	Ac=Oc=Os&As	Ac=Os&As		As&Os
	Oc=Os&Ac&As	As=Os&Oc	Ac=As=Os&Oc	Ac=Os&Oc		Oc&Os
	Ac=As&Oc&Os	As=Os&Ac&Oc	As=Oc=Os&Ac	Oc=Os&As		Ac&As
		Ac=Oc&As&Os		Ac=Os&As&Oc		Ac&Oc
				As=Oc&Ac&Os		As&Oc
						Ac&Os

Ac: Actual, cheek
Oc: Observed, cheek
As: Actual, shoulder
Os: Observed, shoulder

=: Matched selectivity pattern
&: Mismatched selectivity pattern

Figure 4.3-figure supplement 1. Complete list of all possible linear models to describe congruency in selectivity patterns across formats.
All 51 possible linear models that could describe an individual neuron's response (selectivity pattern) to the four different touch types), and the congruency of that SP across formats (Ac, As, Oc, and Os) are listed by their mathematical description. The '=' symbol denotes congruency in SP between the format listed before and after the symbol, and the '&' symbol denotes mismatched SPs across the bookend formats of the symbol. The 51 possible models are grouped into 7 categories for interpretative purposes, labeled accordingly. Ac, actual cheek touch; As, actual shoulder touch; Oc, observed cheek touch; Os, observed shoulder touch.

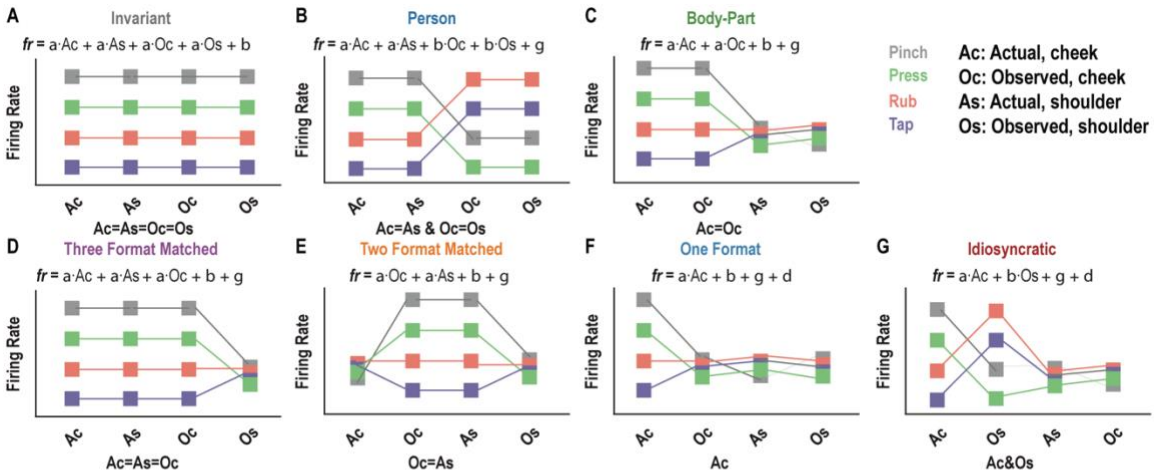

Figure 4.3-figure supplement 2. Schematic illustration of each category of linear models to describe congruency in selectivity patterns across formats.
A-G, Schematic illustrations of each category of the 51 possible linear models describing how a neuron's firing rate response (selectivity pattern, SP) to the four touch types, within each of the four formats (Ac, As, Oc, Os, see legend), can link together. Details of each panel are as in the description for Figure 4.3A-C. Ac, actual cheek touch; As, actual shoulder touch; Oc, observed cheek touch, Os, observed shoulder touch; fr, firing rate

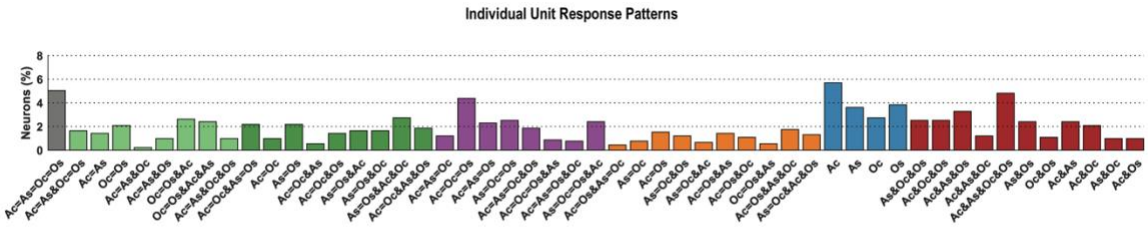

Figure 4.3-figure supplement 3. Breakdown of single unit responses according to all possible response configurations

Histogram showing the percentage of PPC neurons that behaved according to each of the 51 possible linear models. The abbreviated form of the model is listed below each bar. As in Figure 4.3-figure supplement 1, the '=' symbol denotes congruency in SP between the format listed before and after the symbol, and the '&' symbol denotes mismatched SPs across the bookend formats of the symbol. Bars are color coded according to the category they belong within (see Figure 4.3-figure supplement 1). Ac, actual cheek touch; As, actual shoulder touch; Oc, observed cheek touch, Os, observed shoulder touch

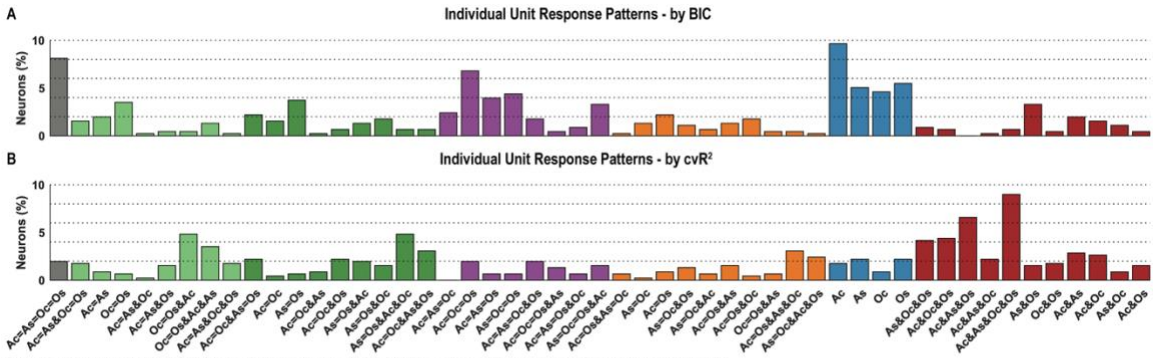

Figure 4.3-figure supplement 4. Breakdown of single unit responses according to all possible response configurations, split by metric

A, Histogram showing the percentage of PPC neurons that behaved according to each of the 51 possible linear models, as in Figure 4.3-figure supplement 3. Here, the breakdown is shown based on using Bayesian Information Criterion (BIC) as the metric for evaluating which of the 51 models was the best-fit for each neuron's behavior. The abbreviated form of the model is listed below each bar. As in Figure 4.3-figure supplement 1, the '=' symbol denotes congruency in SP between the format listed before and after the symbol, and the '&' symbol denotes mismatched SPs across the bookend formats of the symbol. Bars are color coded according to the category they belong within (see Figure 4.3-figure supplement 1). B, As in Panel A, except here the breakdown is shown based on using cross-validated R2 (coefficient of determination) of the model as the metric for evaluating the best-fit for each neuron's behavior. BIC, Bayesian information criterion; cvR2, cross-validated coefficient of determination; Ac, actual cheek touch; As, actual shoulder touch; Oc, observed cheek touch, Os, observed shoulder touch

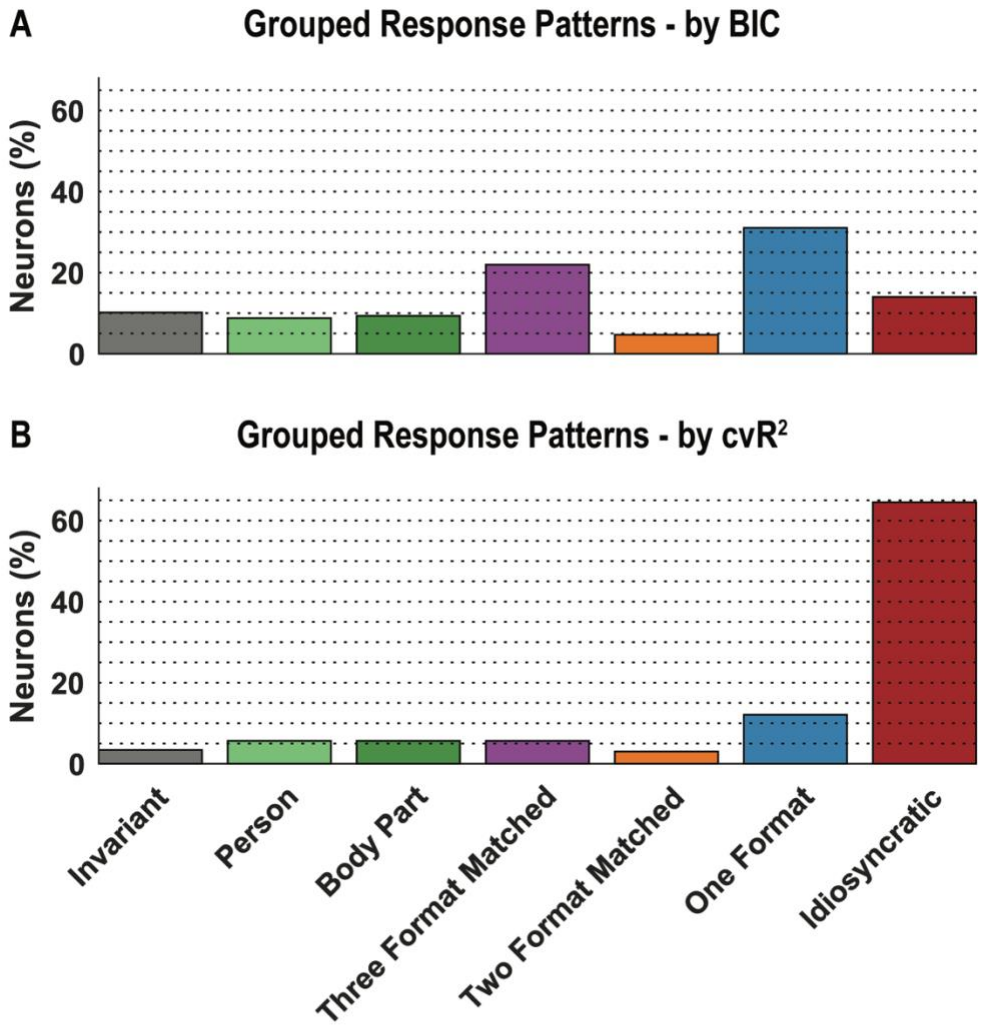

Figure 4.3-figure supplement 5. Breakdown of single unit responses within categories, split by metric

A, A, Histogram showing the percentage of PPC neurons that behaved according to each of 7 categories of linear models (see Figure 4.3-figure supplement 1). Here, the breakdown is shown based on using Bayesian Information Criterion (BIC) as the metric for evaluating which of the 51 models was the best-fit for each neuron's behavior. The category name is listed below each bar. Bars are color coded according to the category they belong within (see Figure 4.3-figure supplement 1). B, As in Panel A, except here the breakdown is shown based on using cross-validated R2 (coefficient of determination) of the model as the metric for evaluating the best-fit for each neuron's behavior. BIC, Bayesian information criterion; cvR2, cross-validated coefficient of determination

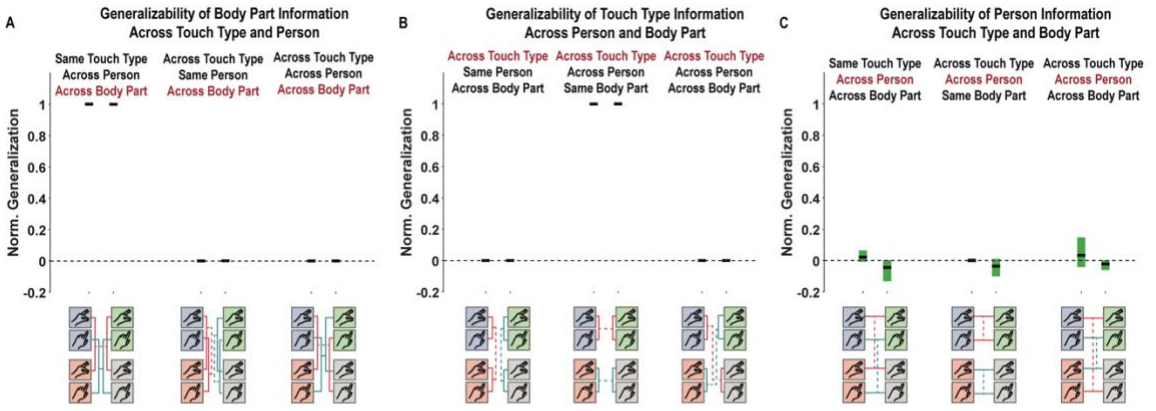

Figure 4.4-figure supplement 1. Synthetically generated mirror neurons do not demonstrate the same generalizability as the PPC neural population.

A, Results of the subspace analysis performed on synthetically generated mirror neurons, when testing how body part information generalizes across the other two dimensions: touch type and person. The normalized generalization (y-axis) is shown for each tested subspace (x-axis). On the x-axis below each group of bars is a condensed schematic (from as in Figure 4.4A), showing the two body parts, the two touch types and two persons. Details are as in Figure 4.4. B, Similar to panel A, except here the generalizability of touch type information is being tested across the other two dimensions (body part, and person). C, Similar to panels A and B, except here the generalizability of person information is being tested across the other two dimensions (body part and touch type). Norm, normalized.

Chapter 5

A Vision of Touch in Human Posterior Parietal Cortex

5.1 INTRODUCTION

Touch is a multisensory perceptual process.[51,52] When an ant crawls up our hand, we often both ***feel*** it and we ***see*** it. In other words, the touch we perceive is constructed from both visual and tactile information.[30-33,82] But what does it mean to ***see*** touch? Are our brains responding to motion on (or near) our body?[30-33,82] To the distance (or lack thereof) between the object and our body? In another setting, the ant may not even be on our hand but on a blade of grass.[36-39] Do our brains represent ***observed*** touch to objects (like the grass) in the same way as to our body? The capacity for our brains to process ***seen*** touch to objects (including ourselves) independent of ***felt*** touch is an essential, yet poorly understood, element of touch perception. More broadly, because touch is ubiquitous in our daily lives, understanding how and where our brains process ***seen*** touch is critical to understanding how we perceive the world in general.

There is considerable evidence that posterior parietal cortex (PPC) is involved. In non-human primates (NHPs), PPC neurons are modulated by visual information: they respond to touch to a monkey differently depending on whether the touch is only ***felt*** or also ***seen***.[140-143] Moreover, some PPC neurons have been identified that respond congruently not only when a monkey ***feels*** the touch but also when it ***sees*** touch to another monkey on the corresponding body part.[140] Such neurons have typically been interpreted as *mirror neurons* and taken as evidence of *mirroring*: using one's own body as a reference for perceiving touch to others.[36-39,57,144] However, this view is restrictive. Human neuroimaging studies indicate that not only ***seen*** touch to other individuals, but also to inanimate objects, activate overlapping regions of PPC.[36,57] While these studies cannot clarify the neuron-level mechanisms of such shared activation, they highlight that PPC processing of ***seen*** touch (and touch in general) is richer than mirror neurons alone can support.

Our lab's recent work has provided fresh insight into the computational basis for behavior in human PPC. In an ongoing brain-machine interface (BMI) clinical trial, we recorded from populations of single neurons in human PPC of a tetraplegic participant and found that they are simultaneously engaged during many forms motor, sensory, and cognitive behaviors.[2-4] This PPC population encodes variables such as ***observed*** actions, ***observed*** touch, motor and tactile ***imagery***, ***felt*** touch, and others.[2-4] Individual neurons are highly variable, including a small fraction that behave like mirror neurons. The population, however, forms rich associations across neurons and domains, representing basic-level, primitive elements that can be repurposed for many different behaviors, in a functional organization that we termed partially mixed selectivity.[4,43] For example, within the sensory space, we found that this population of heterogeneous PPC neurons encodes ***felt*** touch, at short latency, through primitive elements such as which body part (e.g., cheek, shoulder) is touched and in what

form (e.g., pinch, press).[5] These elements are *compositional,* combining in myriad ways to inform diverse touch conditions (e.g., cheek pinch, shoulder press). They are also *generalizable*, able to explain novel (held out) data. Moreover, they generalize to other forms of touch processing as well, supporting not only ***felt*** touch, but also ***imagined*** touch to oneself, and ***seen*** touch to others.

Here we test a hypothesis that this PPC population provides a versatile substrate that can account for ***seen*** touch, not just to others but more universally, also to oneself (independent of tactile input) and to objects in the world. In a unique opportunity, we record single-unit activity in the same study participant as above (insensate below approximately the shoulder-level) during various forms of ***seen*** touch.[5] We find that this PPC population represents ***seen*** touch to all tested objects, including the participant's insensate arm regions. Touch is preferentially encoded over visually matched controls (such as motion near the object but without contact), and invariant to gaze. ***Seen*** touch to all regions of objects is encoded, but touch responses are sensitive to object attributes: similar object features evoke more similar PPC responses to touch. As in previous studies, individual neurons are complex.[3] Yet, meaningful information can be decoded from population representations such as the precise location along an object that is touched, and whether touch was ***seen*** as the stimulus moved toward the object or away from the object. Our current results are a novel, neuron-level understanding of ***seen*** touch, thus far untested in animal models or in humans. Taken together with our previous work in PPC, we speculate that this PPC substrate encodes an *internal model* for touch, a mental blueprint for touch and all its related aspects, that is flexibly engaged whether the touch is ***felt***, ***imagined***, or ***seen***.

5.2. METHODS

5.2.1 Subject details

All data were recorded from a 62-year-old tetraplegic female participant in a brain-machine interface (BMI) clinical trial. She has a high-cervical spinal cord injury (SCI) between cervical levels three and four, sustained approximately 10 years prior to the study, and with no preserved sensory or motor function below the shoulder-level. She was implanted with two 96-channel Neuroport Arrays (Blackrock microsystems model numbers 4382 and 4383) 6 years post-injury, in the left hemisphere. Informed consent was obtained, and she understood the nature, objectives, and potential risks, of the surgical procedure and the subsequent clinical studies. All procedures were approved by the Institutional Review Boards (IRBs) at the California Institute of Technology (IRB #18-0401), the University of California, Los Angeles (IRB #13-000576-AM-00027), and Casa Colina Hospital and Centers for Healthcare (IRB #00002372).

5.2.2 Experimental setup

All experiments were conducted at Casa Colina Hospital and Centers for Healthcare. NS was seated in a motorized wheelchair in a well-lit room. For all tasks, a wheel-chair

compatible drawing board was placed in front of the subject. Task details are presented below, but briefly, various objects (including the subject's arms) were placed on the drawing board such that touch to the objects was clearly visible to the subject at a convenient head position. A 27-inch LCD monitor was positioned behind NS (visible to the experimenters but not to NS) too cue the experimenter for the presentation of stimulus. Cue presentation was controlled by the psychophysics toolbox (Brainard, 1997) for MATLAB (MathWorks).[67]

5.2.3 Physiological recordings

NS was implanted with one Neuroport array at the junction of the intraparietal sulcus (IPS) and postcentral sulcus (PCS), a region we refer to as PC-IP. The other was in the left superior parietal lobule (SPL). Following surgery, the SPL implant did not function. Only data recorded from PC-IP were used in this study. Both arrays were explanted approximately one year after data in this study were collected.

Neural activity recorded from the array was amplified, digitized, and sampled at 30 kHz using a Neuroport neural signal processor. This system has received Food and Drug Administration (FDA) clearance for <30 days of recordings. We received an investigational device exemption (IDE) from the FDA (IDE #G120096, G120287) to extend the implant duration for the purposes of the BMI clinical study.

We have previously published our sorting algorithm.[4] Putative neuron action potentials were detected at threshold crossings of -3.5 times the root-mean-square of the high-pass filtered (250 Hz full bandwidth signal. Each waveform was made of 48 samples (1.6 ms), with 10 samples prior to triggering and 38 samples after. Single- and multi-unit activity was sorted using Guassian mixture modeling on the first three principal components of the detected waveforms. Well-isolated single and multi-units were pooled across recording sessions. To ensure that such pooling did not bias the conclusions of the paper, we performed core analyses on single-units alone, potential multi-units alone, and all units together. The results of these analyses, shown as supplemental figures for key results, and generally demonstrate that our results were robust to the pooling of all sorted units together.

5.2.4 Task procedures

5.2.4.1 Basic touch task ***(BTT; Figure 5.1; Figure 5.1-figure supplement 1, Figure 5.1-figure supplement 4, Figure 5.1-figure supplement 5, Figure 5.1-figure supplement 6, Figure 5.1-figure supplement 7).*** The purpose of this task was to explore the selectivity of PPC neurons to touch compared to touch-like controls. NS viewed stimuli (brush strokes at 1 Hz, for 2 seconds) to one of four objects (**Figure 5.1A**). These were the hand region (insensate) of the subject's right arm (**Native**), the same region of an experimenter's right arm (**Exp-Right**), a **banana** and a **sponge**. Stimuli included ***seen*** touch (**Touch**), touch-like motion overlapping, but not touching, the object (**Overlap**), motion adjacent (**Adjacent**), and motion far (**Far**) from the object. Objects were placed on a drawing board placed in front of

the subject, with only the tested object visible on any time. The stimulated site on each object consistently subtended a consistent angle on the eyes.

Each trial began with a cue to the experimenter indicating the object and stimulus, followed by a one second pause, then a two second stimulation phase, and ended with a one second intertrial interval. Eight blocks of pseudorandomly interleaved trials were conducted. All conditions were tested within each block, prior to the next block commencing. To test whether gaze fixation influenced the neural responses, two versions of this task were performed. In one, fixation was always at the Touch site (**Figure 5.1**), and in the other, it was at the **Adjacent** site (**Figure 1-figure supplement 4**). In total, we recorded from 382 neurons over six recording sessions.

5.2.4.2 Multiple object task ***(MOT; Figure 5.2, Figure 5.2-figure 2 supplement 1, Figure 5.1-figure supplement 2, Figure 5.1-figure supplement 3).*** We performed this task to examine the sensitivity of PPC touch representations to object attributes. Here, NS viewed touch stimuli (brush strokes at 1Hz for 2 seconds) presented to one of four objects: insensate regions of the subject's right arm (**Native)**, an experimenter's right arm (**Exp-Right**), the same experimenter's left arm (**Exp-Left**), and a **sponge**. Three fields were defined on each of the arms: **Field 1** on the hand, **Field 2** on the distal forearm, and **Field 3** on the proximal forearm. Three similarly separated fields were also defined on the sponge. Corresponding fields on all objects subtended a consistent angle on the eyes. The subject was instructed to fixate gaze at the field to which touch was being applied. The trial structures and durations were as described in the previous task. As earlier, only the tested object was visible at any time. In total, we recorded from 672 neurons over seven recording sessions.

Another version of this task was performed to test whether ***seen*** touch to inanimate objects other than the sponge was also represented. As in the previous paragraph, NS viewed touch to different fields on multiple objects. Here, 7 objects were tested to sample various household categories. These included: the subject's right arm (**Native**), **scissors**, **hammer**, **shoe**, **toy** alligator, **banana,** and **sponge**. Two fields were defined on each object, consistently separated. **Field 1** was away from NS; **Field 2** was closer to NS. As earlier, only the tested object was visible at any time. NS was instructed to always gaze at the tested touch site. In total, in this version, we recorded from 545 neurons over six recording sessions.

5.2.4.3 Spatial coverage task ***(SCT; Figure 5.3).*** This task was performed to evaluate the spatial coverage of neural response fields and their properties. A magnetic tracking probe (Ascension trakSTAR, Ascension Technology, Corp.) wrapped in cotton and attached to a six-inch coffee stirrer was used to provide the tactile stimulus. Two objects were tested: the subject's right and left arms. For this task, the arm that was being tested was placed in front of the subject on a drawing board with the palm downward (i.e., the dorsal arm facing upward and being tested). Each recording session began with the calibration of this sensor to 14 pre-defined points on each of the two arms: the fingertips, the knuckles, the medial and lateral margins of the wrist, and the medial and lateral margins of the elbow. Calibration

was followed by the actual task. The task involved drawing the probe along each arm and along each digit, so that across all five digits, the probe covered the entire two-dimensional vector space of the arm from fingertips to elbows. Each trial began with a 2.5 s cue to the experimenter indicating the arm (left or right) to be tested, a finger (thumb, index, middle, ring, pinky) and a direction (*proximal* or *distal*). Following a 1 s delay, the stimulation phase began, lasting a total of 8 s. During the stimulation phase, the trakSTAR was drawn along the arm from the indicated fingertip to elbow and back to fingertip (if the cued direction was *proximal*) and from elbow to fingertip and back to elbow (if the cued direction was *distal*). The full factorial of the task thus comprised 5 digits x 2 directions x 2 objects. Because the probe was pre-calibrated, the neural firing rate response to touch could be precisely correlated to spatial location along the arm. During the task, gaze was always fixated at the contact point, tracking the probe. We ran four sessions of this task on non-consecutive days, recording from a total of 421 neurons.

5.2.4.4 Reference frame task ***(RFT; Figure 5.4, Figure 5.4-figure supplement 1, Figure 5.4-figure supplement 2, Figure 5.4-figure supplement 3).*** This task was performed to evaluate how changes in gaze position (**Gaze**, or **G**) or touch location (**Field**, or **F**) affect neural firing rate. Two objects were tested: the subject's right arm (**Native**) and a sponge (**Sponge**). Because neurons could prefer one position of an object over others, we tested two positions for each object, acknowledging that this cannot adequately sample the infinite possible positions but only give a preliminary sense of how position may affect responses. Objects were placed on a drawing board that stood in front of the subject's wheelchair such that the subject could rest her arms conveniently on the drawing board. The drawing board was approximately 3 feet x 2 feet and was placed such that the longer edges were parallel to the floor and facing the subject. **Position 1** consisted of the object at approximately a 45° angle to the lower, longer edge of the table. **Position 2** consisted of the object at approximately a 90° angle to the same edge. Four locations or fields were defined along each object, approximately 3 cm apart. On the arm, these were on the knuckles, on the wrist, and two along the forearm. They were similarly spaced along the sponge. Each trial began with presentation of the effector to be tested and its position (over 5 s). The effector of interest was placed in the correct position at this time. Next, over 1 s, the location for gaze was cued as an audible sound (*1,2,3,* or *4* with *1* being farthest from the body and *4* being nearest*).* Simultaneously, the location for touch was cued on the screen (LCD monitor placed adjacent to the subject) visible to the experimenter. There was a brief 1 s pause followed by the onset of the stimulus phase, lasting 2 s. During the stimulus phase, touch was presented to the cued location, as brush strokes, at 1 Hz, for 2 s. The stimulus phase was followed by an inter-trial interval of 0.5 s. A schematic representation of the task is shown in **Figure 4A.** Four trials were performed for each unique combination of gaze and touch location. This task was performed in 6 sessions, during which a total of 589 neurons were recorded.

5.2.5 Quantification and statistical analysis

5.2.5.1 Linear analysis (***Figure 5.1B***, ***Figure 5.2B, Figure 5.1-figure supplement 1, Figure 5.1-figure supplement 2, Figure 5.1-figure supplement 3, Figure 5.1-figure supplement 4, Figure 5.1-figure supplement 5***). For each unit, we fit a linear model describing its firing rate as a function of response to each test condition. Response was defined as the mean firing rate between 0.5 and 2.5 s after stimulus onset. These times were chosen to minimize the influence of experimenter delays in presenting the stimulus. The baseline was defined as firing rate during the 1 s prior to stimulus presentations. The linear model was written as:

$$FR = \sum_c \beta_c \, \mathrm{X}_c + \beta_0$$

Where FR is the firing rate, X_c is the vector indicator variable for test condition *c*, β_c is the estimated scalar weighting coefficient for each condition, and β_0 is a constant offset term. A neuron was considered responsive to a particular condition if the t-statistic for its associated beta coefficient was significant (p<0.05, false discovery rate (FDR) corrected for multiple comparisons).

5.2.5.2 Discriminability index (***Figure 5.1C, Figure 5.2B, Figure 5.1-figure supplement 2, Figure 5.1-figure supplement 4***). To quantify how well neural activity can be discriminated from baseline (pre-stimulus) activity, we developed a cross-validated measure: discriminability index (DI). As with the linear analysis described above, the stimulation phase window was defined as 0.5 to 2.5 s after the stimulus onset, and baseline was defined as the 1 s prior to stimulus presentation). The firing rate of all recorded neurons was concatenated into a vector, denoted by A. The firing rate of each neuron during the baseline phase was similarly concatenated to form a vector, denoted by B. Next, a non-dimensional sensitivity index was computed as:

$$DI = \frac{\bar{A} - \bar{B}}{\sqrt{\frac{{\sigma_A}^2 + {\sigma_B}^2}{2}}}$$

Where $\bar{A}$ is the mean of the firing rate vector A, $\bar{B}$ is the mean of the firing rate vector B, σ_A is the standard deviation of the vector A, and σ_B is the standard deviation of the vector B. In **Figure 5.1C**, this calculation was performed in a time resolved manner, using a sliding time window (300 ms each, stepped at 10 ms intervals), beginning 1.5 s prior to the stimulus onset. Mean and bootstrapped 95% confidence intervals were computed for each time bin from the values computed for each session. The same baseline window was used as in the static version.

5.2.5.3 Classification (***Figure 5.1D***, ***Figure 5.2D***) *and time-resolved classification* (***Figure 5.1E***). Classification was performed using linear discriminant analysis (LDA) with the following parameter choices: (1) only the mean firing rates differ for unit activity in response to each touch location (covariance of the normal distributions are the same for each condition) and (2) firing rates for each unit are independent (covariance of the normal distribution is diagonal). The classifier took as input a matrix of firing rates for all sorted units.

The analysis was not limited to significantly modulated units to avoid 'peeking' effects.[69] Classification performance is reported as prediction accuracy of a stratified leave-one-out cross-validation analysis. The analysis was performed independently for each recording session, and results were then averaged across days. In **Figure 5.1E**, this analysis was performed in a sliding-time window manner (300 ms each window, stepped at 10 ms intervals), beginning 0.5 s prior to the stimulation onset.

5.2.5.4 Correlation (***Figure 5.2C***). To perform the population correlation analyses, we quantified the neural representations as a vector of firing rates, one vector for each condition with each vector element summarizing the response of an individual unit. As before, neural activity was summarized as the mean firing rate during the stimulation phase window, defined as 0.5–2.5 s after the onset of stimulus presentation. Firing rate vectors were constructed by averaging the responses across 50–50 splits of trial repetitions. The mean responses across different splits were correlated within and across conditions, then the splits were regenerated, and the correlation computed 250 times. To test whether the difference between any pair of conditions was statistically significant, we used a shuffle permutation test applied to the correlations computed over the 250 random splits.

5.2.5.5 Event related averages (***Figure 5.1-figure supplement 5, Figure 5.2-figure supplement 1***). For each unit, neural activity was averaged within 750 ms intervals starting from 0.5 s prior stimulation onset, stepping to 2.5 s after, in 100 ms step intervals. Responses were grouped by condition, and a mean and standard error on the mean (SEM) were computed for each time window and for each condition.

5.2.5.6 Selectivity curve analysis (***Figure 5.1-figure supplement 7***). This analysis was performed to evaluate whether the neural coding for touch relative to controls (overlap, adjacent, far stimuli). For each unit tuned to at least one condition ($p<0.05$, FDR corrected, see above), repetitions for each condition were split in half to create training and test splits of the data. Repetitions were then averaged to create a single value per condition for each of the training and the test sets. Training set data were rank ordered from the object resulting in the highest firing rate to the object resulting in the lowest firing rate. This computed order was used to sort the test data, and the firing rates for each of the four stimuli (touch, overlap, adjacent, far) to objects was noted. This process was repeated 500 times and the results averaged across folds. The result is a cross-validated measure of the response of each unit as a function of rank. Responses were normalized between 0 and 1, before averaging across the population of tuned units. Confidence intervals were estimated using a bootstrap procedure. Both the mean and the 95% confidence interval are shown.

5.2.6.7 Spatial coverage of tactile responsiveness (***Figure 5.3B, Figure 5.3C***). For analyses, the insensate arm region was divided into a grid of size 5 x 4 (5 fingers x 4 regions), where the 4 regions were the fingers, the palm, the distal forearm, and the proximal forearm. Each trial consisted of the magnetic probe touching a particular region twice (once moving from elbow to fingertips, and once moving the other way). Thus, across 12 trials, there were 24 times that the probe traversed each region. The neural response to a region

was summarized as the mean firing rate within a region (i.e., across all the times the firing rate was sampled as the tracker was drawn along the region), such that there a total of 5 fingers x 4 regions x 24 points in all. The baseline firing rate response was defined as the mean firing rate for that region in the 1 s immediately prior to the stimulus onset. For each unit, we performed a cross-validated partial least squares regression in which we evaluated for a systematic variation in the unit's response to each of the grid locations. A unit was considered significantly tuned if the F-statistic for this regression model was significant (p<0.05, FDR corrected for multiple comparisons). For comprehensive examination, this process was repeated with only those firing rates from when the probe was moving in one direction (e.g., fingertips to elbow) and not the other, and vice-versa.

For graphically illustrating the spatial coverage of touch responsiveness, the firing rate at each grid location was averaged such that there were two points now within each of the 5 x 4 grid, one for proximal motion of the probe (i.e., fingertips to elbow) and one for distal motion (i.e., elbow to fingertips). This grid of firing rates was first plotted as a heatmap in MATLAB. A computer graphics program (Blender, Inc.) was used to color a vector graphics model of a human arm with the heatmap generated in MATLAB. This could precisely be performed by dividing the arm into individual pixels and defining the color of the pixel based on the heatmap.

Finally, a PCA was performed on all the firing rate responses across all significant units together. First, all significant units were identified, independently for the right and the left arms. Their firing rates across the 5 x 4 grid were concatenated to form a vector. These vectors were then concatenated, column wise. A PCA was applied to this matrix (one for each of the right and left arms). The resulting principal components were transformed back into a 5 x 4 grid. These described the spatial coverage within each component.

5.2.6.8 Spatial localization of seen touch response (***Figure 5.3D***). For this analysis, we divided the entire length of the arm from fingertips to elbows into 10 evenly separated sites along each finger, for a total of 50 points. The magnetic tracking probe sampled spatial locations near each of these sites, for all of which there was a corresponding firing rate response. We first grouped all the sampled points (and corresponding firing rates), on each trial, by the site to which it was nearest. For each trial, the firing rate response of all the points attached to each site was averaged. Thus, there were 50 such firing rate measures for each trial. We performed an analysis in which we trained a linear model using gaussian process regression, based on the firing rates across the 50 sites, in a stratified leave-one-out cross-validation analysis. All trials but one was used to train the linear model. The model was tested by asking it to predict the contact site using data from the held-out trial. This process was repeated across all trials. The prediction accuracies are reported as mean along with the 95% confidence interval on the mean and categorized by the digit along which the site was. This analysis was done using held-out data from the same laterality (i.e., the same arm from which data were used to train the model), as well as across-laterality (using data from the arm not used to train the model). The analysis was performed separately for each recording session, and results averaged across days.

5.2.6.9 Stimulus direction prediction (***Figure 5.3E***). In this analysis, used LDA classification to predict the direction of motion of the probe, i.e., proximal, or distal. The classifier was given as input a matrix of firing rates across the 5 x 4 x 24 grid of firing rates (see Methods: *Spatial coverage of tactile responsiveness*), along with a label indicating the direction of the probe during the recorded firing rate. All sorted units were used in this analysis. The classifier predicted, in a stratified leave-one-out cross-validation, the probe motion direction on held-out data. The analysis was performed separately for each recording session, and results averaged across days.

5.2.6.10 Reference frame evaluation (***Figure 5.4B, Figure 5.4-figure supplement 2***). We used a gradient analysis to quantify how changes in the behavioral variables (field location (F), gaze or eye location (E), object position) influenced the neural response to touch. We created a four-by-four matrix (response matrix) representing the neural activity for each unique combination of touch location (F) and gaze position (E), in each of the two object positions, and for each object (Native, Sponge). Gradients were then determined using the gradient function in MATLAB. For each gradient, a resultant angle and length was computed to summarize the net direction and magnitude of change across the entire response matrix (idealized responses shown in ***Figure 5.4-figure supplement 1***). However, gradients often show a symmetrical pattern that would result in cancellation of symmetrical angles. To avoid this, we double each angle in the matrix and represent each angle from 0° to ±180°. Therefore, the summed resultant angle is represented by 0° for gradients oriented left and right, ±180° for gradients oriented up and down, and −90° for gradients oriented along the diagonal. The summed resultant angle and length however cannot be mapped directly onto the response matrix; thus, we have notated the appropriate variable and combinations of variables to help with interpretation. For example, in ***Figure 5.4-figure supplement 1***, gaze only (E) modulation would be found at ±180°, Field only (F) modulation is seen at 0°, E+F at 90°, and E-F at −90°. Therefore, we can use the angle of the resultant angle as a proxy for overall orientation bias for a variable or variable pair. This analysis was performed for each object, in both positions of the object. The position in which the object displayed a greater gradient was denoted the preferred position for the object. **Figure *5*.4B** and ***Figure 5.4-figure supplement 2*** show a distribution of orientations for the Native and the Sponge objects, respectively, in their preferred positions.

5.2.6.11 Correlation in neural response to ***seen*** *touch across positions* (***Figure 5.4-figure supplement 3***). We wished to examine whether there was symmetry in the response pattern for each unit to seen touch to the various combinations of gaze and touch location, across preferred and non-preferred response positions. For each unit, on each object, we concatenated the responses to all unique combinations of gaze and touch location (in each position) into a vector. This vector was correlated across preferred and non-preferred positions. The distribution of correlation coefficients across the neural population for the Native arm, and the sponge, are shown in ***Figure v4-figure supplement 3***. To test whether the mean correlation coefficient of this histogram distribution was statistically significant, we used a shuffle permutation test applied to the correlations computed over 250 random 50-

50 splits of trials. The distribution of shuffled differences served as the null distribution, against which we compared the true mean correlation coefficient to determine significance.

5.3 RESULTS

We recorded populations of single neurons from NS, a tetraplegic human participant (spinal cord injury at cervical levels 3-4; C3/4) in a BMI clinical trial. NS has no sensation or motor function below the shoulder-level. Recordings were made from a microelectrode array implanted in the left hemisphere at the junction of the post-central and intraparietal sulcus, during four tasks that probed neural responsiveness to ***seen*** touch to insensate regions of her arms, to other human arms, or to inanimate objects. Waveform sorting resulted in both well-isolated singe-units and multi-unit groupings. The main figures aggregate across sorted channels, but key analyses are performed separately for well-isolated and multi-unit activity (shown in supplemental figures) to ensure that pooling units did not bias our results.

5.3.1 Human PPC encodes seen touch (visual estimation of touch) to objects

We first examined the hypothesis that PPC neurons encode ***seen*** touch to more than just the human form. In a basic touch task (BTT), the subject viewed stimuli (brush strokes at 1Hz, for 2s) to one of four objects (**Figure 5.1A**). Stimuli included ***seen*** touch (Touch) and three controls: touch-like motion overlapping but not touching the object (Overlap), motion adjacent (Adjacent) and motion far (Far) from the object. Four objects were tested including the subject's insensate right hand (Native), a matched region on an experimenter's right hand (Exp-Right), a banana, and a sponge. There were thus two animate objects (arms) and two inanimate objects (banana, sponge). The touched site on each object subtended a consistent angle on the subject's eyes. Gaze (monitored by pupil tracking) was always directed at the stimulus location.

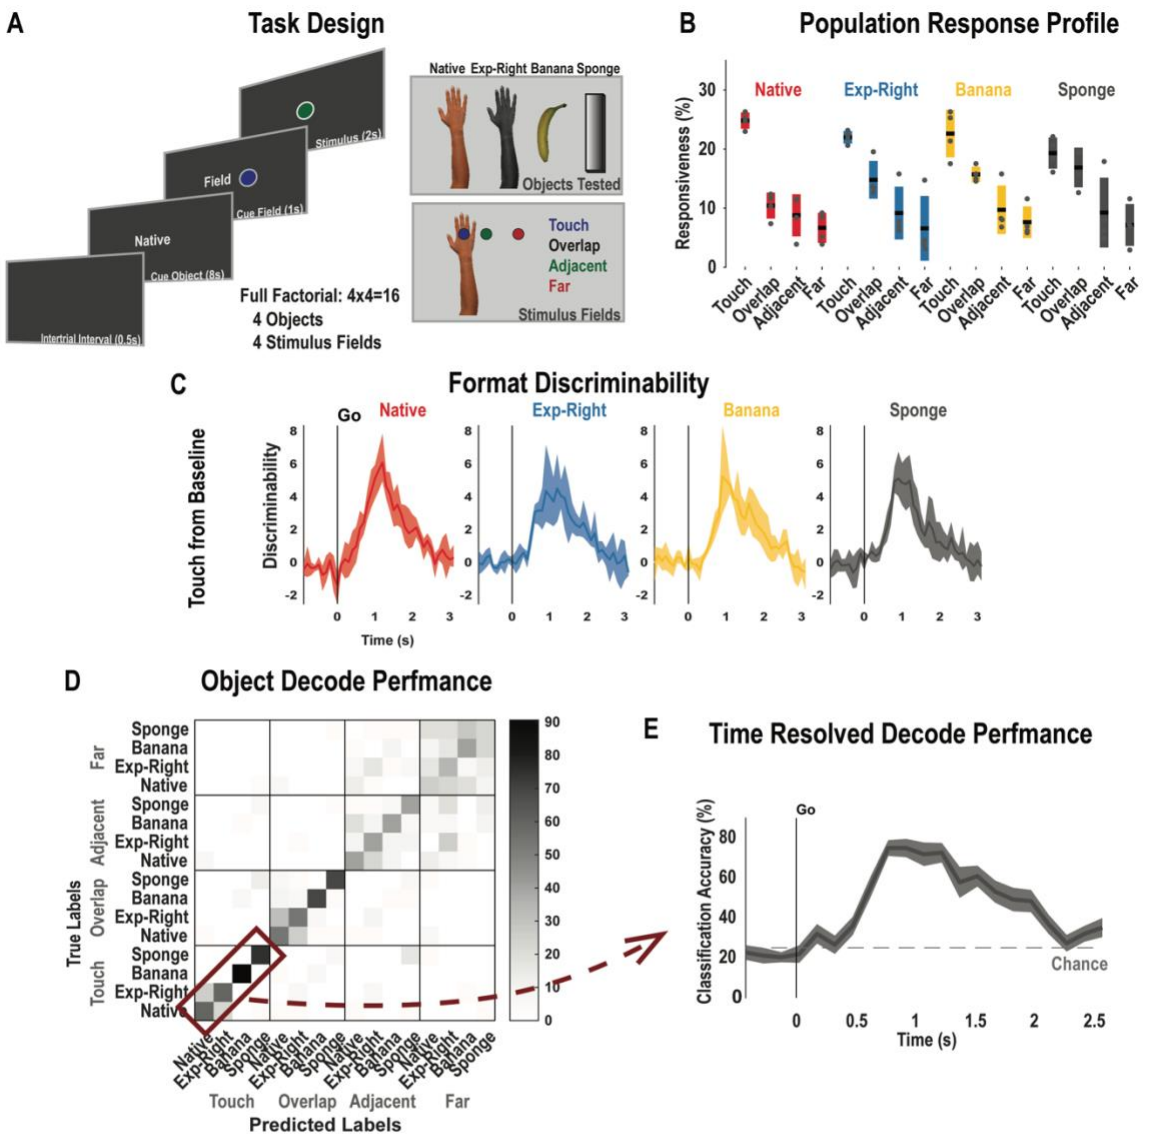

Figure 5.1. PPC preferentially encodes touch over visually matched controls.

A, Schematic illustrating the 4x4 factorial task design, in which we tested four stimuli applied to four objects: the subject's insensate right hand (Native), an experimenter's right hand (Exp-Right), a banana, and a sponge. The four stimuli were Touch, touch-like motion overlapping but not contacting the object (Overlap), touch-like motion adjacent (Adjacent) to the object, and touch-like motion far (Far) from the object. The phases of each trial are shown, including a cue indicating the object to test (Cue Object), the stimulus to provide (Cue Field), the actual stimulation phase (Stimulus), and an intertrial interval. The timings for each phase are indicated. B, Percent of the PPC neuronal population that demonstrated significant modulation relative to baseline for each tested condition ($p<0.05$, FDR corrected, $n=382$ units). Results are shown as the mean percentage (horizontal black line) of the population ± bootstrapped 95% confidence interval (CI; bar height). Population results were not affected by pooling together single and potential multi-units (Figure 1-figure supplement 1). Gray dots represent single session results. C, Discriminability of neural population activity during seen touch to the four objects, from baseline neural activity. Each column depicts the time-resolved (x-axis) discriminability (y-axis) for one object (labeled), from 1 s prior to stimulation onset, to 3 seconds after. Discriminability is shown as the mean ± standard error on the mean (SEM) over $n=10$ trials. The vertical line depicts the onset of the stimulus phase. D, Confusion matrix of the cross-validated classification accuracy (as percentage) for predicting conditions from population neural data. Shades of gray indicates that cross-validated accuracy, as in the scale. The matrix is an average of the confusion matrices computed for each recording day individual. E, Time-resolved classification accuracy (y-axis) for the touch stimulus to the four different objects is shown as a function of time (x-axis). Accuracy is shown as mean ± 95% CI. Chance accuracy is shown as a dashed gray line. The vertical line indicates the onset of the stimulation phase. Exp-Right, experimenter's right hand; s, seconds

*5.3.1.1 PPC neurons encode **seen** touch to all tested objects.* For each neuron, we fit a linear model that explained firing rate as a function of responses to each stimulus location. Neural response to a stimulus was considered significant if the t-statistic for the associated beta coefficient was significant ($p<0.05$, false discovery rate [FDR] corrected for multiple

comparisons). Of 382 total neurons recorded in this task, 227 responded to at least one stimulus (ANOVA). Across objects, the PPC population preferentially encoded ***seen*** touch over other stimuli ($\chi^2(1)$**=**3175.44, p<0.05; **Figure 5.1B**; **Figure 5.1-figure supplement 1**), although all conditions were significantly represented. The responsiveness to touch did not differ across objects. Moreover, the strength of the neural response, measured by the discriminability of neural response from baseline activity, did not differ across objects (**Figure 5.1C**). We separately found, in a multiple object task (see Methods: MOT) that the PPC response to ***seen*** touch is expansive (**Figure 5.1-figure supplement 2; Figure 5.1-figure supplement 3**) and includes many categories of objects such as tools (scissors, hammer), household items (shoes, toys), fruit (banana) and of course, the sponge. These results are consistent with the PPC response to ***seen*** touch being broad, if not universal.

5.3.1.2 Seen touch is preferred over controls even when it is not the target of gaze. In a control version of the task (**Figure 5.1-figure supplement 4**), gaze was directed at the adjacent site during all stimuli. As in the version when gaze was directed at the touch site, the PPC responsiveness to ***seen*** touch was greater than to the controls, indicating that the preferential encoding of touch by this population is not simply due to touch being at the focus of gaze. In further support that gaze was not the main driver of the results, the response profile to the controls did not differ either, between the two versions of the task.

5.3.1.3 Is the PPC response a measure of stimulus distance from the body? If this were true, we would expect that the population responsiveness to the overlap condition should be greater than to adjacent, which in turn should be greater than to the far condition. This is, however, not the case. Although a significantly larger fraction of the PPC population responds to overlap than to the other controls in some conditions (e.g., when gaze is at the touch site), this is not consistent across gaze positions. Moreover, there is no significant difference between responsiveness to adjacent and far stimuli. Together, these findings suggest that this PPC population prefers seen touch and may also encode visual stimuli (Overlap) in the immediate vicinity of the body (i.e., peri-personal space) but does not encode stimuli as a function of distance from the body.

Individual neurons were heterogeneous, representing one or more stimuli, to one or more objects (**Figure 5.1-figure supplement 5**). Representative examples of neurons highlighting this variability are shown in **Figure 5.1-figure supplement 6**. To understand single unit responsiveness, we performed an analysis in which for each unit, we first rank ordered the objects based on the maximal mean response they evoked (across the four stimuli; **Figure 5.1-figure supplement 7**). Across all such ordered unit responses, we found that within each object rank, touch evoked a greater mean response than control stimuli, paralleling the population activity (as in **Figure 5.1B**). This is confirmed also by a principal component analysis (PCA), in which most of the variance in neural activity across rank ordered units is explained by touch (**Figure 5.1-figure supplement 7**).

5.3.1.4 PPC discriminates the stimulus-object response but not object identity alone. As expected, a population of heterogeneous neurons enabled accurate cross-validated

classification of the PPC responses to stimuli (**Figure 5.1D**). Consistent with touch representations being strongest in this population, the classification accuracy for touch was higher than for the controls. The greatest confusion within this classification was for the two arms (the subject's and the experimenter's). This indicates that the PPC response is sensitive to object identity, and that similar object features may evoke more similar responses to ***seen*** touch, a point which we expand upon next. Of note, in a sliding time version of this analysis, classification accuracy prior to stimulus onset is at chance level (**Figure 5.1E**), indicating that the discriminability is for the stimulus-object response and not for object identity alone.

5.3.2 PPC responses to seen touch are sensitive to object attributes

To understand the sensitivity of touch responses to object features further, we performed a task (MOT; **Figure 5.2A**) in which the subject viewed ***seen*** touch to one of three locations on one of four objects. Three objects were arms: the subject's insensate right arm (Native), matched regions of an experimenter's right arm (Exp-Right), and the same experimenter's left arm (Exp-Left). The fourth object, a control, was the sponge. Three fields (touch locations) were defined on each object. On the arms, these were the hand (field 1), the distal forearm (field 2), and the proximal forearm (field 3).

Of 672 neurons recorded in this task, 478 discriminated touch locations on the four objects (**Figure 5.2B**). All touch locations are robustly represented, and comparably discriminable from baseline (**Figure 5.2B**). Some example neurons are shown in **Figure 5.2-figure supplement 1**, illustrating response to specific conditions. At the level of the neural population, a cross-validated correlation analysis (**Figure 5.2C**) shows that ***seen*** touch responses to the three arms are more similar than to each other than to the sponge. This is true in every field, and especially prominent for field 1 and field 3. For example, the mean correlation coefficient for the representation in field 1 of the native arm to the two experimenter arms (exp-right and exp-left) is 0.84 and 0.77, significantly greater than the between the native arm and the sponge (0.6; $p<0.05$ for both comparisons).

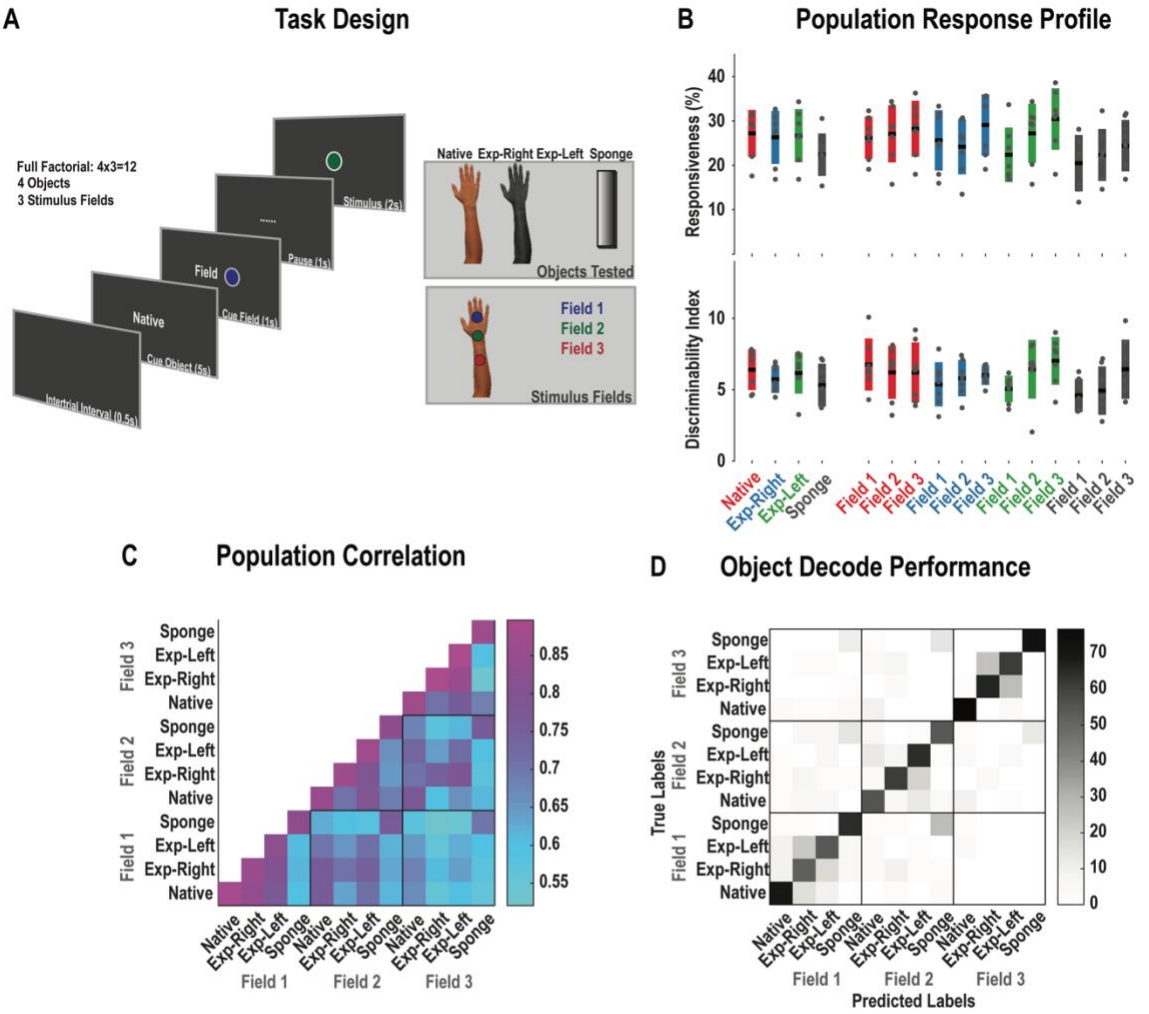

Figure 5.2. Seen touch responses in PPC are sensitive to object identity

A, Schematic illustrating the 4x3 factorial task design, in which we tested seen touch applied to three fields on 4 objects: the subject's insensate right arm (Native), an experimenter's right arm (Exp-Right), the same experimenter's left arm (Exp-Left), and a Sponge. The three evenly spaced fields were chosen such that Field 1 was on the hand, Field 2 was on the distal forearm, and Field 3 was on the proximal forearm, and similarly positioned on the sponge. The phases of each trial are shown, including a cue indicating the object to test (Cue Object), the field to touch (Cue Field), a small delay (Pause) the actual stimulation phase (Stimulus), and an intertrial interval. The timings for each phase are indicated. B, (Top), Percent of the PPC neuronal population that demonstrated significant modulation relative to baseline for each tested condition (p<0.05, FDR corrected, n=545 units). Results are shown as the mean percentage (horizontal black line) of the population ± bootstrapped 95% confidence interval (CI; bar height), for all fields together (on the left) and then divided by field (on the right). Population results were not affected by pooling together single and potential multi-units (Figure 1-figure supplement 3). B, (Bottom), Discriminability of neural population activity during seen touch to the seven objects, from baseline neural activity. As in the top panel, results are shown as mean percentage (horizontal black line) of the population ± bootstrapped 95% confidence interval (CI; bar height), for both fields together on the left and then divided by field on the right. In both panels, individual dots represent single-session results. C, Cross-validated population correlation within and across test conditions. Colors indicate strength of correlation as in the scale. The matrix is an average of the results computed for each day. D, Confusion matrix of the cross-validated classification accuracy (as percentage) for predicting conditions from population neural data. Shades of gray indicates that cross-validated accuracy, as in the scale. The matrix is an average of the confusion matrices computed for each recording day. Hz, Hertz; Exp-Right, experimenter's right hand; s, seconds.

The object sensitivity manifested also in a cross-validated classification analysis (**Figure 5.2D**). All conditions could be accurately discriminated. However, as in the correlation analysis, representations of the three arms are similar, resulting in confusion. For example, in field 1, the native arm tends to be confused for the experimenter's right arm, less so for the experimenter's left arms, and almost negligibly to the sponge. This supports that similar object features evoke more similar PPC responses to ***seen*** touch. Of note, this analysis shows that the experimenter's right arm tends to be confused more often for the

experimenter's left arm, than for the subject's right arm, suggesting that the feature space that drives touch responses may be more complex (e.g., involving skin color, texture, hair, etc.) rather than just shape alone.

5.3.3 PPC response to seen touch is spatially sensitive but not spatially restricted

The preceding task showed that ***seen*** touch responses depend upon the features of the touched object. But are all regions of objects represented? How do individual units enable the population response? We performed a spatial coverage task (SCT; **Figure 5.3A**) in which we recorded neural activity while the subject viewed touch along the full length (insensate) of her right and left forearm. Touch was provided by a magnetic tracking probe that was pre-calibrated to defined points along the two arms (such as fingertips and elbows). It was then traced, along each finger, either proximally (from fingertip to elbow) or distally (from elbow to fingertip), such that the entire arm was covered, and the neural response to touch could be precisely correlated to the spatial location of the probe. Manipulating the starting direction of the probe additionally allowed us to evaluate whether there was directionality to the touch responses (i.e., whether the probe moving distal or proximal modulated the response).

For analyses, we divided the forearm into four regions: the fingers, the palm, the distal and the proximal forearm. In all, 189 or 399 units recorded in this task were significantly modulated (**Figure 5.3-figure supplement 1**) by ***seen*** touch to at least one region (and in at least one direction of the probe), on the right arm (146) or the left arm (143). The individual unit responses on both left and right arms were strongly correlated ($r=0.83$, $p<0.05$). Similarly, for each unit, the response during proximal probe motion and distal motion were strongly correlated, both for the right arm ($r=0.77$, $p<0.05$) and for the left arm ($r=0.72$, $p<0.05$). These findings together indicate that unit responses to touch are consistent however touch is applied (i.e., direction). Moreover, consistent with our previous studies, information generalizes across similar contexts (here, from left arm to right arm).

Individual units exhibited response fields to ***seen*** touch that spanned multiple regions. The response fields for some example units are shown in **Figure 5.3B**, during both proximal and distal probe motion, and on both left and right arms. A PCA performed on the pooled response across all units (**Figure 5.3C**) showed that the population collectively provides full spatial coverage for ***seen*** touch to both left and right arms. ***Seen*** touch to the hands (on both the left and right sides) accounted for most of the variance in population neural activity. The projection of the first three components of the PCA during the distal and proximal probe motions is shown in **Figure 5.3C**.

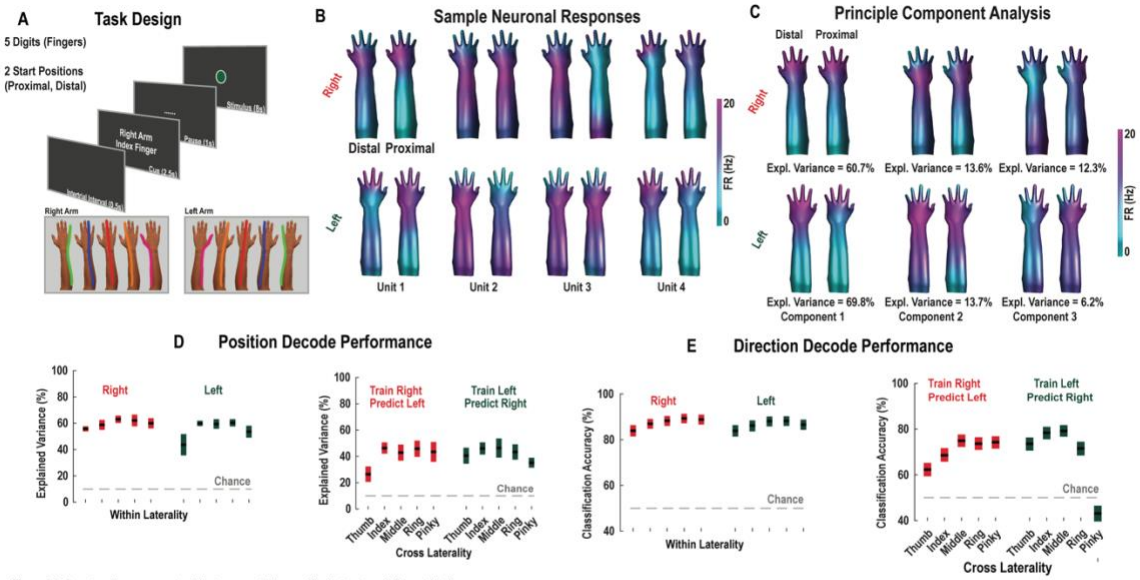

Figure 5.3. Seen touch responses to objects are spatially specific, but not spatially restricted.
A, Schematic illustration of the task design. After the tracking probe was calibrated to pre-defined points on the subject's insensate arms (bilaterally), each trial (of 12 trials) began with a cue to the experimenter indicating the finger and the arm along which to draw the probe (Cue), followed by a brief pause (Delay), then the stimulation phase (Go), followed by an intertrial interval. The paths along which the probe was drawn along the arms are shown in the panels at the bottom (for right arm on the right and left arm on the left). B, Example responses of individual neurons. Each panel of two arms (e.g., top row, left most two arms) shows a single unit's response to probe motion as it was being drawn distally from elbow to fingertips (distal; left) and proximally from fingertips to elbow (proximal; right). Each column of panels is a single unit's response to seen touch to the right arm (top row) and left arm (bottom row), shown in a color as in the scale on the right. C, The projection of population response along the first three components of a PCA on all units together, but separately for the right (top row) and left (bottom row) arms, is shown. The first three components are shown (left to right), and the explained variance along each component is indicated. The strength of the projection is as in the color scale. D, A linear model was trained on population data and tested on held-out data, to predict spatial location of seen touch responses. In the left panel, the prediction accuracy (y-axis; explained variance) is shown for a model trained on data from the right arm, used to predict on test data from the right arm (left half), and similarly trained on data from the left arm and used to predict on test data from the left arm (right half). The right panel shows a similar analysis but performed across-laterality (trained on one arm, tested on the other). The prediction accuracy for test data drawn from each finger is shown (x-axis). Chance accuracy is shown as the gray dashed line. E, Similar to panel D. A linear classifier was trained on population data and tested on held-out data, to predict direction of probe motion (proximal or distal) of seen touch responses. In the left panel, the classification accuracy (y-axis) is shown for a model trained on data from the right arm, used to predict on test data from the right arm (left half), and similarly trained on data from the left arm and used to predict on test data from the left arm (right half). The right panel shows a similar analysis but performed across-laterality (trained on one arm, tested on the other). The classifications accuracy for test data drawn from each finger is shown (x-axis). Chance accuracy is shown as the gray dashed line. Expl, explained; Hz, Hertz.

We next asked: in a population that provides full spatial coverage of ***seen*** touch to objects, yet is sensitive to object features, can the precise spatial location of touch be decoded from the population activity? To answer this, we first divided the arm length into a grid of size 5 x20 (5 fingers, 20 evenly sized divisions from fingertip to elbow). We found that a trained model could accurately predict the precise location of held out data (**Figure 5.3D**) on both the right arm (59.95%, $p<0.05$) and the left arm (55.43%, $p<0.05$). Moreover, information generalizes across left and right arms. A model trained on data from the right arm could predict test data from the left arm (41.02%, $p<0.05$), and vice versa (42.33, $p<0.05$). We found similarly that not only touch location, but even the direction of probe motion could be accurately decoded from population data (**Figure 5.3E**), within laterality (right: 87.44%; left: 86.57%), as well as when trained on data from right and tested on left (70.73%) and vice versa (72.79%).

5.3.4 PPC responses to <u>seen</u> touch are encoded in a variety of reference frames

Our findings thus far indicate that ***seen*** touch responses in PPC are sensitive to object features, and preferentially encoded over matched controls (e.g., overlapping but non-contacting motion). But does changing eye position modulate the response to touch? Moreover, if the object changes position in space, does the PPC response to touch remain constant? In other words, does information about touch to an object in one position generalize to that in another position?

We performed a reference frame task (RFT; **Figure 5.4A**) to evaluate whether responses to ***seen*** touch are object centric, or relative to gaze (i.e., eye position). We systematically

manipulated touch and gaze position across 1) 2 objects (subject's right arm and sponge), in 2) 2 positions each, with 3) 4 locations on each object across which touch and gaze were independently varied. This task allowed us to identify whether units behaved according to one of three types (**Figure 5.4-figure supplement 1**): 1) eye position units that responded maximally to a particular eye position regardless of where touch is applied, 2) field units that responded maximally to ***seen*** touch to a particular field on an object regardless of eye position, or 3) eye-field units responded maximally when gaze was targeted at the touched field.

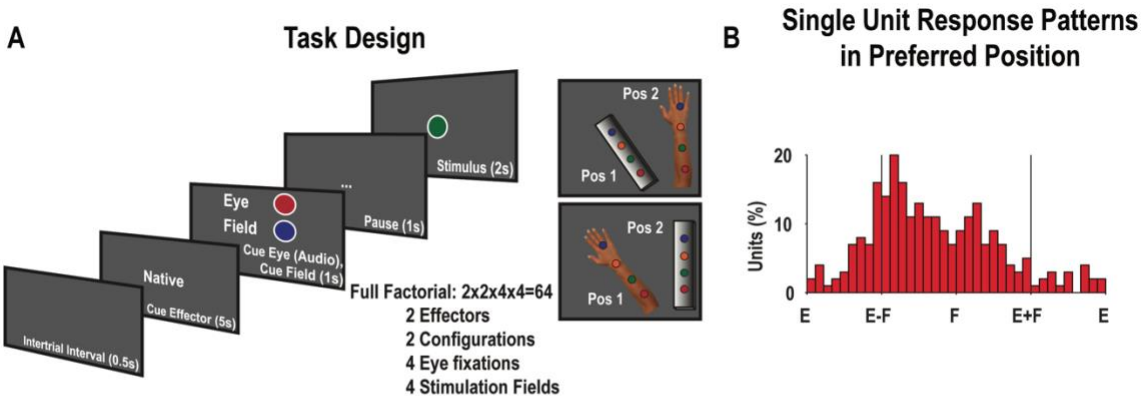

Figure 5.4. Seen touch is encoded in a mixture of reference frames.
A, Illustration of the task design. Overlaid panels show example trial epochs, while the right panels show possible configurations. B, PPC population reference frames for the native arm in the two labeled positions (shown in panel A). For each unit, gradients, and overall resultant angles were computed as described. The distribution of angles for the preferred position is plotted as a histogram. Any unit with a significant modulation (by linear regression) to touch in any condition was included. The number of such significantly modulated units is indicated as a percentage of the recorded population. The different types of reference frames are indicated: E (encoding eye position only), E-F and E+F (encoding touch as a function of eye position), and F (encoding touch to a particular field only). S, seconds.

For each unit, we identified its preferred response position (for each object) as that evoking the maximal mean response across all 16 conditions (4 touch locations and 4 gaze positions). On both objects (**Figure 5.4B** and **Figure 5.4-figure supplement 2**), there was a heterogeneity in the behavior of individual units (in their preferred position for that object). Most units were eye-field units that encoded ***seen*** touch relative to eye position, although a number were field units whose maximal response was tied to features of the object (a particular field) regardless of eye position. Across units, the response patterns generalized across preferred and non-preferred positions (**Figure 5.4-figure supplement 3**).

5.4 DISCUSSION

Touch is a ubiquitous perceptual process. When we are touched on ourselves, our brains typically receive both visual and tactile information regarding the stimulus.[30-32] Just as often, however, we perceive touch through visual information alone, such as when we ***see*** touch to objects around us. How does our brain ***see*** touch we can't ***feel***?

In recent work, we identified a population of human PPC neurons that processes many forms of touch: ***felt*** touch to oneself, ***imagined*** touch to oneself, and ***observed*** touch to others.[5] Although individual neurons are complex, the population creates rich associations

across neurons, representing basic-level, elemental information that relates to touch types, or body parts.[2] This information is flexible, generalizing across diverse contexts, including during cognitive touch processing (in the absence of tactile input).

Our current study extends these findings, advancing our understanding of touch perception. We show, through human single-unit recordings, that 1) PPC neurons encode ***seen*** touch even to oneself, independent of ***felt*** touch, 2) these neurons more broadly encode ***seen*** touch to objects around us, not just animate but also inanimate, and 3) human PPC also encodes touch related variables such as the spatial point of contact, and direction of motion of stimulus. These results represent novel findings, previously unstudied in humans or in animal models.

5.4.1 PPC provides a neural basis for seen touch to oneself

In NHPs, PPC neurons are often multimodal: they respond to visual, tactile, and other sensory stimuli.[30-33,145] During touch to oneself, multiple studies suggest that PPC neurons estimate the stimulus in all available modalities (e.g., visual, tactile) and combine estimates into a unified representation.[34,54,141-143,146] In support, for example, the perceptual accuracy of touch is enhanced when visual information is available, and concordant, with tactile information. Conversely, it can confound touch perception when it is discordant with tactile information.[147]

Our study is the first neuron-level characterization of the human PPC response to ***seen*** touch, completely disentangled from simultaneous tactile information. We make use of a unique opportunity to record from a tetraplegic human subject with intact cortex while applying touch to insensate arm regions. We find that PPC neurons encode spatially structured receptive fields to ***seen*** touch, bilaterally, and provide full spatial coverage. Individual units are heterogeneous, with response fields organized by body part (e.g., hand, or forearm). At both the level of individual units and the neural population, however, information generalizes across body sides. This mirror symmetry is consistent with our previous findings for ***felt*** touch representations, and ***imagined*** touch, within this same neural population.[5] It reflects a broad theme in our current study: the PPC population supports similar behavioral contexts in similar ways. In other words, touch to objects with similar attributes evokes similar PPC neural activity, whether the objects are one's own left and right arms, or others' arms.

5.4.2 Seen touch is encoded broadly, even to objects other than oneself

Foundational studies in NHPs have documented that some PPC neurons encode not only stimuli to the monkey's body but also to others', and to realistic substitutes.[32] For example, some neurons have been identified that respond to position of a monkey's arm covered from view, as well as to the ***seen*** position of another, and to a realistic look-alike.[32] In our own previous work, we found that the same neural population we recorded from here, responded to both ***felt*** touch to oneself, and ***seen*** touch to other individuals.[5] Our current results extend

this work. We find here that the response to ***seen*** touch within this PPC population is not limited to oneself, or to other humans, but also includes objects with no resemblance to the human form (like a shoe). These findings are consistent with recent human neuroimaging studies that indicate similarly, that ***seen*** touch to animate and inanimate objects share cortical activation with felt touch to oneself.[57] We find that on all objects, touch is preferentially encoded over other visual stimuli (such as motion near the object). Yet, the touch responses are sensitive to object identity. As we mentioned above, similar objects (here, the arms), engage the neural population more similarly.

PPC neurons may also be sensitive to where touch occurs relative to one's body. In **Figure 5.2B**, for example, across objects, PPC touch responsiveness (and discriminability) to field 3 (nearest to the subject's body) was higher than to fields 1 (farthest) and 2, although all fields were represented. A similar result has been described within PPC in the motor neurophysiology literature. When we manipulate objects, for example, actions that move objects away from ourselves (e.g., push) evoke activity different from those actions that bring objects closer toward us (e.g., grasping).[148] One hypothesis is that PPC represents processes (including ***seen*** touch) in our immediate vicinity (our *peri-personal space*) differently than those that occur further away. Consistent with such a view, although touch is preferentially encoded over controls in our task, the overlap condition tended to be represented more strongly than the near or far conditions.

5.4.3 Could neural activity in PPC during <u>seen</u> touch arise from alternate cognitive processes?

Processing ***seen*** touch could engage resources involved in action observation (i.e., the action that provides the stimulus), or allocation of attention, or other processes. In previous work from our lab, for example, we found that this PPC population encoded manipulative actions (e.g., drag, drop, push, grasp) that the subject ***observed*** being performed by others.[2] Indeed, PPC has been described to belong to an action observation network (AON), a system of distributed brain regions that contributes to our perception of various classes of actions.[136] Yet, in the current study, it appears unlikely that the coding for ***seen*** touch simply reflects the general phenomena of action observation for at least three reasons. First, we found that visually matched controls (touch like motion without contact at various distances) did not evoke the same response as to touch. Second, we found that this same substrate encodes many other forms of touch including ***imagined*** touch to oneself and ***felt*** touch. Moreover, ***felt*** touch is encoded at short latency, suggesting that this PPC region is involved in the rapid somatosensory processing of touch.[5] Third, the structure of ***seen*** touch responses to the subject's insensate arm regions shared a similar body-part organization and left-right mirror symmetry to ***felt*** touch representations in our previous work.[4,5] The recruitment of a substrate involved in ***felt*** touch representations, with shared structure across ***seen*** and ***felt*** touch to oneself, and its engagement in ***seen*** touch more broadly, supports a view that this substrate represents touch and that ***seen*** touch responses do not simply reflect neural activity related to action observation.

Gaze can direct attention to a particular target, and attention has been shown to enhance stimulus response.[124-126,149] Thus, the preferential encoding of touch over controls when touch was the target of gaze could have been a gain modulation due to attention. However, even when gaze was fixated at an alternate site (the adjacent site), we found the same preferential response for touch over controls. We note, however, that from the reference frame task, a substantial fraction of neurons in this PPC population encoded ***seen*** touch relative to gaze. This suggests that the population represents touch preferentially over other stimuli even when it is not the direct target of gaze, but the strength of response to touch is modulated by gaze, possibly mediated by directing attention to the touch location.

5.4.4 Unpacking the functional organization of a shared substrate across diverse forms of sensory, motor, and cognitive behaviors

Aggregating across our current and previous work, this PPC substrate processes ***attempted*** and ***observed*** actions, ***semantic*** processing of action words***,*** motor, and tactile ***imagery***, ***felt*** touch to oneself, ***seen*** touch, and other variables.[2-5] Classically, shared representations across behaviors have been interpreted through *mirror neurons.*[13,24,27,150] These are neurons that were first discovered in high-level cortex in monkeys and activate *congruently* (i.e., identically) both when a monkey performs a *specific* movement (e.g., grasping an apple) and when it observes someone else performing the corresponding movement.[35,36,151-153] Their *specificity* and *congruency* inspired a *mirror hypothesis:* we understand others' intentions by mapping the visual representations of their actions onto neurons that process our own movements.[154] Similar neurons have been postulated to enable understanding in other domains as well, such as for emotions or sensations, collectively forming a system that supports how we understand the world.[12,28] Yet, the mirror neuron hypothesis is restrictive. If understanding actions, emotions, or sensations comes from activating neurons that process our own, mirror neurons alone cannot explain how we understand actions we cannot perform (e.g., flying), or sensations or emotions we may never have experienced before.[29] It is also difficult to understand how a mirror neuron that processes ***felt*** touch to ourselves may activate to touch we ***see*** to a shoe, or a toy.

Mounting recent literature points to a *neural population doctrine*, in which the fundamental units of computation in our brains are not individual neurons, but rather populations of neurons.[1] Consistent with this, in our previous and current work, we have found that individual neurons exhibit variability and heterogeneity in the variables they respond to. At a population level, however, rich associations develop across neurons, enabling meaningful relationships to be encoded. We termed this functional organization *partially mixed selectivity*.[4] This organizational structure enables many related variables (or behaviors) to be simultaneously represented by a shared substrate.

Within the sensory domain, we previously found that this substrate encodes many forms of touch. Subspaces of population activity codify meaningful, basic-level information such as the type of touch, where on the body is touched. These can combine in diverse ways to inform touch behaviors, and generalize across individuals, such that this substrate can

support not only touch when ***felt*** but also when ***imagined*** or ***observed*** occurring to others. At the neuron-level, this is made possible by the partially mixed functional organization. At the population level, this organization is consistent also with the cognitive behavioral view that at the heart of human cognition are flexible internal models of the world. These models are automatically engaged during all our behavior, including cognition, and are functionally a blueprint of the world around us that can generalize across diverse conditions to inform our repertoire.[6,14] Our current results accord with such a view that PPC encodes a flexible internal model (here, for touch) that supports diverse contexts of touch, including ***seen*** touch to objects in the world, not limited to the human form.[6,14]

5.5 CONCLUSION

We provide the first human neuron-level evidence that PPC encodes seen touch to oneself, to other individuals and to inanimate objects. The population encodes seen touch in a manner sensitive to object identity, yet the touch response is broad, if not universal. Our work adds to previous findings that this population of neurons also supports felt touch and imagined touch to oneself. Collectively this work supports a view that populations of neurons in shared high-level brain regions like PPC encode information in a versatile manner that can support diverse forms of behavior including cognition, as we show here through touch and tactile cognition.

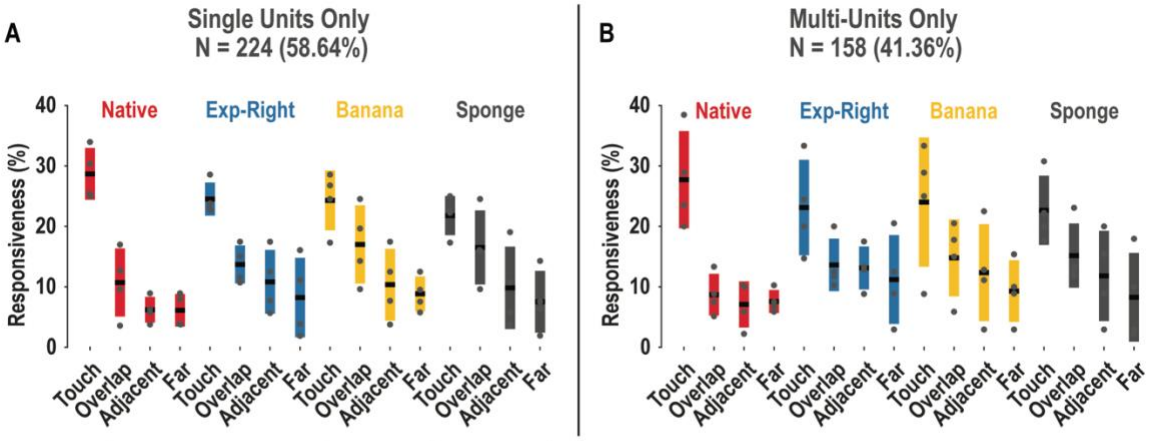

Figure 5.1-figure supplement 1. Preferential PPC encoding of touch over controls is not changed by pooling together single and multi-units.

A, Percetage of the PPC population that demonstrated significant modulation to each condition ($p<0.05$, FDR corrected) when considering only high-quality single units. As in Fig 1A, the black horizontal line represents the mean, and the height of the bar represents the 95% confidence interval. Gray dots represent single session results. B, Similar to A, but considering potential multi-units. N, sample size.

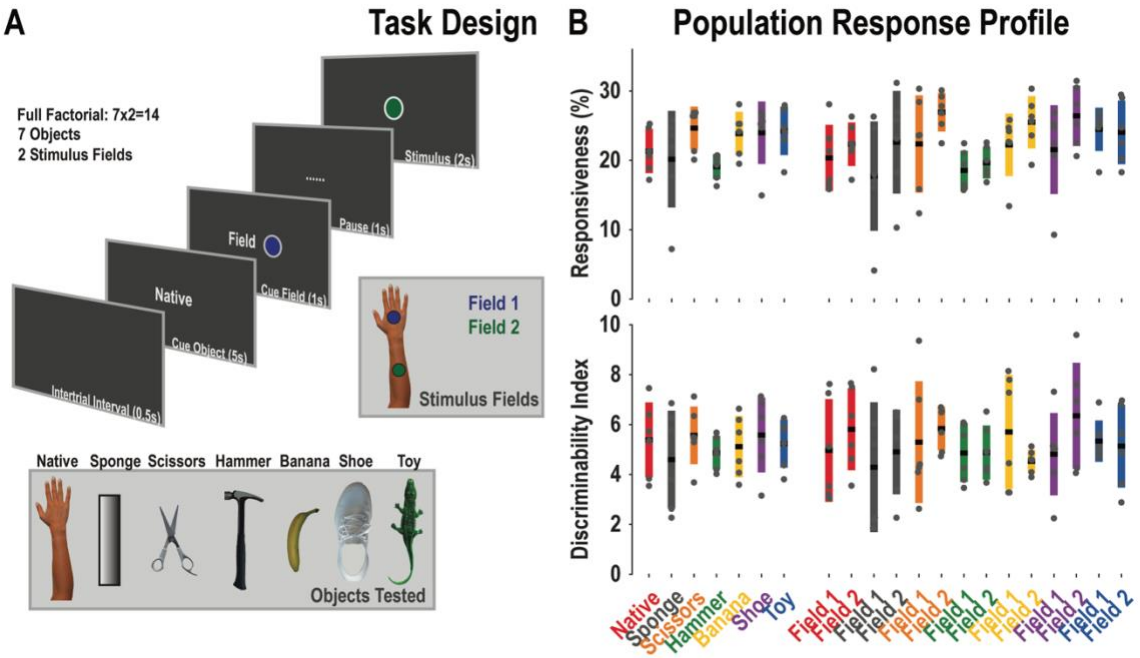

Figure 5.1-figure supplement 2. PPC represents seen touch to all tested animate and inanimate objects.

A, A, Schematic illustrating the 7x2 factorial task design, in which we tested seen touch applied to two fields on 7 objects: the subject's insensate right hand (Native), a sponge, scissors, a hammer, a banana, a shoe and a toy alligator. The two fields were equally separated on all objects. Field 1 was distal to the subject. The phases of each trial are shown, including a cue indicating the object to test (Cue Object), the field to touch (Cue Field), a small delay (Pause) the actual stimulation phase (Stimulus), and an intertrial interval. The timings for each phase are indicated. In the panel on the bottom, the orientation of each object relative to the subject are shown. The right panel shows the orientation of the two fields, on the Native arm, as an example. B, (Top), Percent of the PPC neuronal population that demonstrated significant modulation relative to baseline for each tested condition ($p<0.05$, FDR corrected, $n=545$ units). Results are shown as the mean percentage (horizontal black line) of the population ± bootstrapped 95% confidence interval (CI; bar height), for both fields together on the left and then divided by field on the right. Population results were not affected by pooling together single and potential multi-units (Figure 1-figure supplement 3). Gray dots represent single session results. B, (Bottom), Discriminability of neural population activity during seen touch to the seven objects, from baseline neural activity. As in the top panel, results are shown as mean percentage (horizontal black line) of the population ± bootstrapped 95% confidence interval (CI; bar height), for both fields together on the left and then divided by field on the right.

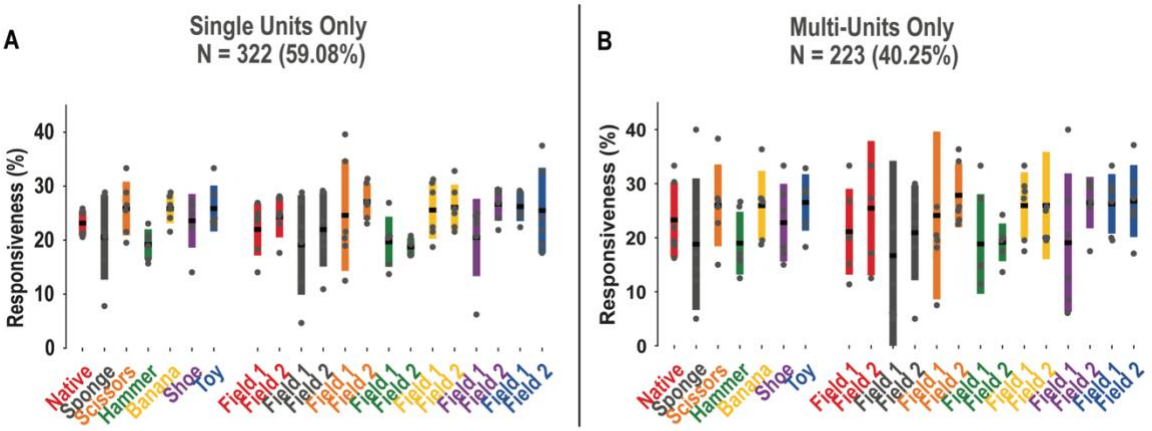

Figure 5.1-figure supplement 3. PPC coding of seen touch to multiple objects is not changed by pooling together single and multi-units.

A, Percentage of the PPC population that demonstrated significant modulation to each condition ($p<0.05$, FDR corrected) when considering only high-quality single units. The black horizontal line represents the mean, and the height of the bar represents the 95% confidence interval. Gray dots represent single session results. B, Similar to panel A, but considering potential multi-units. N, sample size.

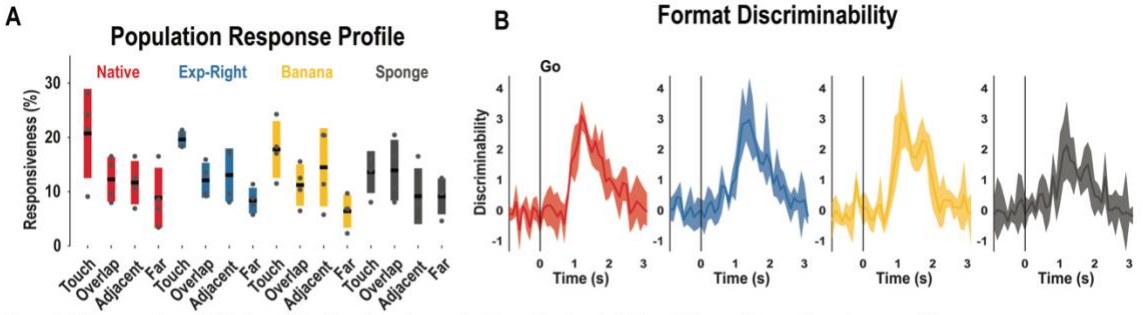

Figure 5.1-figure supplement 4. Preferential coding of touch over visually-matched controls is not influenced by an alternate gaze position.

A, Percent of the PPC neuronal population that demonstrated significant modulation relative to baseline for each tested condition ($p<0.05$, FDR corrected, n=382 units). Experimental design identical to Fig 1A, except here gaze was always targeted at the Adjacent site, rather than at the Touch site. Results are shown as the mean percentage (horizontal black line) of the population ± bootstrapped 95% confidence interval (CI; bar height). Gray dots represent single session results. B, Discriminability of neural population activity during seen touch to the four objects, from baseline neural activity. Each column depicts the time-resolved (x-axis) discriminability (y-axis) for one object (labeled), from 1 s prior to stimulation onset, to 3 seconds after. Discriminability is shown as the mean ± standard error on the mean (SEM) over n=10 trials. The vertical line depicts the onset of the stimulus phase.

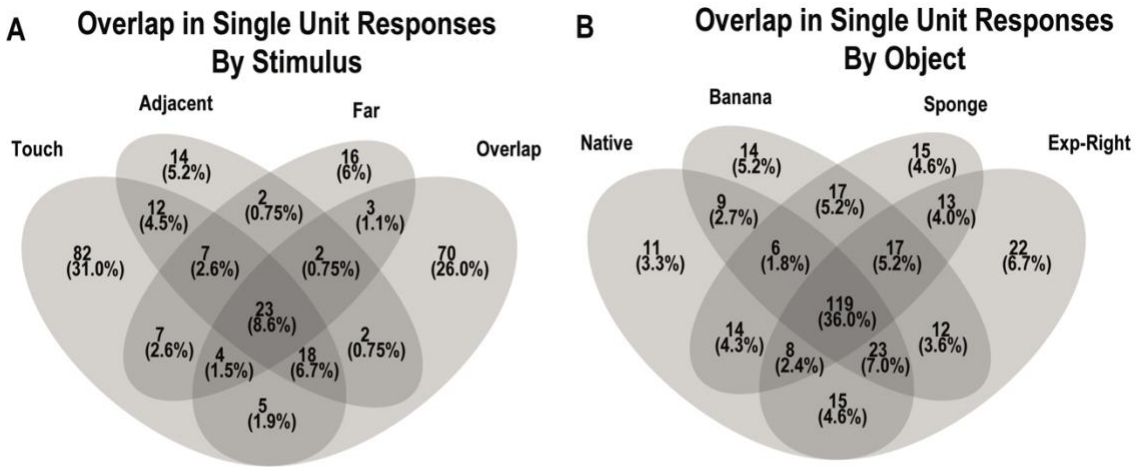

Figure 5.1-figure supplement 5. Individual neurons within the PPC population exhibit responsiveness to many diverse conditions.
A, Venn diagram illustrating the number of PPC neurons that were significantly modulated to each form of stimulus, to any object. The number in parentheses is the percentage of all significantly tuned units. B, Venn diagram illustrating the number of PPC neurons that were significantly modulated to any form of stimulus, to the four different objects. As in panel A, the number in parentheses is the percentage of all significantly tuned units. Exp-Right, experimenter's right arm.

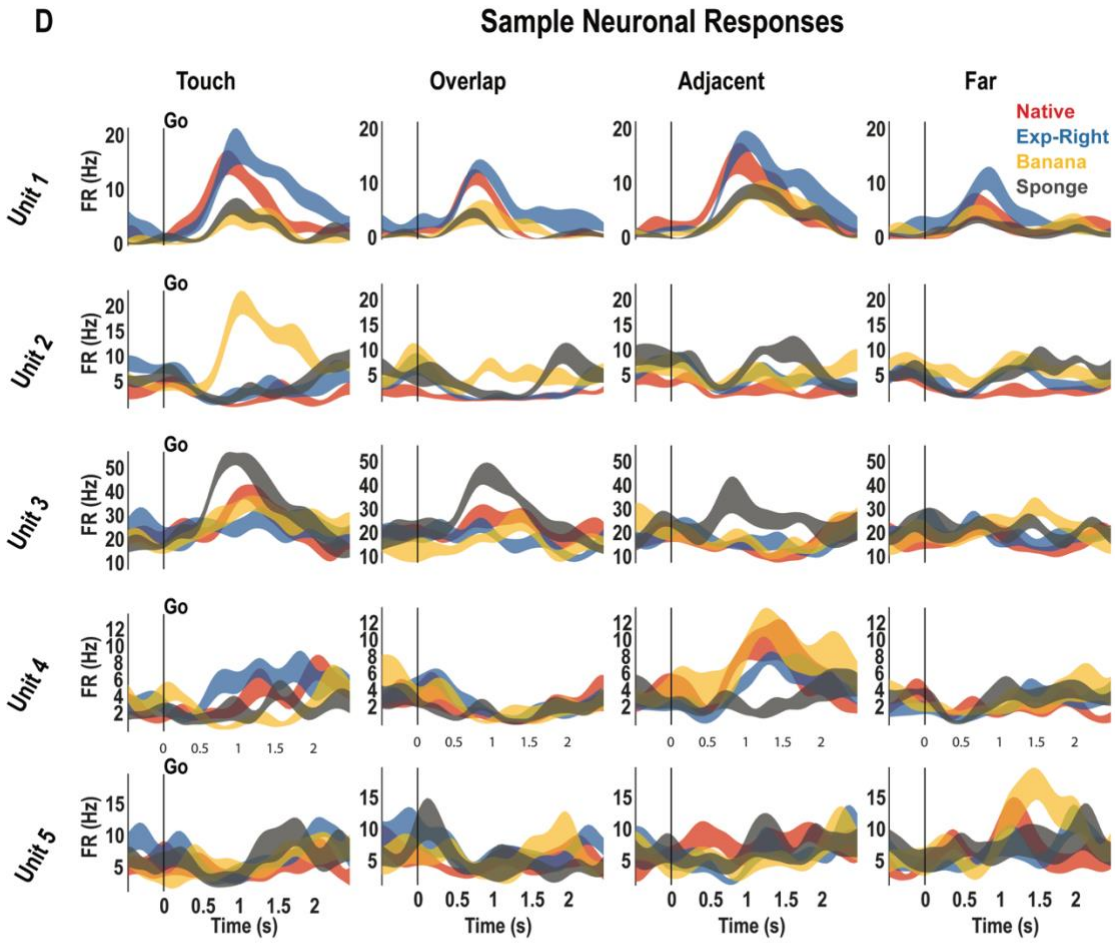

Figure 5.1-figure supplement 6. Example PPC neuronal responses to visual stimuli.

Representative neuronal responses illustrating discrimination of the various stimuli to the four objects. Each row of panels depicts the response for one neuron to touch (first column), overlap (second column), adjacent (third column), or far (fourth column) motion. Within each panel, the stimulus-response to all four objects is indicated, color-coded as in the legend. Each panel shows the neural response (mean firing rate in Hz ± standard error on the mean, n = 10 trials) as a function of time (x-axis). The vertical line labeled 'Go' indicates the start of the stimulus phase. Hz, Hertz; Exp-Right, experimenter's right hand; s, seconds.

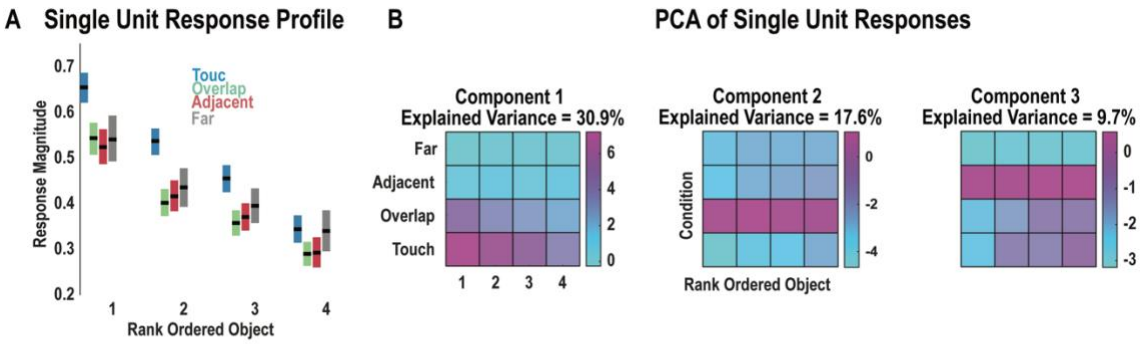

Figure 5.1-figure supplement 7. Single unit tuning profile to the four different stimuli show that touch is preferentially encoded.

A, Mean (horizontal black line) ± 95% CI (bar height) for normalized, cross-validated, rank-ordered responses (across objects) for each of the four different visual stimuli (color coded). Responses show that touch is preferentially encoded over other controls on all objects. B, First three components of a PCA performed on the pooled response across all single units. Each matrix shows the projection within the labeled component (top of panel) from each stimulus (rows) for rank-ordered objects (columns). The strength of the projection is color coded as in the scale. The fraction of population variance explained by each component is indicated above each panel.

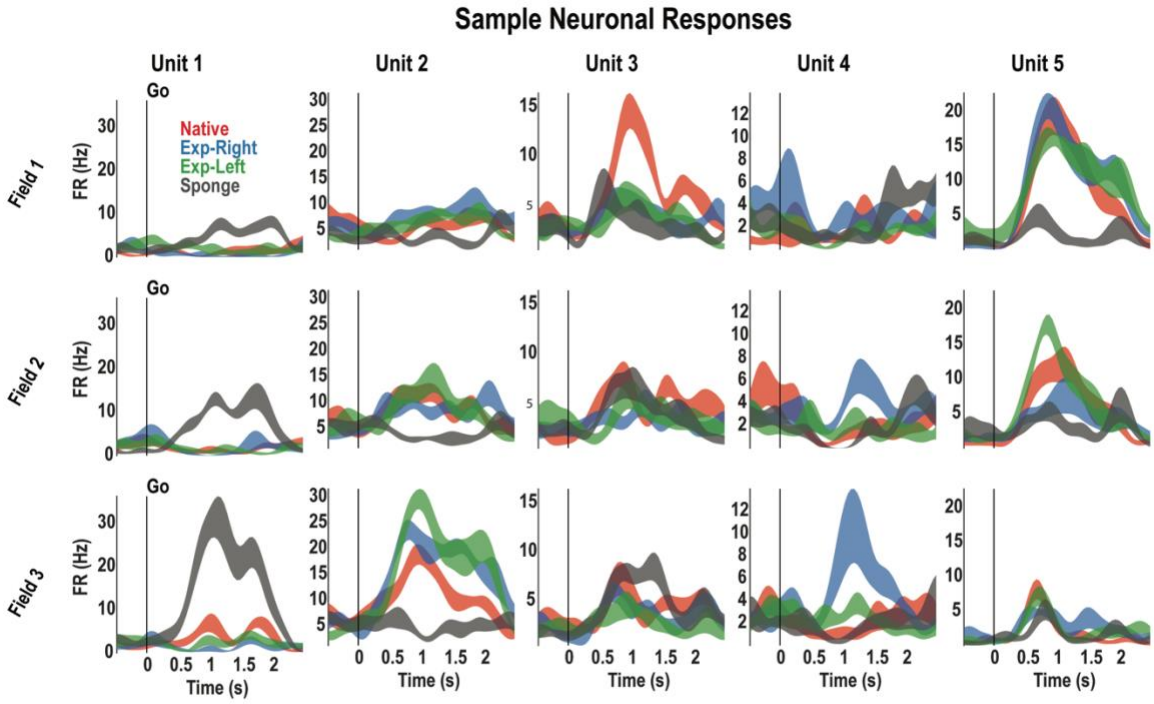

Figure 5.2-figure supplement 1. Neurons code similar object similarly yet exhibit respond broadly to multiple fields on multiple objects. *Representative neuronal responses. Each column of panels depicts the response for one neuron to seen touch to the four objects (color-coded), to Field 1-3 (in rows 1-3). Each panel shows the neural response (mean firing rate in Hz ± standard error on the mean, n = 10 trials) as a function of time (x-axis). The vertical line labeled 'Go' indicates the start of the stimulus phase.*

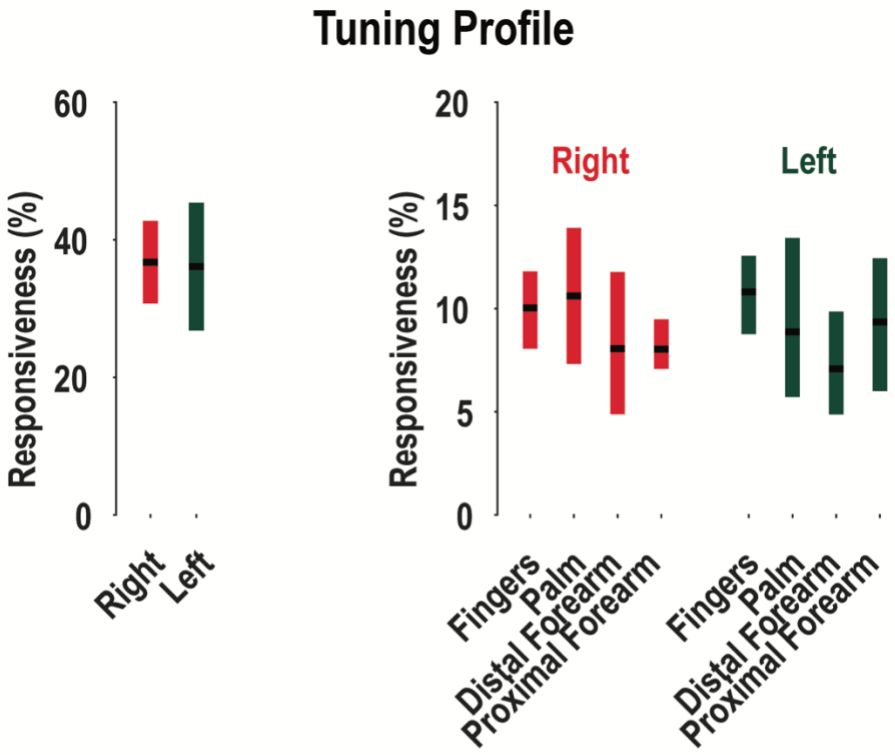

Figure 5.3-figure supplement 1. Single unit tuning profile to the four different stimuli show that touch is preferentially encoded.

Percent of the PPC neuronal population that demonstrated significant modulation relative to baseline for each tested condition ($p<0.05$, FDR corrected, $n=421$ units). Results are shown as the mean percentage (horizontal black line) of the population ± bootstrapped 95% confidence interval (CI; bar height). The left panel shows the overall tuning profile, whereas the right panel shows the breakdown by region of modulation.

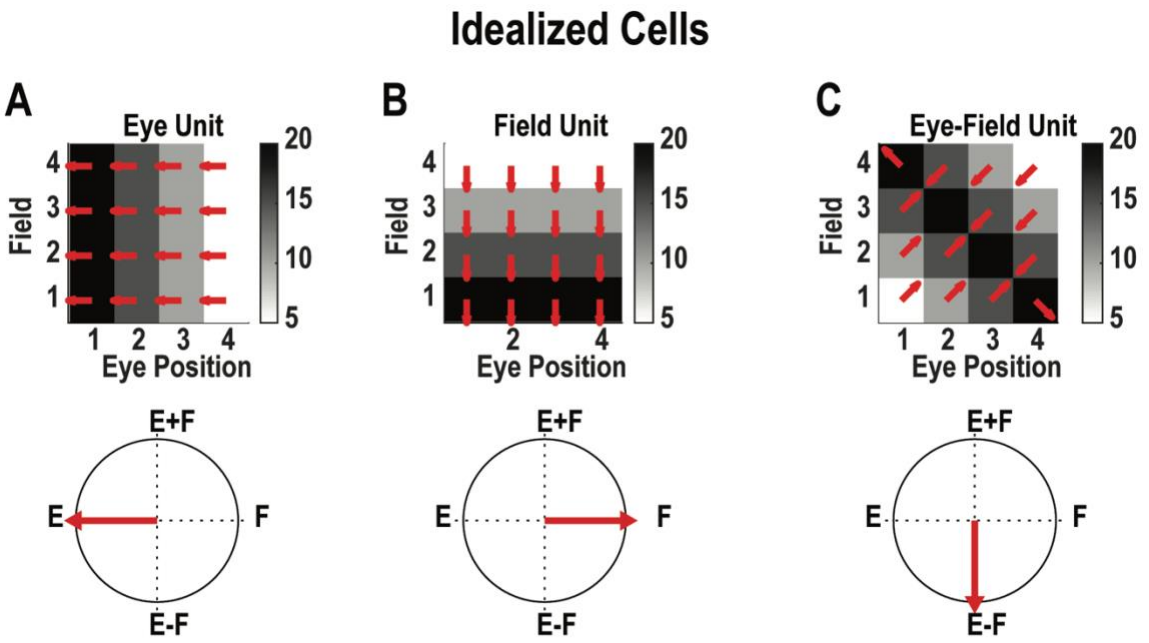

Figure 5.4-figure supplement 1. Idealized single unit responses illustrating the different ways seen touch may be represented.

Across panels A-C, response matrices (top row) for idealized units showing eye-centered coding (panel A), field-centric coding (panel B), and eye-field coding (panel C). X-axis and y-axis labels indicate possible values for each variable. Gradients estimated using MATLAB are shown in red arrows. Arrow length and direction indicate the relative importance of each variable to the unit's firing rate. More details are in Buneo et al., 2002. Summing the resultant angles over the response matrix generates the unit's reference frame, as in the circular plots (panels A-C; bottom row) for the corresponding sample unit (in the same column). Here, 0° represents field-centric coding, 180° (eye-centric) and +90° and -90° degrees represent combination eye-field coding.

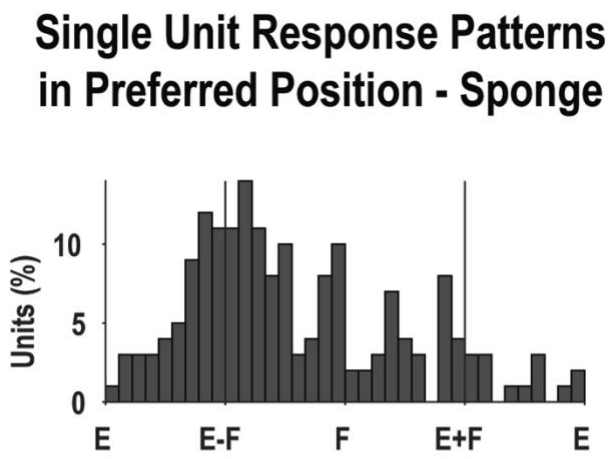

Figure 5.4-figure supplement 2. On the sponge too, seen touch is encoded in a mixture of reference frames, but mostly relative to eye position.

BPPC population reference frames for the sponge in the two labeled positions (as in Fig 4A). For each unit, gradients, and overall resultant angles were computed as described. The distribution of angles for the preferred position is plotted as a histogram. Any unit with a significant modulation (by linear regression) to touch in any condition was included. The number of such significantly modulated units is indicated as a percentage of the recorded population. The different types of reference frames are indicated: E (encoding eye position only), E-F and E+F (encoding touch as a function of eye position), and F (encoding touch to a particular field only).

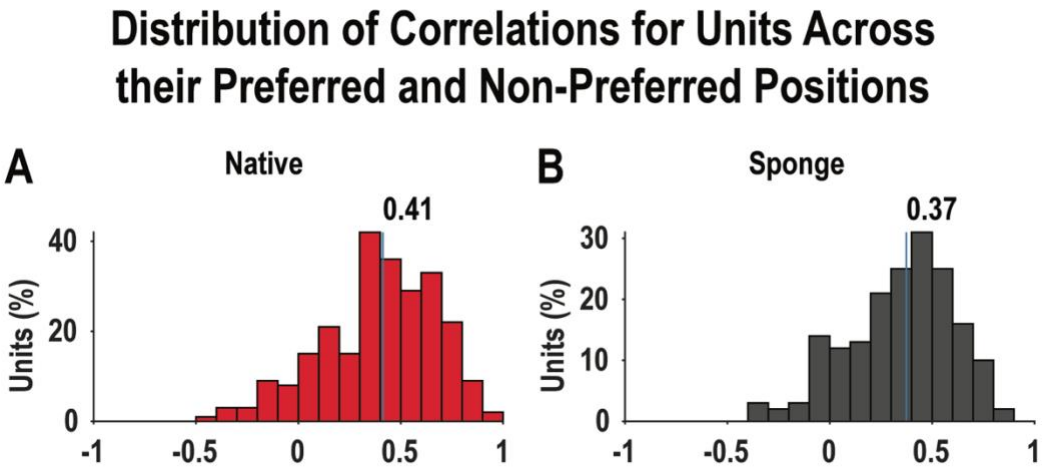

Figure 5.4-figure supplement 3. Neural information regarding seen touch generalizes across preferred and non-preferred positions.

A, Correlation across all units (y-axis; number as a percentage) demonstrating significant modulation to seen touch in reference frame encoding across preferred and non-preferred positions. The correlation strength (r) is shown on the x-axis. This panel displays correlation only for the Native arm (subject's insensate right arm). The labeled vertical black line indicates the median correlation coefficient across the population. B, As in panel A, but for the sponge.

Chapter 6

CONCLUSION

This dissertation is a novel, neuron-level demonstration that shared populations of neurons in human PPC support many aspects of sensory cognition. Previous studies from the Andersen lab showed that this substrate encoded multiple variables related to motor cognition including ***executed*** movements, ***attempted*** movements, ***observed*** actions and the ***semantic*** processing of action verbs.[2-5] Here, we build upon this work, extending it into the sensory domain. We show that the same substrate also encodes ***actual*** touch and ***imagined*** touch to oneself, ***observed*** touch to others, as well as ***observed*** touch to objects in our environment. As in the motor domain, individual neurons exhibit highly variable responses.[2] At the population-level, however, we find that the diverse touch conditions are explained by a small number of subspaces (meaningful groupings of neurons) that encode basic-level, elemental information such as touch location, and touch type. This suggests a *compositional* basis in PPC, such that various touch conditions are encoded through diverse combinations of common primitive elements. Moreover, these subspaces are *generalizable*, able to explain novel (held out) data. These principles of compositionality and generalizability suggest a basis by which PPC may support cognitive behaviors, even in situations that extend beyond our experiences.

Understanding the neuron-level basis for behavior has, to date, been challenged by most conventional research modalities to date. On the one hand, human neuroimaging studies suggest that multiple behaviors often activate shared cortex, including within PPC, but cannot clarify the neuron-level mechanisms that underlie such shared cortical activation.[36-39,115] On the other hand, neurophysiological studies in animal models offer high (neuron-level) resolution but are limited in the scope of behavior that can be examined, especially cognitive behavior.[30-33] Our current work takes advantage of recent technological advances that allow examining populations of human neurons at the high resolution of single units over multiple sessions, an ability that has not always been available.[1] We showed that shared populations of neurons are similarly engaged during many forms of touch, whether ***actual***, ***imagined*** or ***seen***. The principles of compositionality and generalizability we demonstrate provide a way by which populations of neurons in PPC may flexibly be repurposed for diverse cognitive forms. We speculate that our findings might represent the neural implementation of a system of cognitive internal models.[6,14] Consistent with such a view, this substrate supports not only sensory cognition (as we showed here), but also motor cognition (as in previous studies from the Andersen lab).[2-5] Moreover, these diverse forms of sensory and motor cognition share similar neural information organized in similar manners. In further support that our findings might reflect internal models in PPC, we found that touch is encoded not only to oneself or to other humans, but also to inanimate objects. In other words, neural information regarding variables encoded can be repurposed to a variety of behavioral contexts.

6.1 Future directions

Our work opens several exciting avenues for future experimentation and exploration. These range from further understanding the neuron-level underpinnings of cognition, to clinical applications and interdisciplinary efforts that can help countless individuals afflicted with various conditions. A few possible extensions of this work are outlined below.

6.1.1. Cognitive neurophysiology. In this work, we studied the functional organization of neural information within a population of neurons in human PPC that supports diverse behaviors. We found that population representations appear consistent with a compositional and generalizable framework of basic-level, primitive elemental information. A natural question that arises is: is such a functional organization widespread across cortex? In other words, do other high-level brain regions (such as prefrontal cortex) follow a similar informational framework? Moreover, we speculated that our findings might be consistent with a framework of internal models implemented in PPC. Do other high-level brain regions also implement internal models? Or do they access those internal models in PPC? In either case, how do different brain regions coordinate and coherently function to enable our behavioral repertoire? Many brain regions (including high-level regions) appear specialized for certain behaviors over others. For example, PPC appears to support many sensory-motor behaviors, but face recognition or place recognition appears to be housed in the temporal lobe. If all high-level brain regions encode many variables, and can support diverse behaviors, what provides them their functional specialization? Another line of questions concerns how compositionality and generalizability might lie at the basis of dreams, or memory. This work has immediate clinical relevance because disorders of such processes could lead to pathologies such, potentially such as post-traumatic stress disorder through impaired (or exaggerated) generalization of memories (especially fear memories), sleep disorders, or many other neuropsychiatric conditions. In short, understanding the answers to these questions at the neuron-level will lead to a previously unavailable, comprehensive understanding of the brain basis for human behavior.

6.1.2 Translational science. The work presented here has tremendous clinical significance. One example of a potential avenue is in the BMI arena. BMIs are devices that record neural activity from the brain, decode the activity and utilize that activity in a prosthetic device. Examples of BMIs might be driving a cursor on a screen through brain activity or moving a robotic arm through brain activity.[3,49,50] Such prosthetics could be life changing for patients with paralysis, for example. Within this space, our work in which we understand better how brain regions encode different behaviors could inform better decoder design. Knowing that neural information is organized in a compositional manner, for example, might enable decoders that can make use of this structure not only in making sense of motor signals (e.g., during movement by a robotic arm), but also in sensory processing (e.g., when the robotic arm hits an obstacle). Moreover, understanding how information is represented within a substrate can enable more efficiently transferring skills learned during one BMI application (such as a robotic limb) to another (such as the cursor-on-screen). Such BMI implementations of our findings could tremendously increase the efficiency and practicality of BMI design. Extending this line of reasoning further, BMIs not only provide a means to

augment or to restore function to patients, but also provide an invaluable way to study how populations of neurons in human brains learn, and how this learning is represented. Another potential future direction is to implement our findings *in silico*. A neural network model that mimics human-like cognition can provide a powerful paradigm in which to study how cognition unfolds. It provides an invaluable mechanism to study how pathologies (like strokes) might cause perturbations in the network, leading to impairments of cognition. To close the loop, it can finally provide a means to study how rehabilitation (such as learning) can affect the network, hopefully restoring lost abilities.

Taken together the work presented here not only enriches our understanding of human brain function, offering insight into the ways in which our remarkable repertoire is enabled, but also has broad applications for the successful implementation of future brain machine interfaces, decoding algorithms, and neuro-prosthetics.

BIBLIOGRAPHY

1 Saxena, S. & Cunningham, J. P. Towards the neural population doctrine. *Curr Opin Neurobiol* **55**, 103-111, doi:10.1016/j.conb.2019.02.002 (2019).

2 Aflalo, T., Zhang, C. Y., Rosario, E., Pouratian, N., Orban, G. A., Andersen, R. A. A shared neural substrate for action verbs and observed actions in human posterior parietal cortex. *BioRxiv*, doi:10.1101/2020.04.20.039529 (2020).

3 Aflalo, T. *et al.* Decoding motor imagery from the posterior parietal cortex of a tetraplegic human. *Science* **348**, 906-910, doi:10.1126/science.aaa5417 (2015).

4 Zhang, C. Y. *et al.* Partially Mixed Selectivity in Human Posterior Parietal Association Cortex. *Neuron* **95**, 697-708 e694, doi:10.1016/j.neuron.2017.06.040 (2017).

5 Chivukula, S. *et al.* Neural Encoding of Felt and Imagined Touch Within Human Posterior Parietal Cortex. *BioRxiv* (2020).

6 Lake, B. M., Ullman, T. D., Tenenbaum, J. B. & Gershman, S. J. Building machines that learn and think like people. *Behav Brain Sci* **40**, e253, doi:10.1017/S0140525X16001837 (2017).

7 Perry, A. *et al.* Mirroring in the Human Brain: Deciphering the Spatial-Temporal Patterns of the Human Mirror Neuron System. *Cereb Cortex* **28**, 1039-1048, doi:10.1093/cercor/bhx013 (2018).

8 Blakemore, S. J., Bristow, D., Bird, G., Frith, C. & Ward, J. Somatosensory activations during the observation of touch and a case of vision-touch synaesthesia. *Brain* **128**, 1571-1583, doi:10.1093/brain/awh500 (2005).

9 Culham, J. C. *et al.* Visually guided grasping produces fMRI activation in dorsal but not ventral stream brain areas. *Exp Brain Res* **153**, 180-189, doi:10.1007/s00221-003-1591-5 (2003).

10 Lotze, M. & Halsband, U. Motor imagery. *J Physiol Paris* **99**, 386-395, doi:10.1016/j.jphysparis.2006.03.012 (2006).

11 McManus, J. N., Li, W. & Gilbert, C. D. Adaptive shape processing in primary visual cortex. *Proc Natl Acad Sci U S A* **108**, 9739-9746, doi:10.1073/pnas.1105855108 (2011).

12 Carrillo, M. *et al.* Emotional Mirror Neurons in the Rat's Anterior Cingulate Cortex. *Curr Biol* **29**, 2104, doi:10.1016/j.cub.2019.05.064 (2019).

13 Rizzolatti, G. *et al.* Functional organization of inferior area 6 in the macaque monkey. II. Area F5 and the control of distal movements. *Exp Brain Res* **71**, 491-507, doi:10.1007/BF00248742 (1988).

14 Schwettmann, S., Tenenbaum, J. B. & Kanwisher, N. Invariant representations of mass in the human brain. *Elife* **8**, doi:10.7554/eLife.46619 (2019).

15 Andersen, R. A. & Cui, H. Intention, action planning, and decision making in parietal-frontal circuits. *Neuron* **63**, 568-583, doi:10.1016/j.neuron.2009.08.028 (2009).

16 Holmes, G. Disturbances of Visual Orientation. *Br J Ophthalmol* **2**, 449-468, doi:10.1136/bjo.2.9.449 (1918).

17 Mountcastle, V. B., Lynch, J. C., Georgopoulos, A., Sakata, H. & Acuna, C. Posterior parietal association cortex of the monkey: command functions for operations within

extrapersonal space. *J Neurophysiol* **38**, 871-908, doi:10.1152/jn.1975.38.4.871 (1975).

18 Mulliken, G. H., Musallam, S. & Andersen, R. A. Forward estimation of movement state in posterior parietal cortex. *Proc Natl Acad Sci U S A* **105**, 8170-8177, doi:10.1073/pnas.0802602105 (2008).

19 Mulliken, G. H., Musallam, S. & Andersen, R. A. Decoding trajectories from posterior parietal cortex ensembles. *J Neurosci* **28**, 12913-12926, doi:10.1523/JNEUROSCI.1463-08.2008 (2008).

20 Andersen, R. A., Aflalo, T. & Kellis, S. From thought to action: The brain-machine interface in posterior parietal cortex. *Proc Natl Acad Sci U S A*, doi:10.1073/pnas.1902276116 (2019).

21 Andersen, R. A. & Buneo, C. A. Intentional maps in posterior parietal cortex. *Annu Rev Neurosci* **25**, 189-220, doi:10.1146/annurev.neuro.25.112701.142922 (2002).

22 Brang, D., Taich, Z. J., Hillyard, S. A., Grabowecky, M. & Ramachandran, V. S. Parietal connectivity mediates multisensory facilitation. *Neuroimage* **78**, 396-401, doi:10.1016/j.neuroimage.2013.04.047 (2013).

23 Chivukula, S., Jafari, M., Aflalo, T., Yong, N. A. & Pouratian, N. Cognition in Sensorimotor Control: Interfacing With the Posterior Parietal Cortex. *Front Neurosci* **13**, 140, doi:10.3389/fnins.2019.00140 (2019).

24 Rizzolatti, G. & Craighero, L. The mirror-neuron system. *Annu Rev Neurosci* **27**, 169-192, doi:10.1146/annurev.neuro.27.070203.144230 (2004).

25 Rizzolatti, G., Fogassi, L. & Gallese, V. Neurophysiological mechanisms underlying the understanding and imitation of action. *Nat Rev Neurosci* **2**, 661-670, doi:10.1038/35090060 (2001).

26 Graziano, M. S. & Webb, T. W. The attention schema theory: a mechanistic account of subjective awareness. *Front Psychol* **6**, 500, doi:10.3389/fpsyg.2015.00500 (2015).

27 Mukamel, R., Ekstrom, A. D., Kaplan, J., Iacoboni, M. & Fried, I. Single-neuron responses in humans during execution and observation of actions. *Curr Biol* **20**, 750-756, doi:10.1016/j.cub.2010.02.045 (2010).

28 Bastiaansen, J. A., Thioux, M. & Keysers, C. Evidence for mirror systems in emotions. *Philos Trans R Soc Lond B Biol Sci* **364**, 2391-2404, doi:10.1098/rstb.2009.0058 (2009).

29 Hickok, G. Eight problems for the mirror neuron theory of action understanding in monkeys and humans. *J Cogn Neurosci* **21**, 1229-1243, doi:10.1162/jocn.2009.21189 (2009).

30 Graziano, M. S. Where is my arm? The relative role of vision and proprioception in the neuronal representation of limb position. *Proc Natl Acad Sci U S A* **96**, 10418-10421, doi:10.1073/pnas.96.18.10418 (1999).

31 Graziano, M. S. A system of multimodal areas in the primate brain. *Neuron* **29**, 4-6 (2001).

32 Graziano, M. S., Cooke, D. F. & Taylor, C. S. Coding the location of the arm by sight. *Science* **290**, 1782-1786, doi:10.1126/science.290.5497.1782 (2000).

33 Graziano, M. S. & Gross, C. G. A bimodal map of space: somatosensory receptive fields in the macaque putamen with corresponding visual receptive fields. *Exp Brain Res* **97**, 96-109 (1993).

34 Holmes, N. P. & Spence, C. The body schema and the multisensory representation(s) of peripersonal space. *Cogn Process* **5**, 94-105, doi:10.1007/s10339-004-0013-3 (2004).

35 Ishida, H., Nakajima, K., Inase, M. & Murata, A. Shared mapping of own and others' bodies in visuotactile bimodal area of monkey parietal cortex. *J Cogn Neurosci* **22**, 83-96, doi:10.1162/jocn.2009.21185 (2010).

36 Keysers, C. & Gazzola, V. Expanding the mirror: vicarious activity for actions, emotions, and sensations. *Curr Opin Neurobiol* **19**, 666-671, doi:10.1016/j.conb.2009.10.006 (2009).

37 Keysers, C. & Gazzola, V. Dissociating the ability and propensity for empathy. *Trends Cogn Sci* **18**, 163-166, doi:10.1016/j.tics.2013.12.011 (2014).

38 Keysers, C., Kaas, J. H. & Gazzola, V. Somatosensation in social perception. *Nat Rev Neurosci* **11**, 417-428, doi:10.1038/nrn2833 (2010).

39 Keysers, C. *et al.* A touching sight: SII/PV activation during the observation and experience of touch. *Neuron* **42**, 335-346 (2004).

40 Jeannerod, M. & Decety, J. Mental motor imagery: a window into the representational stages of action. *Curr Opin Neurobiol* **5**, 727-732, doi:10.1016/0959-4388(95)80099-9 (1995).

41 Andersen, R. A., Hwang, E. J. & Mulliken, G. H. Cognitive neural prosthetics. *Annu Rev Psychol* **61**, 169-190, C161-163, doi:10.1146/annurev.psych.093008.100503 (2010).

42 Rigotti, M. *et al.* The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585-590, doi:10.1038/nature12160 (2013).

43 Zhang, C. Y. *et al.* Preservation of Partially Mixed Selectivity in Human Posterior Parietal Cortex across Changes in Task Context. *eNeuro* **7**, doi:10.1523/ENEURO.0222-19.2019 (2020).

44 Liu, J., Harris, A. & Kanwisher, N. Perception of face parts and face configurations: an FMRI study. *J Cogn Neurosci* **22**, 203-211, doi:10.1162/jocn.2009.21203 (2010).

45 Pilgramm, S. *et al.* Motor imagery of hand actions: Decoding the content of motor imagery from brain activity in frontal and parietal motor areas. *Hum Brain Mapp* **37**, 81-93, doi:10.1002/hbm.23015 (2016).

46 Sugarman, M. A. *et al.* Functional magnetic resonance imaging of semantic memory as a presymptomatic biomarker of Alzheimer's disease risk. *Biochim Biophys Acta* **1822**, 442-456, doi:10.1016/j.bbadis.2011.09.016 (2012).

47 Buzsaki, G., Anastassiou, C. A. & Koch, C. The origin of extracellular fields and currents--EEG, ECoG, LFP and spikes. *Nat Rev Neurosci* **13**, 407-420, doi:10.1038/nrn3241 (2012).

48 Jarosiewicz, B. *et al.* Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface. *Sci Transl Med* **7**, 313ra179, doi:10.1126/scitranslmed.aac7328 (2015).

49 Leuthardt, E. C., Schalk, G., Wolpaw, J. R., Ojemann, J. G. & Moran, D. W. A brain-computer interface using electrocorticographic signals in humans. *J Neural Eng* **1**, 63-71, doi:10.1088/1741-2560/1/2/001 (2004).

50 Pandarinath, C. *et al.* Neural population dynamics in human motor cortex during movements in people with ALS. *Elife* **4**, e07436, doi:10.7554/eLife.07436 (2015).

51 de Haan, E. H. F. & Dijkerman, H. C. Somatosensation in the Brain: A Theoretical Re-evaluation and a New Model. *Trends Cogn Sci* **24**, 529-541, doi:10.1016/j.tics.2020.04.003 (2020).

52 de Lafuente, V. & Romo, R. Neural correlate of subjective sensory experience gradually builds up across cortical areas. *Proc Natl Acad Sci U S A* **103**, 14266-14271, doi:10.1073/pnas.0605826103 (2006).

53 Avillac, M., Ben Hamed, S. & Duhamel, J. R. Multisensory integration in the ventral intraparietal area of the macaque monkey. *J Neurosci* **27**, 1922-1932, doi:10.1523/JNEUROSCI.2646-06.2007 (2007).

54 Hwang, E. J., Hauschild, M., Wilke, M. & Andersen, R. A. Spatial and temporal eye-hand coordination relies on the parietal reach region. *J Neurosci* **34**, 12884-12892, doi:10.1523/JNEUROSCI.3719-13.2014 (2014).

55 Seelke, A. M. *et al.* Topographic Maps within Brodmann's Area 5 of macaque monkeys. *Cereb Cortex* **22**, 1834-1850, doi:10.1093/cercor/bhr257 (2012).

56 Sereno, M. I. & Huang, R. S. Multisensory maps in parietal cortex. *Curr Opin Neurobiol* **24**, 39-46, doi:10.1016/j.conb.2013.08.014 (2014).

57 Chan, A. W. & Baker, C. I. Seeing is not feeling: posterior parietal but not somatosensory cortex engagement during touch observation. *J Neurosci* **35**, 1468-1480, doi:10.1523/JNEUROSCI.3621-14.2015 (2015).

58 Lucas, M. V., Anderson, L. C., Bolling, D. Z., Pelphrey, K. A. & Kaiser, M. D. Dissociating the Neural Correlates of Experiencing and Imagining Affective Touch. *Cerebral Cortex* **25**, 2623-2630, doi:10.1093/cercor/bhu061 (2015).

59 Rutishauser, U., Aflalo, T., Rosario, E. R., Pouratian, N. & Andersen, R. A. Single-Neuron Representation of Memory Strength and Recognition Confidence in Left Human Posterior Parietal Cortex. *Neuron* **97**, 209-220 e203, doi:10.1016/j.neuron.2017.11.029 (2018).

60 Sakata, H., Takaoka, Y., Kawarasaki, A. & Shibutani, H. Somatosensory properties of neurons in the superior parietal cortex (area 5) of the rhesus monkey. *Brain Res* **64**, 85-102 (1973).

61 Huang, R. S., Chen, C. F. & Sereno, M. I. Spatiotemporal integration of looming visual and tactile stimuli near the face. *Human Brain Mapping* **39**, 2156-2176, doi:10.1002/hbm.23995 (2018).

62 Sereno, M. I. & Huang, R. S. Multisensory maps in parietal cortex. *Current Opinion in Neurobiology* **24**, 39-46, doi:10.1016/j.conb.2013.08.014 (2014).

63 Huang, R. S., Chen, C. F., Tran, A. T., Holstein, K. L. & Sereno, M. I. Mapping multisensory parietal face and body areas in humans. *Proceedings of the National Academy of Sciences of the United States of America* **109**, 18114-18119, doi:10.1073/pnas.1207946109 (2012).

64 Caramazza, A., Anzellotti, S., Strnad, L. & Lingnau, A. in *Annual Review of Neuroscience* Vol. 37 *Annual Review of Neuroscience* (ed S. E. Hyman) 1-15 (2014).

65 Gallivan, J. P., McLean, D. A., Valyear, K. F., Pettypiece, C. E. & Culham, J. C. Decoding action intentions from preparatory brain activity in human parieto-frontal networks. *J Neurosci* **31**, 9599-9610, doi:10.1523/JNEUROSCI.0080-11.2011 (2011).

66 Gallivan, J. P., McLean, D. A., Smith, F. W. & Culham, J. C. Decoding effector-dependent and effector-independent movement intentions from human parieto-frontal brain activity. *J Neurosci* **31**, 17149-17168, doi:10.1523/JNEUROSCI.1058-11.2011 (2011).

67 Brainard, D. H. The Psychophysics Toolbox. *Spat Vis* **10**, 433-436 (1997).

68 Harris, K. D., Quiroga, R. Q., Freeman, J. & Smith, S. L. Improving data quality in neuronal population recordings. *Nat Neurosci* **19**, 1165-1174, doi:10.1038/nn.4365 (2016).

69 Smialowski, P., Frishman, D. & Kramer, S. Pitfalls of supervised feature selection. *Bioinformatics* **26**, 440-443, doi:10.1093/bioinformatics/btp621 (2010).

70 Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat Neurosci* **17**, 1500-1509, doi:10.1038/nn.3776 (2014).

71 Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience* **17**, 1500-1509, doi:10.1038/nn.3776 (2014).

72 Sakellaridi, S. *et al.* Intrinsic Variable Learning for Brain-Machine Interface Control by Human Anterior Intraparietal Cortex. *Neuron* **102**, 694-705 e693, doi:10.1016/j.neuron.2019.02.012 (2019).

73 Cunningham, J. P. & Ghahramani, Z. Linear Dimensionality Reduction: Survey, Insights, and Generalizations. *J Mach Learn Res* **16**, 2859-2900 (2015).

74 Rosenbaum, D. A. in *Memory and Control of Action* (ed R.A. Magill) 231–274 (1983).

75 Lecas, J. C., Requin, J., Anger, C. & Vitton, N. Changes in neuronal activity of the monkey precentral cortex during preparation for movement. *J Neurophysiol* **56**, 1680-1702, doi:10.1152/jn.1986.56.6.1680 (1986).

76 Ames, K. C., Ryu, S. I. & Shenoy, K. V. Simultaneous motor preparation and execution in a last-moment reach correction task. *Nat Commun* **10**, 2718, doi:10.1038/s41467-019-10772-2 (2019).

77 Kaas, J. H. What, if anything, is SI? Organization of first somatosensory area of cortex. *Physiol Rev* **63**, 206-231, doi:10.1152/physrev.1983.63.1.206 (1983).

78 Kaas, J. H., Nelson, R. J., Sur, M., Lin, C. S. & Merzenich, M. M. Multiple representations of the body within the primary somatosensory cortex of primates. *Science* **204**, 521-523, doi:10.1126/science.107591 (1979).

79 Sur, M., Merzenich, M. M. & Kaas, J. H. Magnification, receptive-field area, and "hypercolumn" size in areas 3b and 1 of somatosensory cortex in owl monkeys. *J Neurophysiol* **44**, 295-311, doi:10.1152/jn.1980.44.2.295 (1980).

80 Ferezou, I. *et al.* Spatiotemporal dynamics of cortical sensorimotor integration in behaving mice. *Neuron* **56**, 907-923, doi:10.1016/j.neuron.2007.10.007 (2007).

81 Geyer, S., Schormann, T., Mohlberg, H. & Zilles, K. Areas 3a, 3b, and 1 of human primary somatosensory cortex. Part 2. Spatial normalization to standard anatomical space. *Neuroimage* **11**, 684-696, doi:10.1006/nimg.2000.0548 (2000).

82 Iwamura, Y. Hierarchical somatosensory processing. *Curr Opin Neurobiol* **8**, 522-528 (1998).

83 Jiang, W., Tremblay, F. & Chapman, C. E. Neuronal encoding of texture changes in the primary and the secondary somatosensory cortical areas of monkeys during passive texture discrimination. *J Neurophysiol* **77**, 1656-1662, doi:10.1152/jn.1997.77.3.1656 (1997).

84 Ruben, J. *et al.* Somatotopic organization of human secondary somatosensory cortex. *Cereb Cortex* **11**, 463-473, doi:10.1093/cercor/11.5.463 (2001).

85 Schnitzler, A., Salmelin, R., Salenius, S., Jousmaki, V. & Hari, R. Tactile information from the human hand reaches the ipsilateral primary somatosensory cortex. *Neurosci Lett* **200**, 25-28, doi:10.1016/0304-3940(95)12065-c (1995).

86 Tame, L., Braun, C., Holmes, N. P., Farne, A. & Pavani, F. Bilateral representations of touch in the primary somatosensory cortex. *Cogn Neuropsychol* **33**, 48-66, doi:10.1080/02643294.2016.1159547 (2016).

87 Tame, L., Pavani, F., Papadelis, C., Farne, A. & Braun, C. Early integration of bilateral touch in the primary somatosensory cortex. *Hum Brain Mapp* **36**, 1506-1523, doi:10.1002/hbm.22719 (2015).

88 Burton, H. & Sinclair, R. J. Second somatosensory cortical area in macaque monkeys. I. Neuronal responses to controlled, punctate indentations of glabrous skin on the hand. *Brain Res* **520**, 262-271, doi:10.1016/0006-8993(90)91714-r (1990).

89 Dykes, R. W. Parallel processing of somatosensory information: a theory. *Brain Res* **287**, 47-115, doi:10.1016/0165-0173(83)90004-8 (1983).

90 Garraghty, P. E., Florence, S. L. & Kaas, J. H. Ablations of areas 3a and 3b of monkey somatosensory cortex abolish cutaneous responsivity in area 1. *Brain Res* **528**, 165-169, doi:10.1016/0006-8993(90)90213-u (1990).

91 Pei, Y. C., Denchev, P. V., Hsiao, S. S., Craig, J. C. & Bensmaia, S. J. Convergence of submodality-specific input onto neurons in primary somatosensory cortex. *J Neurophysiol* **102**, 1843-1853, doi:10.1152/jn.00235.2009 (2009).

92 Pons, T. P., Garraghty, P. E., Friedman, D. P. & Mishkin, M. Physiological evidence for serial processing in somatosensory cortex. *Science* **237**, 417-420, doi:10.1126/science.3603028 (1987).

93 Saal, H. P., Harvey, M. A. & Bensmaia, S. J. Rate and timing of cortical responses driven by separate sensory channels. *Elife* **4**, e10450, doi:10.7554/eLife.10450 (2015).

94 Soso, M. J. & Fetz, E. E. Responses of identified cells in postcentral cortex of awake monkeys during comparable active and passive joint movements. *J Neurophysiol* **43**, 1090-1110, doi:10.1152/jn.1980.43.4.1090 (1980).

95 Zhu, Z., Disbrow, E. A., Zumer, J. M., McGonigle, D. J. & Nagarajan, S. S. Spatiotemporal integration of tactile information in human somatosensory cortex. *BMC Neurosci* **8**, 21, doi:10.1186/1471-2202-8-21 (2007).

96 Reed, J. L. *et al.* Response properties of neurons in primary somatosensory cortex of owl monkeys reflect widespread spatiotemporal integration. *J Neurophysiol* **103**, 2139-2157, doi:10.1152/jn.00709.2009 (2010).

97 Vazquez, Y., Zainos, A., Alvarez, M., Salinas, E. & Romo, R. Neural coding and perceptual detection in the primate somatosensory thalamus. *Proc Natl Acad Sci U S A* **109**, 15006-15011, doi:10.1073/pnas.1212535109 (2012).

98 Bisley, J. W., Krishna, B. S. & Goldberg, M. E. A rapid and precise on-response in posterior parietal cortex. *J Neurosci* **24**, 1833-1838, doi:10.1523/JNEUROSCI.5007-03.2004 (2004).

99 Regev, T. I., Winawer, J., Gerber, E. M., Knight, R. T. & Deouell, L. Y. Human posterior parietal cortex responds to visual stimuli as early as peristriate occipital cortex. *Eur J Neurosci* **48**, 3567-3582, doi:10.1111/ejn.14164 (2018).

100 Rosenbaum, D. A. The movement precuing technique: assumptions, applications and extensions. *Advances in Psychology* **12**, 231-274 (1983).

101 Klaes, C. *et al.* Hand Shape Representations in the Human Posterior Parietal Cortex. *J Neurosci* **35**, 15466-15476, doi:10.1523/JNEUROSCI.2747-15.2015 (2015).

102 Binder, J. R. & Desai, R. H. The neurobiology of semantic memory. *Trends Cogn Sci* **15**, 527-536, doi:10.1016/j.tics.2011.10.001 (2011).

103 Meyer, K. & Damasio, A. Convergence and divergence in a neural architecture for recognition and memory. *Trends Neurosci* **32**, 376-382, doi:10.1016/j.tins.2009.04.002 (2009).

104 Miyashita, Y. Perirhinal circuits for memory processing. *Nat Rev Neurosci* **20**, 577-592, doi:10.1038/s41583-019-0213-6 (2019).

105 Patterson, K., Nestor, P. J. & Rogers, T. T. Where do you know what you know? The representation of semantic knowledge in the human brain. *Nat Rev Neurosci* **8**, 976-987, doi:10.1038/nrn2277 (2007).

106 Ralph, M. A., Jefferies, E., Patterson, K. & Rogers, T. T. The neural and computational bases of semantic cognition. *Nat Rev Neurosci* **18**, 42-55, doi:10.1038/nrn.2016.150 (2017).

107 Kilteni, K., Andersson, B. J., Houborg, C. & Ehrsson, H. H. Motor imagery involves predicting the sensory consequences of the imagined movement. *Nat Commun* **9**, 1617, doi:10.1038/s41467-018-03989-0 (2018).

108 Schmidt, T. T. & Blankenburg, F. The Somatotopy of Mental Tactile Imagery. *Front Hum Neurosci* **13**, 10, doi:10.3389/fnhum.2019.00010 (2019).

109 Decety, J., Jeannerod, M., Durozard, D. & Baverel, G. Central activation of autonomic effectors during mental simulation of motor actions in man. *J Physiol* **461**, 549-563, doi:10.1113/jphysiol.1993.sp019528 (1993).

110 Decety, J., Jeannerod, M., Germain, M. & Pastene, J. Vegetative response during imagined movement is proportional to mental effort. *Behav Brain Res* **42**, 1-5, doi:10.1016/s0166-4328(05)80033-6 (1991).

111 Decety, J., Jeannerod, M. & Prablanc, C. The timing of mentally represented actions. *Behav Brain Res* **34**, 35-42, doi:10.1016/s0166-4328(89)80088-9 (1989).

112 Decety, J. & Michel, F. Comparative analysis of actual and mental movement times in two graphic tasks. *Brain Cogn* **11**, 87-97, doi:10.1016/0278-2626(89)90007-9 (1989).

113 Papaxanthis, C., Pozzo, T., Skoura, X. & Schieppati, M. Does order and timing in performance of imagined and actual movements affect the motor imagery process? The duration of walking and writing task. *Behav Brain Res* **134**, 209-215, doi:10.1016/s0166-4328(02)00030-x (2002).

114 Papaxanthis, C., Schieppati, M., Gentili, R. & Pozzo, T. Imagined and actual arm movements have similar durations when performed under different conditions of direction and mass. *Exp Brain Res* **143**, 447-452, doi:10.1007/s00221-002-1012-1 (2002).

115 Lucas, M. V., Anderson, L. C., Bolling, D. Z., Pelphrey, K. A. & Kaiser, M. D. Dissociating the Neural Correlates of Experiencing and Imagining Affective Touch. *Cereb Cortex* **25**, 2623-2630, doi:10.1093/cercor/bhu061 (2015).

116 Wise, N. J., Frangos, E. & Komisaruk, B. R. Activation of sensory cortex by imagined genital stimulation: an fMRI analysis. *Socioaffect Neurosci Psychol* **6**, 31481, doi:10.3402/snp.v6.31481 (2016).

117 Huth, A. G., de Heer, W. A., Griffiths, T. L., Theunissen, F. E. & Gallant, J. L. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature* **532**, 453-458, doi:10.1038/nature17637 (2016).

118 Martin, A. GRAPES-Grounding representations in action, perception, and emotion systems: How object properties and categories are represented in the human brain. *Psychon Bull Rev* **23**, 979-990, doi:10.3758/s13423-015-0842-3 (2016).

119 Pulvermuller, F. How neurons make meaning: brain mechanisms for embodied and abstract-symbolic semantics. *Trends Cogn Sci* **17**, 458-470, doi:10.1016/j.tics.2013.06.004 (2013).

120 Johansen-Berg, H., Christensen, V., Woolrich, M. & Matthews, P. M. Attention to touch modulates activity in both primary and secondary somatosensory areas. *Neuroreport* **11**, 1237-1241, doi:10.1097/00001756-200004270-00019 (2000).

121 Hamalainen, H., Hiltunen, J. & Titievskaja, I. fMRI activations of SI and SII cortices during tactile stimulation depend on attention. *Neuroreport* **11**, 1673-1676, doi:10.1097/00001756-200006050-00016 (2000).

122 Puckett, A. M., Bollmann, S., Barth, M. & Cunnington, R. Measuring the effects of attention to individual fingertips in somatosensory cortex using ultra-high field (7T) fMRI. *Neuroimage* **161**, 179-187, doi:10.1016/j.neuroimage.2017.08.014 (2017).

123 Roland, P. E. Somatotopical tuning of postcentral gyrus during focal attention in man. A regional cerebral blood flow study. *J Neurophysiol* **46**, 744-754, doi:10.1152/jn.1981.46.4.744 (1981).

124 Boynton, G. M. A framework for describing the effects of attention on visual responses. *Vision Res* **49**, 1129-1143, doi:10.1016/j.visres.2008.11.001 (2009).

125 Williford, T. & Maunsell, J. H. Effects of spatial attention on contrast response functions in macaque area V4. *J Neurophysiol* **96**, 40-54, doi:10.1152/jn.01207.2005 (2006).

126 Snyder, A. C., Yu, B. M. & Smith, M. A. Distinct population codes for attention in the absence and presence of visual stimulation. *Nat Commun* **9**, 4382, doi:10.1038/s41467-018-06754-5 (2018).
127 Tandon, S., Kambi, N., Lazar, L., Mohammed, H. & Jain, N. Large-scale expansion of the face representation in somatosensory areas of the lateral sulcus after spinal cord injuries in monkeys. *J Neurosci* **29**, 12009-12019, doi:10.1523/JNEUROSCI.2118-09.2009 (2009).
128 Mohammed, H. & Hollis, E. R., 2nd. Cortical Reorganization of Sensorimotor Systems and the Role of Intracortical Circuits After Spinal Cord Injury. *Neurotherapeutics* **15**, 588-603, doi:10.1007/s13311-018-0638-z (2018).
129 Kikkert, S. *et al.* Revealing the neural fingerprints of a missing hand. *Elife* **5**, doi:10.7554/eLife.15292 (2016).
130 Makin, T. R. & Flor, H. Brain (re)organisation following amputation: Implications for phantom limb pain. *Neuroimage* **218**, 116943, doi:10.1016/j.neuroimage.2020.116943 (2020).
131 Makin, T. R. & Bensmaia, S. J. Stability of Sensory Topographies in Adult Cortex. *Trends Cogn Sci* **21**, 195-204, doi:10.1016/j.tics.2017.01.002 (2017).
132 Armenta Salas, M. *et al.* Proprioceptive and cutaneous sensations in humans elicited by intracortical microstimulation. *Elife* **7**, doi:10.7554/eLife.32904 (2018).
133 Flesher, S. N. *et al.* Intracortical microstimulation of human somatosensory cortex. *Sci Transl Med* **8**, 361ra141, doi:10.1126/scitranslmed.aaf8083 (2016).
134 Wang, X., Men, W., Gao, J., Caramazza, A. & Bi, Y. Two Forms of Knowledge Representations in the Human Brain. *Neuron* **107**, 383-393 e385, doi:10.1016/j.neuron.2020.04.010 (2020).
135 Lake, B. M., Salakhutdinov, R. & Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science* **350**, 1332-1338, doi:10.1126/science.aab3050 (2015).
136 Orban, G. A., Lanzilotto, M. & Bonini, L. From Observed Action Identity to Social Affordances. *Trends Cogn Sci* **25**, 493-505, doi:10.1016/j.tics.2021.02.012 (2021).
137 Kobak, D. *et al.* Demixed principal component analysis of neural population data. *Elife* **5**, doi:10.7554/eLife.10989 (2016).
138 Churchland, P. S. *Braintrust: What neuroscience tells us about morality*. (Princeton University Press, 2011).
139 Chang, L. & Tsao, D. Y. The Code for Facial Identity in the Primate Brain. *Cell* **169**, 1013-1028 e1014, doi:10.1016/j.cell.2017.05.011 (2017).
140 Iwamura, Y., Iriki, A. & Tanaka, M. Bilateral hand representation in the postcentral somatosensory cortex. *Nature* **369**, 554-556, doi:10.1038/369554a0 (1994).
141 Makin, T. R., Holmes, N. P. & Ehrsson, H. H. On the other hand: dummy hands and peripersonal space. *Behav Brain Res* **191**, 1-10, doi:10.1016/j.bbr.2008.02.041 (2008).
142 Makin, T. R., Holmes, N. P. & Zohary, E. Is that near my hand? Multisensory representation of peripersonal space in human intraparietal sulcus. *J Neurosci* **27**, 731-740, doi:10.1523/JNEUROSCI.3653-06.2007 (2007).

143 Pascual-Leone, A. & Hamilton, R. The metamodal organization of the brain. *Prog Brain Res* **134**, 427-445, doi:10.1016/s0079-6123(01)34028-1 (2001).

144 Delhaye, B. P., Long, K. H. & Bensmaia, S. J. Neural Basis of Touch and Proprioception in Primate Cortex. *Compr Physiol* **8**, 1575-1602, doi:10.1002/cphy.c170033 (2018).

145 Duhamel, J. R., Colby, C. L. & Goldberg, M. E. Ventral intraparietal area of the macaque: congruent visual and somatic response properties. *J Neurophysiol* **79**, 126-136, doi:10.1152/jn.1998.79.1.126 (1998).

146 Iwamura, Y., Tanaka, M., Iriki, A., Taoka, M. & Toda, T. Processing of tactile and kinesthetic signals from bilateral sides of the body in the postcentral gyrus of awake monkeys. *Behav Brain Res* **135**, 185-190 (2002).

147 Pasalar, S., Ro, T. & Beauchamp, M. S. TMS of posterior parietal cortex disrupts visual tactile multisensory integration. *Eur J Neurosci* **31**, 1783-1790, doi:10.1111/j.1460-9568.2010.07193.x (2010).

148 Gardner, E. P. *et al.* Neurophysiology of prehension. I. Posterior parietal cortex and object-oriented hand behaviors. *J Neurophysiol* **97**, 387-406, doi:10.1152/jn.00558.2006 (2007).

149 Shomstein, S. Cognitive functions of the posterior parietal cortex: top-down and bottom-up attentional control. *Front Integr Neurosci* **6**, 38, doi:10.3389/fnint.2012.00038 (2012).

150 Jiang, X., Saggar, H., Ryu, S. I., Shenoy, K. V. & Kao, J. C. Structure in Neural Activity during Observed and Executed Movements Is Shared at the Neural Population Level, Not in Single Neurons. *Cell Rep* **32**, 108148, doi:10.1016/j.celrep.2020.108148 (2020).

151 Ebisch, S. J. *et al.* The sense of touch: embodied simulation in a visuotactile mirroring mechanism for observed animate or inanimate touch. *J Cogn Neurosci* **20**, 1611-1623, dois:10.1162/jocn.2008.20111 (2008).

152 Ferrari, P. F. & Rizzolatti, G. Mirror neuron research: the past and the future. *Philos Trans R Soc Lond B Biol Sci* **369**, 20130169, doi:10.1098/rstb.2013.0169 (2014).

153 Schaefer, M., Xu, B., Flor, H. & Cohen, L. G. Effects of different viewing perspectives on somatosensory activations during observation of touch. *Hum Brain Mapp* **30**, 2722-2730, doi:10.1002/hbm.20701 (2009).

154 Schaefer, M., Heinze, H. J. & Rotte, M. Embodied empathy for tactile events: Interindividual differences and vicarious somatosensory responses during touch observation. *Neuroimage* **60**, 952-957, doi:10.1016/j.neuroimage.2012.01.112 (2012).