

Neural dynamics of adaptive value computation in the human brain

Thesis by
Aniek Fransen

In Partial Fulfillment of the Requirements for the
Degree of
Social and Decision Neuroscience



CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2026
Defended May 18th 2026

© 2026

Aniek Fransen

ORCID: 0009-0004-8385-4665 

All rights reserved

ACKNOWLEDGEMENTS

This work should start with my expression of gratitude to my parents, first and foremost. Thank you both for your unwavering support, even as plans changed and an eight month internship has turned into an 8+ year journey instead. You have encouraged me to take leaps of faith that have helped shape my life in the best ways. Another huge thank you to my amazing husband, you have been the unexpected highlight of this journey. I'm so grateful for your ability to provide laughter and perspective to my downs, and cheer me on to reach for the highs. Lastly, thank you to my family and friends in The Netherlands for the long calls and visits home that helped me reset and re-energize.

Next, I would like to express my gratitude to all the friends and colleagues at Caltech. It's been such a pleasure to learn with and from you. Thank you to the SDN-group specifically. I had not expected to find a group of peers so willing to raise each other up. I'm grateful for your friendships and the many dinners we shared together.

Lastly, thank you to my advisors for the opportunity to work on science that still excites me today. And, of course, a thank you to my committee members for the feedback, enthusiasm and perspective as this thesis took shape over the years.

ABSTRACT

Adaptive decision-making requires the brain to flexibly compute subjective values for choice options across changing environmental contexts and distinct operational domains, such as visual stimuli versus motor actions. This dissertation elucidates the neural architectures supporting adaptive valuation across multiple scales of analysis, leveraging functional magnetic resonance imaging (fMRI) and human single-unit electrophysiology.

First, investigating the valuation of multi-attribute stimuli under shifting goals reveals a hierarchical valuation process. While visual cortices represent static, context-independent stimulus attributes, regions within the prefrontal cortex (i.e., ventromedial prefrontal cortex (vmPFC) and orbitofrontal cortex (OFC)) transform these into context-sensitive “attributes in value space”. These intermediate representations are then integrated into a unified subjective value signal along the dorsomedial prefrontal cortex (dmPFC).

Second, single-neuron recordings during structurally parallel action- and stimulus-based tasks uncover a temporally shifting representational architecture. Prior to choice, neurons across the vmPFC, anterior cingulate cortex (ACC), and pre-supplementary motor area (preSMA) broadly track available pre-decision values as to facilitate comparison across both stimuli and actions. However when probing choice-dependent valuation, the network segregates: the ACC specializes in action-based chosen value and the vmPFC tracks stimulus-based chosen value. In contrast to the specialization the preSMA encodes value and chosen identity across both choice domains.

Together, these findings demonstrate that human valuation relies on a dynamic sequence of transformations. This prefrontal network balances the abstraction required to compare disparate options with the specificity needed for accurate goal-directed behavior.

PUBLISHED CONTENT AND CONTRIBUTIONS

Fransen, A. et al. (2026a). “Choice-driven remapping of action- and stimulus-anchored value in human single neurons”.

Fransen, A. et al. (2026b). “The neural computations underlying context dependent attribute-based valuation of complex stimuli”. In: doi: <https://doi.org/10.1101/2025.10.10.681681>.

Contributed to the task design across both projects, implemented the tasks, collected the data, ran all behavioral and neural analysis, and wrote the papers based on feedback from co-authors.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
Published Content and Contributions	v
Table of Contents	v
List of Illustrations	vii
List of Tables	viii
Chapter I: General introduction	1
1.1 The challenge of adaptive valuation	3
1.2 Valuing multi-attribute objects	4
1.3 Disentangling stimulus- and action-based choice mechanisms	5
1.4 From whole brain to single neurons	7
Chapter II: The neural computations underlying context dependent attribute- based valuation of complex stimuli	12
Chapter III: Choice-driven remapping of action- and stimulus-anchored value in human single neurons	45
Chapter IV: General discussion	81
4.1 Summary of main findings	81
4.2 ‘General’ value computation	82
4.3 ‘Specific’ value computation	85
4.4 Concluding remarks	87
Bibliography	91
Appendix A: Supplementary material - Chapter 2	92
Appendix B: Supplementary material - Chapter 3	99

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
2.1 Task and hypothesis schematic.	17
2.2 Behavioral results.	20
2.3 Static attribute results.	23
2.4 Generalization analysis for attributes in value space.	26
2.5 Integrated value results.	29
2.6 Integrated value vs attributes in value space encoding.	32
3.1 Task and behavioral results.	49
3.2 Decision value coding.	55
3.3 Chosen value encoding.	58
3.4 Chosen identity coding.	62
A.1 Supplement: Online experiment design.	92
A.2 Supplement: Rating distributions.	93
A.3 Supplement: Behavioral model fitting results.	95
A.4 Supplement: ROI locations as used in generalization analysis.	95
A.5 Supplement: Dorsal visual stream attribute encoding.	96
A.6 Supplement: Generalization results additional attributes.	97
A.7 Supplement: Integrated value decoding.	98
B.1 Supplement: electrode locations.	100
B.2 Supplement: model performance.	100
B.3 Supplement: Consistency in coding direction of the largest beta.	101
B.4 Supplement: chosen identity during outcome.	102

LIST OF TABLES

<i>Number</i>	<i>Page</i>
A.1 Attribute main effects across subjects within each goal-context. . . .	94
A.2 Statistics of Kruskal - Wallis tests for coefficients fit using the Logistic model.	94
B.1 Patient demographics.	99
B.2 Total BIC scores across subjects.	99

Chapter 1

GENERAL INTRODUCTION

Every day, humans make choices that dictate how they allocate their limited time, energy, and resources. Unlike perceptual decisions, which rely on identifying objective physical attributes of the external environment (e.g., determining which of two visual stimuli is brighter), value-based decisions are driven by subjective, internally generated preferences (Cabanac, 1992). To make adaptive choices, the brain must continuously estimate the expected desirability (i.e., expected value) of potential candidates (Rangel, Camerer, and Montague, 2008). Attaching such value judgments to choice options allows the decision-maker to evaluate and compare a variety of different options on a subjective scale (McNamee, Rangel, and O’Doherty, 2013). Through the computation and comparison of these subjective (expected) values, choices can be made to maximize rewards and minimize punishments. Ultimately, by translating internal preferences into actionable values, this system forms the foundation of goal-directed behavior.

A rich body of literature has sought to map the neural substrates responsible for generating these value signals that underlie choice. Foundational work using single-unit electrophysiology in non-human primates established that individual neurons in the orbitofrontal cortex (i.e., OFC) encode the subjective value of available options parametrically; crucially, this coding occurs in an action-independent manner, abstracting away from both the sensory properties of the goods and the physical movements needed to execute the choice (Padoa-Schioppa and Assad, 2006). In humans, quantitative meta-analyses of functional neuroimaging studies have confirmed the presence of a valuation system centered on the ventromedial prefrontal cortex (i.e., vmPFC) and ventral striatum, which reliably represents value across vastly different categories of rewards (Bartra, McGuire, and Kable, 2013). Targeted fMRI paradigms have demonstrated that especially the prefrontal value-coding areas (e.g., vmPFC, OFC) track the expected utility of primary reinforcers like flavors and liquid delivery (O’Doherty et al., 2001), as well as secondary reinforcers like monetary gambles (Tom et al., 2007). This valuation architecture appears to underpin valuation of a wide variety of objects that are not primary or monetary in nature (McNamee, Rangel, and O’Doherty, 2013). Remarkably, these regions have also been found to process not only tangible rewards but also abstract, high-level

stimuli such as visual aesthetic beauty and compliance with complex social norms (O'Doherty et al., 2003; Ruff and Fehr, 2014).

For these value representations to successfully guide goal-directed behavior, however, the decision-maker must select and execute specific physical actions. A central debate in decision neuroscience centers on whether this choice computation is fundamentally framed in the space of actions or the space of stimuli and goods. Under an action-based choice framework, value judgments are converted to the space of actions and options are evaluated directly as alternative behavioral responses (Cisek, 2007; Kennerley et al., 2006). Conversely, a goods-based choice framework posits that options are first compared on an abstract, effector-independent value scale within ventral prefrontal networks, entirely isolated from motor preparation (Padoa-Schioppa and Assad, 2006). If the choice is not made over actions directly but is instead resolved in the space of goods, a downstream goods-to-action transformation becomes a structural necessity. Since a physical action is ultimately required for even a goods-based choice the abstractly selected option must be dynamically mapped onto effector-specific motor networks to execute the physical response (Wunderlich, Rangel, and O'Doherty, 2009; Cai and Padoa-Schioppa, 2014). These types of hypothesized mechanisms are necessary for bridging the gap between economic preference and behavioral output.

To guide behavior adaptively over time, the valuation system must also solve the fundamental problem of credit assignment (i.e., linking downstream outcomes to the specific choices or actions that caused them). This dynamic updating is largely orchestrated by subcortical structures working in tight loops with the cortex. Midbrain dopaminergic neurons broadcast reward prediction errors (RPEs), which signal the discrepancy between expected and actual outcomes (Schultz, Dayan, and Montague, 1997). Direct signals of these types of value updates have been additionally found in lateral PFC, basal-ganglia and the midbrain (e.g., Valentin and O'Doherty, 2009; Asaad and Eskandar, 2011; Lak, Stauffer, and Schultz, 2014; Man et al., 2024). These prediction errors can systematically revise the value estimates maintained in prefrontal regions based on received outcomes. This process is of particular importance in changing environments, where RPE-based updating needs to ensure value estimates are a reflection of the current state of the world.

This collective body of literature provides robust evidence for an integrated system able to compute value estimates, select choice options, and learn based on RPEs. However, classical experimental paradigms often capture these computations under

relatively stable conditions. Instead, the problem facing the brain’s value-based decision network through interactions with our everyday environment is inherently unstable. A person cannot simply rely on a static lookup table that maps specific choice options (e.g., objects or actions) to fixed value judgments. Both our external environment and our internal physiological states are highly dynamic, meaning that a choice that is highly advantageous in one moment may be detrimental in the next. Often the timescale of these environmental changes could make a complete re-learning of values based on RPEs inefficient. Consequently, the computation of value must be a profoundly flexible process, capable of rapidly shifting in accordance to the continuous stream of information to guide behavior adaptively.

1.1 The challenge of adaptive valuation

This demand for behavioral adaptability imposes significant computational challenges on the brain. Specifically, adaptive decision-making requires the valuation- and choice-systems described above to rapidly adjust along several distinct dimensions. First, we must frequently re-evaluate identical options as our immediate goals or environmental contexts shift (Castegnetti, Zurita, and DeMartino, 2021). This means we cannot rely on static associations between objects and their values (as in e.g., McNamee, Rangel, and O’Doherty, 2013; Suzuki, Cross, and O’Doherty, 2017; Iigaya et al., 2021). Instead, these shifting conditions require a mechanism that can integrate the composite, stimulus-intrinsic attributes of objects into a cohesive, context-sensitive value signal (O’Doherty, Rutishauser, and Iigaya, 2021). Second, the brain must be able to dynamically compute and update value estimates across fundamentally different operational domains. To effectively guide goal-directed behavior, the value computation system must be equipped to estimate value-based preferences for both the environmental stimuli encountered and the specific set of actions available to the decision-maker (Hare et al., 2011).

Understanding how the brain successfully navigates these flexibility demands is the central focus of this dissertation. This work elucidates the neural architectures underlying flexible valuation by investigating how the brain dynamically constructs and updates value across these dimensions of complexity and choice domain. We argue that completely characterizing this system requires a multi-scale approach: while macro-scale network dynamics allow us to track the high-level information routing towards/across the frontal lobe, the actual representation of value likely occurs on the scale of single-neuron computations. By leveraging both macroscopic functional imaging network dynamics and human single-neuron computations, this

research elucidates computations facilitating flexible valuation across various levels of analysis. Ultimately, it reveals how the frontal lobe and associated circuits construct, differentiate, and integrate value to guide adaptive behavior in changing environments.

1.2 Valuing multi-attribute objects

Real-world choices rarely involve simple, uni-dimensional options. Instead, the value of everyday objects relies on the integration of their composite, stimulus-intrinsic attributes (as hypothesized in Winterfeldt and Fisher, 1975). Recent neuroimaging work has successfully mapped this additive process, demonstrating that the brain tracks individual attributes. For example, a food’s nutritional properties or an object’s visual properties are found to be encoded in distinct cortical areas before combining them into a singular, integrated (i.e., summed) value signal within the ventromedial prefrontal cortex (vmPFC) and medial orbitofrontal cortex (mOFC) (as in Kahnt et al., 2011; Suzuki, Cross, and O’Doherty, 2017; Iigaya et al., 2023). While these models of multi-attribute value computation suggest a straightforward integration of attributes to predict overall value, such architectures struggle to account for the flexibility of human preferences. Specifically, valuation models must reconcile the fact that while an item’s physical properties are constant, its subjective value is highly dependent on the current environmental state (e.g., the warmth of a jacket is invaluable on a ski trip but a burden on a beach).

As a stable weighted sum over attributes cannot accommodate changing overall values, recent theoretical work proposed that context-dependent value construction must rely on one of two potential computational architectures. The first hypothesis posits a highly flexible weight-adjustment scheme, where underlying stimulus attributes are represented independently of context, but their integration weights are dynamically scaled on the fly as external demands or internal states shift (O’Doherty, Rutishauser, and Iigaya, 2021). The second hypothesis avoids the high computational expense of real-time weight recalculation by proposing fixed-weight multiplicative representations (O’Doherty, Rutishauser, and Iigaya, 2021). Under this framework, the brain builds non-linear mixtures of ‘attribute $\times$ context’ combinations, meaning flexibility is achieved not by changing weights, but by selectively gating and activating the specific conjunctive representations relevant to the current environment. Determining which of these mechanisms is employed in the prefrontal cortex remains a critical open question in establishing a truly adaptive model of multi-attribute valuation.

The first part of this dissertation addresses this open question by investigating the neural and computational mechanisms that such flexible multi-attribute valuation. Rather than assuming the valuation system relies on a fixed integration over stimulus attributes, this work tests the hypothesis that the brain realizes an intermediary ‘attribute $\times$ context’ representation. Specifically, chapter 2 investigates the hypothesis that the brain transforms static stimulus attributes into context-sensitive ‘attributes in value space’ before integrating them into a flexible overall stimulus value. By characterizing these ‘attributes in value space’ as an intermediate layer of conjunctive attribute-context coding, this framework could provide a mechanism for how the brain computes dynamically shifting value judgments from invariant stimulus attributes.

1.3 Disentangling stimulus- and action-based choice mechanisms

Beyond adjusting the value of objects in accordance with current goals or contexts, flexible valuation also requires the brain to dynamically adapt to *what* is being valued. A fundamental debate in the study of decision-making centers on whether the neural computations underlying choice are segregated into distinct domains based on the nature of the option (e.g., Rudebeck et al., 2008; Camille, Tsuchida, and Fellows, 2011; Cai and Padoa-Schioppa, 2014; Gläscher, Hampton, and O’Doherty, 2009). Specifically, this literature aims to determine whether the brain maintains separate architectures for computing “stimulus value” (the value of a visual object) versus “action value” (the value of a physical motor plan). The relevance of this debate extends beyond neuroscience as reinforcement learning (RL) frameworks explicitly divide environmental evaluation into distinct computational variables (Sutton and Barto, 1998). In these models, stimulus value conceptually aligns with state-value computations ($V(s)$), representing the expected future reward of encountering a specific stimulus independent of the actions taken. Conversely, action value corresponds to the Q -value ($Q(s, a)$), representing the expected utility of executing a specific physical behavior. This algorithmic distinction directly mirrors the choice frameworks detailed above: a goods-based choice architecture relies on updating stimulus-centric states, whereas an action-based architecture necessitates tracking the specific value attached to motor executions.

While this computational framework provides a precise mathematical language for distinct valuation domains, how the human brain implements these algorithms remains a subject of research. Single-neuron studies in non-human primates have mapped these computational variables to specific neural circuits, suggesting a clear

anatomical segregation between valuation domains. Under this classical view, action values are predominantly represented within dorsal corticostriatal circuits, including the dorsal striatum and the anterior cingulate cortex (ACC) (e.g., Samejima et al., 2005; Hayden and Platt, 2010). In contrast, stimulus values are often mapped to ventral prefrontal regions such as the orbitofrontal cortex (OFC) and ventromedial prefrontal cortex (vmPFC) (Padoa-Schioppa and Assad, 2006). This regional specialization is strongly supported by causal lesion evidence in both humans and non-human primates. These studies reveal a striking double dissociation between these stimulus- and action-based networks. Specifically, focal damage to the OFC or vmPFC selectively disrupts choice behavior when decisions must be guided by competing stimulus values, while leaving action-guided choices fully preserved (Rudebeck et al., 2008; Camille, Tsuchida, and Fellows, 2011). Conversely, focal lesions to the ACC yield the precise opposite deficit, impairing the ability to choose optimally in value-based decisions framed around physical actions while leaving stimulus-guided choice intact (Rudebeck et al., 2008; Camille, Tsuchida, and Fellows, 2011).

These distinct anatomical boundaries suggested by causal models, however, become more blurred when evaluating macro-scale networks in the healthy human brain. Functional magnetic resonance imaging (fMRI) paradigms investigating the spatial organization of these valuation domains frequently reveal significant regional convergence rather than isolated systems. Specifically, the ventromedial prefrontal cortex (vmPFC) and amygdala have been implicated in tracking value-related metrics that support choices across both the stimulus and action domains (Gläscher, Hampton, and O'Doherty, 2009; Wunderlich, Rangel, and O'Doherty, 2010). Furthermore, efforts to map the functional cascade of stimulus-based choices onto behavioral execution demonstrate that this transition relies on a distributed network, including the vmPFC, dorsomedial PFC, and motor cortices, to coordinate stimulus-to-action mapping (Hare et al., 2011).

However, because the blood-oxygen-level-dependent (BOLD) signal aggregates metabolic activity across hundreds of thousands of neurons within any given voxel, macro-scale imaging cannot differentiate between true domain-general multiplexing at the cellular level and the dense, sub-voxel intermixing of discrete, domain-specific neuronal populations. Due to the rare and invasive nature of human intracranial recordings, it remains fundamentally unknown whether individual neurons in the human brain strictly segregate by valuation domain to maintain separate representa-

tional spaces, or if they instead rely on a shared, domain-general cellular substrate.

Furthermore, attempting to map these value computations to their neuronal underpinnings is complicated by the methodological difficulty of separating motor actions from spatial targets. Because some experimental paradigms require subjects to indicate choices by moving toward a specific location, neural signals traditionally interpreted as encoding action value might simultaneously reflect the value of the spatial target, spatial attention, or the directional movement vector itself (for further discussion see Padoa-Schioppa, 2011). This has made it difficult to isolate “pure” motor action value. In consideration of these potential confounds, here I define action value as value tied specifically to the physical effector executing the choice, completely independent of spatial location.

To determine whether the human brain employs a unified or domain-specific valuation architecture, the second part of this dissertation leverages a task specifically designed to separate stimulus- and action-based choice mechanisms. By employing human single-unit electrophysiology, I directly contrast single-neuron activity across structurally parallel, action-based and stimulus-based reinforcement learning tasks. Crucially, this chapter goes beyond initial value representation to probe how the broader decision architecture adapts to these different choice domains. Specifically, by tracking cellular dynamics across the full choice process, Chapter 3 additionally investigates the representation of the motor execution plan and evaluates how these signals may serve as anchors to facilitate outcome-driven credit assignment. Ultimately, this work aims to reveal how the human brain maps valuation, choice execution, and credit assignment across action- and stimulus-based choice domains.

1.4 From whole brain to single neurons

To elucidate the flexibility of the human valuation system, this thesis explores these processes across multiple spatial and computational scales. Functional magnetic resonance imaging (fMRI), utilized in Chapter 2, provides crucial insights into the whole-brain network dynamics of value computation. It allows us to map how the various computations underlying multi-attribute valuation are distributed across the brain, progressing from the visual cortex to the prefrontal cortex. As such, the use of fMRI allows us to gain a holistic perspective on flexible value computation.

However, functional imaging inherently lacks the spatial resolution required to determine how these signals are constructed at the cellular level—a spatial limitation that is critical for our second question. If a single functional voxel correlates with

both stimulus and action value, fMRI cannot determine whether that region contains a genuinely mixed-selective neural population, or simply spatially intermixed, domain-specific neurons. Due to this limitation, we turn to human single-unit electrophysiology to resolve whether the computations of stimulus and action value are integrated or segregated at the fundamental level of the neural circuit.

By combining macroscopic fMRI (Chapter 2) with high-resolution microelectrode recordings (Chapter 3), this dissertation spans these critical levels of analysis. Together, these complementary methodologies provide a comprehensive account of flexible valuation, revealing how the human brain operates both as a distributed macro-network and at the resolution of individual neurons to construct, differentiate, and integrate value to guide dynamic behavior.

References

- Asaad, W. F. and E. N. Eskandar (Dec. 2011). “Encoding of Both Positive and Negative Reward Prediction Errors by Neurons of the Primate Lateral Prefrontal Cortex and Caudate Nucleus”. en. In: *The Journal of Neuroscience* 31.49, pp. 17772–17787. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3793-11.2011.
- Bartra, O., J. T. McGuire, and J. W. Kable (Aug. 2013). “The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value”. en. In: *NeuroImage* 76, pp. 412–427. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2013.02.063.
- Cabanac, M. (Mar. 1992). “Pleasure: the common currency”. en. In: *Journal of Theoretical Biology* 155.2, pp. 173–200. ISSN: 00225193. DOI: 10.1016/S0022-5193(05)80594-6.
- Cai, X. and C. Padoa-Schioppa (Mar. 2014). “Contributions of Orbitofrontal and Lateral Prefrontal Cortices to Economic Choice and the Good-to-Action Transformation”. en. In: *Neuron* 81.5, pp. 1140–1151. ISSN: 08966273. DOI: 10.1016/j.neuron.2014.01.008.
- Camille, N., A. Tsuchida, and L. K. Fellows (Oct. 2011). “Double Dissociation of Stimulus-Value and Action-Value Learning in Humans with Orbitofrontal or Anterior Cingulate Cortex Damage”. en. In: *Journal of Neuroscience* 31.42, pp. 15048–15052. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3164-11.2011.
- Castegnetti, G., M. Zurita, and B. DeMartino (2021). “How usefulness shapes neural representations during goal-directed behavior”. en. In: *Science Advances* 7.15. DOI: 10.1126/sciadv.abd5363.

- Cisek, P. (Sept. 2007). “Cortical mechanisms of action selection: the affordance competition hypothesis”. en. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 362.1485, pp. 1585–1599. ISSN: 0962-8436, 1471-2970. DOI: 10.1098/rstb.2007.2054.
- Gläscher, J., A. N. Hampton, and J. P. O’Doherty (Feb. 2009). “Determining a Role for Ventromedial Prefrontal Cortex in Encoding Action-Based Value Signals During Reward-Related Decision Making”. en. In: *Cerebral Cortex* 19.2, pp. 483–495. ISSN: 1460-2199, 1047-3211. DOI: 10.1093/cercor/bhn098.
- Hare, T. A. et al. (Nov. 2011). “Transformation of stimulus value signals into motor commands during simple choice”. en. In: *Proceedings of the National Academy of Sciences* 108.44, pp. 18120–18125. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1109322108.
- Hayden, B. Y. and M. L. Platt (Mar. 2010). “Neurons in Anterior Cingulate Cortex Multiplex Information about Reward and Action”. en. In: *The Journal of Neuroscience* 30.9, pp. 3339–3346. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.4874-09.2010.
- Iigaya, K. et al. (June 2021). “Aesthetic preference for art can be predicted from a mixture of low- and high-level visual features”. en. In: *Nature Human Behaviour* 5.6, pp. 743–755. ISSN: 2397-3374. DOI: 10.1038/s41562-021-01124-6.
- Iigaya, K. et al. (Jan. 2023). “Neural mechanisms underlying the hierarchical construction of perceived aesthetic value”. en. In: *Nature Communications* 14.1, p. 127. ISSN: 2041-1723. DOI: 10.1038/s41467-022-35654-y.
- Kahnt, T. et al. (May 2011). “Decoding different roles for vmPFC and dlPFC in multi-attribute decision making”. en. In: *NeuroImage* 56.2, pp. 709–715. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2010.05.058.
- Kennerley, S. W. et al. (July 2006). “Optimal decision making and the anterior cingulate cortex”. en. In: *Nature Neuroscience* 9.7, pp. 940–947. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn1724.
- Lak, A., W. R. Stauffer, and W. Schultz (2014). “Dopamine prediction error responses integrate subjective value from different reward dimensions”. en. In: *Proceedings of the National Academy of Sciences* 111.6, pp. 2343–2348. DOI: 10.1073/pnas.1321596111.
- Man, V. et al. (Mar. 2024). “Temporally organized representations of reward and risk in the human brain”. en. In: *Nature Communications* 15.1, p. 2162. ISSN: 2041-1723. DOI: 10.1038/s41467-024-46094-1.
- McNamee, D., A. Rangel, and J. P. O’Doherty (2013). “Category-dependent and category-independent goal-value codes in human ventromedial prefrontal cortex”. en. In: *Nature Neuroscience* 16.4, pp. 479–485. DOI: 10.1038/nn.3337.

- O'Doherty, J. et al. (Mar. 2001). "Representation of Pleasant and Aversive Taste in the Human Brain". en. In: *Journal of Neurophysiology* 85.3, pp. 1315–1321. ISSN: 0022-3077, 1522-1598. DOI: 10.1152/jn.2001.85.3.1315.
- O'Doherty, J. P., U. Rutishauser, and K. Iigaya (Oct. 2021). "The hierarchical construction of value". en. In: *Current Opinion in Behavioral Sciences* 41, pp. 71–77. ISSN: 23521546. DOI: 10.1016/j.cobeha.2021.03.027.
- O'Doherty, J. et al. (Jan. 2003). "Beauty in a smile: the role of medial orbitofrontal cortex in facial attractiveness". en. In: *Neuropsychologia* 41.2, pp. 147–155. ISSN: 00283932. DOI: 10.1016/S0028-3932(02)00145-8.
- Padoa-Schioppa, C. (July 2011). "Neurobiology of Economic Choice: A Good-Based Model". en. In: *Annual Review of Neuroscience* 34.1, pp. 333–359. ISSN: 0147-006X, 1545-4126. DOI: 10.1146/annurev-neuro-061010-113648.
- Padoa-Schioppa, C. and J. A. Assad (May 2006). "Neurons in the orbitofrontal cortex encode economic value". en. In: *Nature* 441.7090, pp. 223–226. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature04676.
- Rangel, A., C. Camerer, and P. R. Montague (July 2008). "A framework for studying the neurobiology of value-based decision making". en. In: *Nature Reviews Neuroscience* 9.7, pp. 545–556. ISSN: 1471-003X, 1471-0048. DOI: 10.1038/nrn2357.
- Rudebeck, P. H. et al. (Dec. 2008). "Frontal Cortex Subregions Play Distinct Roles in Choices between Actions and Stimuli". en. In: *The Journal of Neuroscience* 28.51, pp. 13775–13785. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3541-08.2008.
- Ruff, C. C. and E. Fehr (Aug. 2014). "The neurobiology of rewards and values in social decision making". en. In: *Nature Reviews Neuroscience* 15.8, pp. 549–562. ISSN: 1471-003X, 1471-0048. DOI: 10.1038/nrn3776.
- Samejima, K. et al. (Nov. 2005). "Representation of Action-Specific Reward Values in the Striatum". en. In: *Science* 310.5752, pp. 1337–1340. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1115270.
- Schultz, W., P. Dayan, and P. R. Montague (1997). "A Neural Substrate of Prediction and Reward". en. In: *Science* 275.5306, pp. 1593–1599. DOI: DOI: 10.1126/science.275.5306.1593.
- Sutton, R. S. and A. G. Barto (1998). *Reinforcement Learning: An Introduction*. en. The MIT Press.
- Suzuki, S., L. Cross, and J. P. O'Doherty (2017). "Elucidating the underlying components of food valuation in the human orbitofrontal cortex". en. In: *Nature Neuroscience* 20.12, pp. 1780–1786. DOI: 10.1038/s41593-017-0008-x.
- Tom, S. M. et al. (Jan. 2007). "The Neural Basis of Loss Aversion in Decision-Making Under Risk". en. In: *Science* 315.5811, pp. 515–518. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1134239.

- Valentin, V. V. and J. P. O'Doherty (Dec. 2009). "Overlapping Prediction Errors in Dorsal Striatum During Instrumental Learning With Juice and Money Reward in the Human Brain". en. In: *Journal of Neurophysiology* 102.6, pp. 3384–3391. ISSN: 0022-3077, 1522-1598. DOI: 10.1152/jn.91195.2008.
- Winterfeldt, D. von and G. Fisher (1975). "Multi-attribute utility theory: models and assessment procedures". In: *Wendt, D., Vlek, C.A.J. (Eds.), Utility, probability, and human decision making*. Dordrecht, The Netherlands: Reidel, pp. 47–86.
- Wunderlich, K., A. Rangel, and J. P. O'Doherty (Oct. 2009). "Neural computations underlying action-based decision making in the human brain". en. In: *Proceedings of the National Academy of Sciences* 106.40, pp. 17199–17204. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.0901077106.
- (Aug. 2010). "Economic choices can be made using only stimulus values". en. In: *Proceedings of the National Academy of Sciences* 107.34, pp. 15005–15010. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1002258107.

Chapter 2

THE NEURAL COMPUTATIONS UNDERLYING CONTEXT DEPENDENT ATTRIBUTE-BASED VALUATION OF COMPLEX STIMULI

Fransen, A. et al. (2026). “The neural computations underlying context dependent attribute-based valuation of complex stimuli”. In: doi: <https://doi.org/10.1101/2025.10.10.681681>.

Abstract

Adaptive decision making requires value computations to be flexible as a stimulus might be highly desirable in one context but completely unappealing in another. However, it is not known how the brain can support such flexible modulation of overall stimulus value especially considering the constituent attributes of a stimulus are static. Here we test for the presence of a novel step in the hierarchical framework of flexible multi-attribute value computation to bridge this dichotomy. We hypothesize that individual attributes of a stimulus are converted into contextually dependent ‘attributes in value space’ before being combined into an integrated stimulus value. To test this framework, human participants (online n=95, MRI n=35) provided ratings for 75 unique high-dimensional clothing stimuli under three instructed ‘goal-contexts’ designed to elicit differences in value judgments. Integrated value ratings for each stimulus were found to indeed be goal-context dependent, while individual stimulus-attributes varied markedly in their contribution to value ratings across goal contexts. In the fMRI data, static representations of constituent stimulus attributes were revealed in visual areas. In contrast, encoding of the hypothesized ‘attributes in value space’ was found alongside integrated overall stimulus value representations within distinct regions of prefrontal cortex. More specifically, attributes in value space and integrated stimulus value were found in vmPFC and dmPFC respectively. These findings indicate that the construction of value for high-dimensional stimuli is achieved through the computation of goal-context-dependent attributes in value space, providing mechanistic insight into how the brain can transform static stimulus attributes into flexible stimulus values as contexts change.

Introduction

Models of decision-making propose that the value of multi-attribute objects relies on the integration of their composite stimulus-intrinsic attributes (Kahnt et al. (2011) and Winterfeldt and Fisher (1975)). Most simply, the value of a gamble, for example, is computed by integrating the probability and magnitude of the prize on offer (Platt and Glimcher, 1999). Similar principles have been proposed to facilitate value judgments for objects with high-dimensional attributes (Zysset et al., 2006), where the presence and degree of various stimulus-intrinsic attributes are evaluated and subsequently summed in a weighted fashion to arrive at an overall object value (Winterfeldt and Fisher, 1975). The resulting sum is then suggested to facilitate decision-making through between-stimulus comparisons. Previous research has examined the relationship between stimulus-intrinsic raw attributes and integrated value, demonstrating that a weighted combination of constituent attributes of a stimulus can predict how much value is assigned to a stimulus overall. For instance, the nutritive content of a food item can influence a person's willingness to consume it (Suzuki, Cross, and O'Doherty, 2017), while the visual properties of a piece of art contribute to aesthetic judgments (Iigaya et al., 2021). This influence of composite attributes on value has been observed across various types of stimuli, ranging from t-shirts with aesthetic and semantic prints (Lim, O'Doherty, and Rangel, 2013) to gambles (Farashahi et al., 2019).

However, focusing on only raw attributes and a fixed weighted integration of attributes to compute an integrated value has left an important gap in our understanding of the computations at play. This gap becomes clear when we consider that value judgments are not made in a vacuum but are highly flexible in that they can be heavily influenced by a person's current goals (Castegnetti, Zurita, and DeMartino, 2021), internal homeostatic state, or environmental context (Farashahi et al., 2019). The same (visually identical) item can be valued highly under some circumstances yet considered low in value in other situations. For example, a ski jacket might be deemed highly valuable to an individual when going on a skiing holiday, but the same jacket would be of very low value when going on a beach vacation. A fundamental open question then is how can such highly flexible contextually-specific value computations be accommodated and understood within the framework of an attribute-based integration mechanism for computing overall stimulus value.

The static computation of multi-attribute stimulus value has been studied extensively. In particular, distinct regions of prefrontal cortex have been implicated in different

aspects of this process. Regions of the lateral prefrontal cortex, comprising both dorsolateral prefrontal cortex and lateral orbitofrontal cortex, have been found to encode information about individual stimulus attributes (e.g., Kahnt et al., 2011, Suzuki, Cross, and O’Doherty, 2017, Iigaya et al., 2023). In contrast, the medial frontal cortex, particularly the ventromedial aspect, has been implicated in encoding the combined value of multi-attribute objects (e.g., Lim, O’Doherty, and Rangel, 2013, McNamee, Rangel, and O’Doherty, 2013, Iigaya et al., 2023). However, in all of these studies it is not known how these brain regions can facilitate context-specific (i.e. flexible) value computations through attribute integration. More specifically, these studies have not distinguished representation of raw stimulus attributes (i.e. the amount of attribute present in a stimulus irrespective of context), from the representations of the value of individual attributes. This gap means we do not yet know how multi-attribute value computation can be flexibly adjusted to changes in our goals, while this is a process that is abundantly present in our everyday decision making.

Here, we aim to address this gap by contrasting two potential computational architectures (O’Doherty, Rutishauser, and Iigaya, 2021). Our primary hypothesis posits that there exists a critical intermediate computational process between the representation of the attributes that make up a stimulus, and the overall integrated value. In this intermediate step, the value of individual attributes is computed in a context-sensitive manner. This context-sensitivity implies that the attributes of, say, a ski-jacket might be valued differently in different contexts. Specifically, we propose that these context-sensitive attributes (e.g., its warmth) are evaluated in a context-dependent value space before being integrated into an overall value signal. We will refer to this hypothesized intermediary step as computing “attributes in value space”. Alternatively, the brain might utilize a direct integration mechanism. Under this alternative hypothesis, raw, static attributes are directly integrated into a goal-dependent overall value. Here, the goal-context simply alters the integration weights applied to the raw stimulus features, effectively bypassing the need to explicitly represent context-dependent attribute values. The goal of the present study is to investigate whether the brain represents “attributes in value space”, as the presence of this computation would help adjudicate between these two hypothesized computational mechanisms.

We designed a novel experimental paradigm specifically to address these hypotheses by distinguishing the neuronal representation of each of three distinct components:

1. the degree to which attributes are present in a stimulus (stimulus-intrinsic raw attributes), 2. the hypothesized value these individual attributes confer given the current goal (attributes in value space), and 3. the integrated value of the stimulus. By isolating these components, we aim to adjudicate between these competing architectures, investigating whether the brain relies on a direct goal-weighted integration of raw attributes or utilizes an intermediate representation of attributes in value space to achieve the goal-dependent valuation of multi-attribute stimuli.

Results

Here we set out to investigate the neural mechanisms underlying the computational steps that enable us to form goal-dependent values for multi-attribute stimuli. To achieve this, we designed a novel task; the “shopping task”. In this task, we introduced participants to instructed “goal-contexts”, enabling us to induce and measure changes in the flexible computation of goal-dependent values across goal-contexts. Under each goal-context, participants viewed 75 unique gender matched clothing stimuli. These goals included going on a ski trip, visiting a tropical island and going for a job interview. Participants indicated how much they would like to wear the displayed item under the current goal (see Figure 2.1). Ratings ranged from 1-“do not want to wear at all” to 7-“want to wear a lot” (see Methods for more details). We expected that the subjective integrated value rating of a particular item (operationalized here by how much a participant would like to wear the item in a given goal-context), would be modulated by the goal-context (i.e., to be context-dependent). To reprise a previously introduced example, a heavy winter jacket might be highly valued in the ski trip goal-context (henceforth ‘context’), while being deemed of very low value in the tropical island vacation context. Similarly, a swimsuit might be highly valued in the tropical island vacation context, but of low value in the job interview context.

We leverage this new paradigm in two experiments: one collected in an online sample (Prolific; $n=95$, 57 female) and one conducted in an fMRI scanner ($n=35$, 20 female, age 18-64). As part of both of these experiments, we measured subjective stimulus ratings directly through the described rating task (see Figure 2.1).

Using this task, we firstly aimed to capture the neural representation of the raw attributes that make up a stimulus. These raw attributes quantify the perceived or objective content of a given attribute in the stimulus as a whole (independently of that attribute’s value). For instance, a heavy jacket will be ‘high’ in the attribute

of warmth, regardless of whether it is worn on a ski trip or on a tropical island vacation (see Figure 2.1C). In addition to the neural computations underlying raw attributes, we set out to find neural representations of attributes transformed into value-space. These hypothesized ‘attributes in value space’ can be distinguished from raw attributes, since the warmth attribute will be valued differently depending on the context (e.g., whether an individual is wearing the jacket on a ski-trip or on a tropical island vacation - see Figure 2.1D). Finally, our task allows us to also probe the subjective (integrated) value of such multi-attribute stimuli (see Figure 2.1E). The integrated value is the value of a stimulus (e.g., the heavy jacket) within each context, when taking into account multiple attributes simultaneously.

Given our main objective is testing for evidence of the encoding of attributes in value space, we picked a small set of salient high-level attributes pertinent to value decisions about clothing items: the “warmth”, “formalness”, “comfort” and “festiveness” of the items. While these are a limited subset of a much larger set of possible attributes, they are sufficient to allow us to distinguish raw attribute, attributes in value space and integrated value representations. We also examined “low-level” visual attributes of the clothing items which correspond to elementary visual properties of the stimuli (see Methods for more details and Igaya et al., 2020). We confirm our claims of sufficiency through this set of “low-level” attributes and contrast the neural representation of these lower level attributes with the selected higher level attributes.




Figure 2.1: Task and hypothesis schematic. See continued caption below.

Behavioral contribution of attributes to integrated value ratings.

We started by probing the relationship between attributes and value at a behavioral level. As attributes in value space (see Figure 2.1D) are latent variables, we sought to approximate their effect on integrated value through a weighting of raw attributes. This weighting can be thought of as an estimation of the transformation the raw attributes undergo to hold value-based information (see Figure 3.1C to D).

To verify that our a priori attributes consistently contribute to overall stimulus value across individuals, we evaluated the independent main effect of each attribute within every context. Using a mixed-effects model to control for individual differences in ratings (fMRI sample only; see Methods), we found that all attributes were significant independent predictors of integrated value under every instructed context (see Supplemental Table A.1 for T-test statistics).

To confirm that the goal contexts differentially transform raw attributes into integrated values (stimulus ratings), we extracted attribute weights using context- and participant-specific regression models (see Methods; Supplemental Figure A.2G). Comparing these coefficients across contexts revealed consistent goal-dependent differences for formality (Kruskal-Wallis; fMRI: $H=70$, $p = 6.23e^{-16}$, Online: $H=189$, $p = 1.12e^{-41}$) and warmth (Kruskal-Wallis; fMRI: $H=92$, $p = 1.01e^{-20}$,

Figure 2.1: (continued) **A)** Participants performed a rating task under three instructed goal-contexts: tropical island vacation, ski trip and job interview. Participants indicate how much they would like to wear the clothing item shown (stimuli shown are for illustrative purposes only and are not part of the stimulus set shown - see Methods) given the goal on a scale from 1 through 7 (1; don't want to wear at all, 7; want to wear a lot). We show participants the same item across all three goal-contexts to measure context-dependent shifts in valuation. **B)** One trial consists of four stages. Trials start with an inter-trial interval, followed by the item (t-shirt here) and the current goal-context. After observing the item and goal-contexts, participants are asked to indicate their rating. This rating is reported by navigating around a circular scale (current position indicated with a green highlight) in a clockwise or counter clockwise direction and confirming the rating with a button press. This manner of rating scale separates the magnitude of the rating from the number and direction of button presses. After confirming their rating, trials are concluded by showing participants the recorded rating (e.g., "2"). Image order was fully randomized and trials under a certain goal were presented in a blocked manner (15 trials before switching). Each fMRI run contained exactly one block of each goal. T-shirt shown is for illustration purposes only and was not part of the stimulus set (see Methods for details). **C)** Schematic representation of the three constructs under investigation: raw attributes (pink schematic), attributes in value space (red schematic) and integrated value (green schematic). We show two example stimuli (t-shirt and pants) and schematically represent the transformations these items undergo. Firstly, in raw attribute space (here shown in 3D: "warmth"x"formalness"x"comfort") items location (and thus the raw attribute axes) are context-independent. Item location indicates the presence of a certain attribute in the clothing item. **D)** The raw attribute vectors then undergo a context-dependent transformation to 'attributes in value space'. 'Stimulus @ context' indicates the context-dependence of these new axes. Arrows are shown to schematically indicate this transformation from raw attributes to attributes in value space. Item location in 3D-attribute-in-value-space indicates not only the amount of an attribute that is present in the clothing item but also the current goal-dependent value of this attribute. As such attribute in value space vectors are different across different contexts (here shown island and job contexts). **E)** Lastly, item's attributes-in-value-space vectors are summed together and moved to a common rating-scale. We schematically show this subjective integrated value axis with each of the example items under each goal-context. The integrated value is directly measured in the task through the participant stimulus ratings (see **A** - on a scale from 1 through 7). Note here that the same item shown under different goal-contexts occupy distinct locations along the integrated value axis to indicate the goal-context dependency of the integrated value ratings.

Online: $H=249$, $p = 8.74e^{-55}$). However, across-goal differences were inconsistent for festiveness (Kruskal-Wallis; fMRI: $H=5$, $p = 0.07$, Online: $H=28$, $p = 7.08e^{-7}$) and comfort (Kruskal-Wallis; fMRI: $H=32$, $p = 9.02e^{-8}$, Online: $H=1$, $p = 0.51$). Thus, goal-dependent attribute transformation occurs robustly, but not for every attribute. Note that in the online study (see Figure 3.1C) we did not obtain subjective attribute ratings. In this sample we therefore instead used the stimulus-specific attribute means across fMRI subjects for this analysis. We additionally replicated these results with weights fitted using a logistic regression model (see supplemental Figure A.2G), with the only qualitative difference being a significant effect now also emerging for comfort in the online group ($p = 0.03$, see supplemental Table A.2 for all statistics).

Integrated value predictions based on high-level, but not low-level, attributes generalize to left-out stimulus sets.

Next, we aimed to evaluate the degree to which our a priori (high-level) raw attributes can predict value ratings across individuals and stimuli. If our set of attributes has a generalizable effect on integrated value we can feel confident we have indeed selected a sufficient set for our following analyses. Additionally, across-subject generalization would let us conclude there is a level of communality to the computations underlying valuation across subjects in our sample, which will be helpful in analyzing the neural data specifically. Previous research has shown that low-level visual attributes can also contribute to valuation of multi-attribute stimuli. We therefore also extracted a set of data-driven low-level attributes using computer-vision methods following a previously validated approach (see Methods and Iigaya et al., 2021). For comparison we tested models with 1) the low-level attributes alone, 2) the high-level attributes alone, and 3) both low- and high-level attributes. We hypothesized that high-level attributes, by virtue of their more abstract nature, possess a better capacity for generalization across individuals compared to low-level attributes.

To investigate the distinct contributions of high- and low-level raw attributes, we employed a cross-validation (leave-one-subject-out) elastic net model across both experiments (see supplement Figure A.3A and B for separate results). Across held-out participants, we see that a model including the high-level raw attributes alone, predicts equally well as a model including all raw attributes (i.e., both high and low level) (spearman correlation: mean $\rho=0.698$ and $\rho=0.697$ for all and just including high-level attributes respectively; Mann-Whitney U test: $\text{stat}= 8387$,




Figure 2.2: Behavioral results. See continued caption on next page.

Figure 2.2: (continued) **A)** Schematic representation of the data used for the behavioral model fitting procedure. Participant's ratings (i.e., subjective integrated value) for the same item will vary across goal-contexts while the raw attributes are context-independent. We expect differences in the context-specific weights ($w_{warmth}, w_{formalness}, \dots, w_n$) to capture the differences in individual raw attribute contributions to the overall within context stimulus ratings. As such, we are plotting these raw attribute weights, fitted in each of the three contexts, below. **B)** Attributes' influence on fMRI participant's ($n=35$) ratings varied greatly depending on the goal. For each participant we fit a GLM to predict participants ratings under each goal-condition. We included each of our apriori raw attributes as regressors to probe their goal-specific contribution to the overall integrated value rating reported by the participant. We are plotting the extracted beta-coefficients for each of these raw attribute-regressors. The resulting distributions of beta-coefficients across subjects are shown in box and whisker plots. For each box and whisker plot, the middle line denotes the median, and the bottom and top edges of the box indicate the first quartile (Q1) and third quartile (Q3) of the distribution. The ends of the whiskers represent the maximum and minimum data points not considering outliers. Outliers are shown in open circles and are denoted as such when they surpass either the lower bound ($Q1 - 1.5 * (Q3 - Q1)$) or the upper bound ($Q3 + 1.5 * (Q3 - Q1)$). Across subject means are reflected in bold horizontal lines. Additionally, individual subject model coefficients are plotted with the significance of the coefficient indicated through its color fill (white for non-significant and black for significant regressor coefficients; threshold used $p < 0.05$). **C)** Results for the identical analyses conducted on the online sample ($n=95$). We observed a highly similar pattern in raw attribute contribution across subjects and goal-contexts. Box and whisker plots mirror those described above and individual regressor coefficients are similarly shown as scatter (with fill indicating significance as above). **D)** Model evaluation results for raw attribute-based value prediction. Box and whisker plots reflect the spearman correlation between the model predicted value rating and subjective ratings of a left-out subject. We trained an elastic net regression model to predict subjective value ratings based on varying sets of raw attributes. Darkest boxes reflect model performance for models including all raw attributes. Middle boxes indicate the performance of models trained on the a priori raw attributes, while rightmost light colored boxes show the performance of models exclusively trained on computer vision extracted low-level attributes. Models are trained across all (both samples) but one left out subject, whose ratings are subsequently predicted by the model. This procedure is repeated 130 times for each model, such that each subject is left-out once. Model performance is statistically identical across models including all and high-level a prior attributes (Mann-Whitney U test: ski; stat=8387, $p = 0.92$, island; stat=8457, $p = 0.99$, job; stat=8424, $p = 0.97$). Model performance is however statistically different for models including only low-level computer vision extracted attributes (Mann-Whitney U test: ski; stat=1595, $p = 3.58e^{-35}$, island; stat=13896, $p = 2.66e^{-19}$, job; stat=15262, $p = 2.75e^{-29}$). **E)** Similarly to above we evaluate model performance across models trained on various sets of attributes. Here however, the data is split across genders. The model is trained on one gender, while the left-out subject is from the opposite gender. Importantly, subjects with different genders observe two unique sets of stimuli allowing us to evaluate our models on a completely left-out stimulus set. Spearman correlations are computed for the model predicted value rating versus subjective value ratings for each left-out opposite gender subject. Equivalently to D, each box and whisker plot includes 130 models (each subject is left out once). Across-gender models trained on all raw attributes and those trained on high-level a prior raw attributes perform equally well (Mann-Whitney U test: ski; stat=8616, $p = 0.78$, island; stat=8995, $p = 0.37$, job; stat=8875, $p = 0.48$). We again do observe a significant difference in performance between models trained on high-level raw attributes and those trained on only low-level raw attributes (Mann-Whitney U test: ski; stat=16764, $p = 8.57e^{-43}$, island; stat=16896, $p = 4.16e^{-44}$, job; stat=16900, $p = 3.80e^{-44}$).

$p = 0.92$ for ski; $\text{stat} = 8456.5$, $p = 0.99$ for island; $\text{stat} = 8424$, $p = 0.97$ for job). This stands in contrast to models trained only on the low-level raw attributes, which yield a significantly worse performance than those trained on high-level raw attributes (Mann-Whitney U test: ski; $\text{stat} = 1595$, $p = 3.58e^{-35}$; island; $\text{stat} = 13896$, $p = 2.66e^{-19}$; job; $\text{stat} = 15262$, $p = 2.75e^{-29}$, see Figure 3.1D). As such we conclude that the high-level attributes are better able to predict value than low-level attributes and do so in a manner that captures similarities across participants.

Next, we aimed to further investigate how well a model trained to predict value ratings, based on either high or low-level attributes, can generalize such predictions. Here we test these models' value predictions for both an unseen stimulus set, and in a completely separate group of individuals. We leveraged the fact that male and female identifying participants were presented with two distinct stimulus sets with male and female gendered clothing items respectively. We exploited this design feature by training our model on one gender and testing it on the ratings from individuals in the opposite gender group, who thus saw the other set of stimuli. Comparing the generalizability of high-level and low-level attributes through this approach shows a stark difference. While a model trained on the high-level raw stimulus-attributes generalized well across stimulus sets (and thus genders), a model trained only on the low-level raw attributes failed to generalize (see supplement Figure A.3D and E for separate online and fMRI subject results). This difference was robustly significant (Mann-Whitney U test: ski; $\text{stat} = 16764$, $p = 8.57e^{-43}$; island; $\text{stat} = 16896$, $p = 4.16e^{-44}$; job; $\text{stat} = 16900$, $p = 3.80e^{-44}$, see Figure 3.1E). These findings indicate a key advantage of high-level attributes and highlights their strength in predicting the value of stimuli. By virtue of the abstraction of these attributes from the basic visual properties of a stimulus, they are able to capture relationships between attributes and value that generalize well across stimulus sets and groups of individuals. To gain insight into the neural representation of these raw attributes, we next turned to the fMRI data.

Visual cortex represents low- and high-level stimulus-intrinsic attributes.

We expect the brain to separately represent each of the three constructs under investigation (raw attributes divided into low and high-level, attributes in value space, and integrated value ratings), as they make up distinct component processes underlying flexible valuation of multi-dimensional stimuli. We first tested the encoding models of raw attributes. We hypothesized that, even though high-level attributes

are sufficient to predict valuation, the brain represents both low-level and high-level attributes. Additionally, we expect such representations to be context-independent since the raw attributes are static properties of the stimuli.

The whole brain analysis shows that these raw attribute representations are concentrated in the visual cortices (F-test across high-level attributes shown in red in Figure 2.3B - with max voxel for high-level raw attributes: [56,15,32], cluster corrected FWE<0.05, voxel threshold $p < 0.001$ for all reported clusters). Consistent with previous findings (e.g., Igaya et al., 2023), early visual regions tend to encode low-level attributes more prominently than high-level attributes. In contrast, we observe that higher visual regions more strongly represent high-level visual attributes (see Figure 2.3B and Figure A.5).




Figure 2.3: **A)** Since raw attributes are context-independent we can model their presence (e.g., the amount of warmth provided) in one model across contexts. In practice, we test for such context-independent raw attribute representations through fitting an encoding model. We observe that both computer-vision extracted attributes (F-test across attributes shown in pink) and participant-rated attributes (F-test across attributes shown in red) are represented in overlapping parts of the occipital cortex (cluster corrected FWE<0.05, voxel surpassing $p < 0.001$). **B)** When plotting the proportion of significant voxels (before cluster correction) that represent low or high-level attributes across ventral-visual cortex ROIs we observe a gradually more prominent representation of high-level attributes along the posterior-anterior axis.

Frontal cortex is a candidate region to hold attributes in value space signals.

We next aimed to test the hypothesis that the brain represents attributes in value space (see Figure 2.1D), in addition to the stimulus-intrinsic raw attributes shown to be encoded above. If a brain area encodes raw stimulus attributes then the representation of that attribute would not be expected to change as the goal context changes. In other words, the raw attribute would be represented in the same way across all goals as it should be goal-context independent (see Figure 2.4A). Alternatively, if an attribute is being represented in value space then it would be expected to exhibit differences in its representation as goal-context changes. Such context-based representational shifts are expected especially in the case where the attribute is relevant for the goal-value computation (such as in the case of the warmth attribute when computing the value of a ski-jacket vs a swimsuit) (see Figure 2.4A). To distinguish between these possible coding schemes, we set up an attribute decoding generalization analysis. We first trained a regularized classifier to decode each of our subjectively rated raw attributes under a specific goal. For example, we might train a model to classify whether clothing stimuli shown under the ski goal are high or low in warmth (median split, for more details see methods). We can then test this model on the same goal (ski) or on one of our two other goals (job interview or island vacation).

If we find that a model trained on one goal's raw attribute ratings generalizes well to another goal this means the underlying neural activity pattern across goals is consistent, as expected when representing a raw attribute. For example, the brain signal encoding the raw attribute of high warmth would be similar for both the ski trip and island vacation goal-contexts. Such similarity in the neural representation of a raw attributes across goals would exclude the possibility of the attribute being represented in value space. Alternatively, if a decoder successfully decodes a raw attribute in the training context, yet performs poorly when tested on another goal context, that would indicate that the neural representation of the raw attribute has shifted from one goal to the next (see Figure 2.4A). Such a decrease in decoding accuracy would be expected if the area computes attributes in value space, as this is a goal-dependent signal. Put more concretely, if the model trained on warmth under the ski goal has poor performance when tested on the tropical island goal this would be indicative of distinct neural coding axes for high vs low warmth across these goals. Evidence of distinctive neural representations across goal contexts would be suggestive of the type of goal-context-dependent modulation of neural activity necessary for the brain representation of attributes in value space. We can then

leverage this evidence to further investigate value-dependent neural codes.

Probing the generalization performance of the two attributes that displayed the most homogeneous behavioral pattern across subjects (formalness and warmth, see Figure 3.1), we observe consistent significant decoding both within the training and each of the test-contexts for the two visual ROIs (see Figure 2.4B). However, for the frontal ROIs (LPFC and MPFC) we observe significant differences in accuracy between the training and test contexts (see Figure 2.4C). Additionally, we find that in all test-contexts, this decrease in performance is so drastic the raw attribute can no longer be significantly decoded in the test context. These frontal ROIs thus display a shift in the representation of raw attributes from one goal to another, thereby meeting the criteria we identified as being necessary for brain areas to represent attributes in value space (see Supplemental Figure A.6 for other attributes).

We additionally probe the generalization performance across all train-test context combinations. We find the out-of-context decoding accuracy (see Figure 2.4D&E - in red) is significantly lower than the within context decoding accuracy in lateral and medial prefrontal cortex (see Figure 2.4D&E - Wilcox test; formalness: LPFC: $p = 0.0009$, MPFC: $p = 0.0005$; warmth: LPFC: $p = 2.27e^{-6}$, MPFC: $p = 2.87e^{-5}$). We ran these same generalization analyses for ventral and dorsal visual stream ROIs. Given our previous encoding analysis finding these regions represent raw attributes (see Figure 2.3A), we expect to find good generalization in these ROIs. We indeed observe no differences in decoding performance for raw attributes across the ventral visual-temporal (VT) and dorsal visual-temporal (DT) ROIs (Wilcox test; formalness: VT: $p=0.20$, DT: $p=0.10$; warmth: VT: $p=0.29$, DT: $p=0.10$).

Prefrontal cortex represents attributes in value space through encoding of model estimates.

While the generalization analysis confirms context-dependent attribute representations exist, this does not yet prove that this is a value-related context dependence. We therefore turn to our earlier behavioral models which computed the relationship between raw attributes and participants' stimulus ratings across goal conditions. From these models we obtained estimates (i.e., β weights) of the transformations necessary to compute attributes in value space - i.e., the transformation of raw attributes into value space. We now directly apply these estimated transformations to compute estimates of attributes in value space. To achieve this we multiplied subject and goal specific behavioral GLM β -weights by the subjective raw attribute ratings




Figure 2.4: Generalization analysis for attributes in value space. See below for continued caption.

Figure 2.4: (continued) **A)** Schematic representation of model performance interpretation. If context-independent raw attributes are represented in an ROI we would expect good generalization performance. Since raw attribute representations are context-independent the hyperplane fitted in one context (e.g., separating high and low warmth items under island) will similarly separate raw attributes under another goal-context (i.e., job). However, if instead an ROI represents context-dependent attributes in value space we expect poor generalization performance. As shown in the example schematic, the value axes shift between contexts and the items shift position accordingly. Therefore a hyperplane fit to separate items in one context (e.g., island) will be unable to separate those items under a different context (e.g., job). **B)** Within subject ventral and dorsal visual cortex generalization decoding accuracies for formality and warmth attributes. The goal-context used for training is shown in the left most position (as annotated on the x-axis) and dashed lines show permuted chance level (mean of null) (see Methods). Statistical significance of the context-specific model performance is computed with respect to each context-specific null and any significant results are indicated (*: $p < 0.05$). We additionally compute statistics for the pairwise comparison between the training and each testing context. The significance of these between-context decoding accuracy differences are similarly indicated (*: $p < 0.05$). For both ventral and dorsal visual cortex ROIs we observe significant within and out-of-context model performance in addition to no significant differences in decoding accuracies between training and testing goal-contexts for both formality and warmth attributes. **C)** Similarly, we show the generalization decoding for lateral and medial prefrontal cortex (LPFC, MPFC) ROIs. Here we do observe model performance is not different from chance in the out-of-context test sets. Additionally, we do observe significant differences in decoding accuracies between training and testing goal-contexts. This implies a shift in the underlying neuronal representation of the attribute across goal-contexts. **D)** Formality decoding accuracy computed for each of our ROIs (VT: ventral temporal, DT: dorsal temporal, LPFC: lateral prefrontal cortex, MPFC: medial prefrontal cortex). Decoding accuracy of out-of-context models is shown in red bars with within-context model performance shown in blue for comparison. Importantly, this figure includes model performances across all trained and tested models. Within-context models thus include those for the ski, island and job contexts with each of these models evaluated across the other two contexts. Individual dots show within subject average performance across the three within-context models (under blue bars) and the within subject averaged out-of-context model performance (under red bars). T-tests are used to compare this out-of-context performance with the within-context performance for each ROI (**: $p < 0.01$). **E)** Similarly for the warmth models. We again observe significant differences in average model performance within contexts and out-of-context in frontal areas (LPFC and MPFC). **F)** The average encoding across estimates of attributes in value space is shown to be significant in various areas within our LPFC (outlined in orange) and MPFC (outlined in light blue). These clusters are focussed on areas including the vmPFC stretching into medial OFC, lateral OFC and more dorsal areas of LPFC. All results are cluster corrected within the ROIs for FWE <0.05 with voxels surpassing $p < 0.001$.

to compute estimates of subjective attributes in value space (i.e., $w_w \times \text{warmth}$). Attributes in value space should capture both the presence of an attribute in an item and the current value of that attribute. For example, a sweater will provide a lot of warmth (e.g., subjective raw warmth rating of 8) and when presented with this sweater under the goal-context of a ski trip, this warmth will be a large positive contributor to the overall item value (e.g., subjective β of 0.6, leading to an attribute value of $8 \times 0.6 = 4.8$). While for the same sweater under the island vacation goal-context a negative attribute value would be estimated (e.g., subjective β of -0.5: $8 \times -0.5 = -4$). These multiplications of fitted weights with raw attributes, thus leaves us with direct personalized estimates of attributes in value space. Next we will use these estimates to test for the encoding of attributes in value space using our neural data (see Methods for more details).

We restrict our analysis here to the areas shown to be candidates for representing attributes in value space in the generalization analysis above (LPFC and MPFC). When mapping out the clusters that correlate with our attribute-in-value-space estimates, we found significant clusters in both lateral and medial PFC (mean across attributes, see Figure 2.4F). When zooming into these large ROIs to delineate effects in individual sub-regions of MPFC and LPFC, we found significant encoding of attributes in value space in multiple PFC subregions including the medial and lateral orbitofrontal cortex (OFC) and parts of the ventromedial prefrontal cortex above the orbital surface (vmPFC) (max voxel in LPFC: [24, 80, 42], $z=6.4$). The abundant encoding of attributes in value space confirms that attributes in value space computations are indeed part of the neural mechanism underlying flexible value computation for multi-attribute stimuli.

Subjective value ratings are represented across-contexts in widespread prefrontal cortex ROIs.

We next turn to investigate the neural representation of the last construct under investigation; subjective integrated value ratings. First, we use the context-dependent stimulus value ratings reported by participants to train a decoding model (see Methods). We found that we could significantly decode participant ratings (i.e., integrated value) across all prefrontal sub-ROIs (see Figure 2.5C - and see supplemental Figure A.7B&C for visual ROI decoding results). Additionally, we probed the generalizability of the stimulus value ratings signal. To facilitate decision making (i.e., to make a comparison between options), we expected that the neural representation of




Figure 2.5: Integrated value results. See below for continued caption.

such value ratings should generalize well across goal-contexts, as it has been abstracted away from the specific stimulus and goal under consideration to a common axis for valuation. We indeed find good generalization of the integrated stimulus value signal across contexts in frontal ROIs (see Figure 2.5B - and see supplemental Figure A.7A for visual ROIs). Importantly, the ability of these frontal ROIs to represent a generalizable value rating signal also confirms the goal-specific raw attribute results discussed above do not stem from an inherent property of the ROI to switch representations across contexts.

Similarly to the approach for attributes in value space, we next turned to model the integrated value (operationalized as participant stimulus ratings) in an encoding analysis performed in the LPFC and MPFC. We found participant ratings correlate with the neural signal along the medial wall of the prefrontal cortex stretching towards the anterior cingulate cortex (max voxel in LPFC:[56,81,28], $z=6.5$ - see

Figure 2.5: (continued) **A)** As items' attributes in value space have been summed together to compute integrated values we once again expect good cross-context generalization. Importantly, at the level of integrated value identical items will have different ratings across contexts we expect items with a similar rating to be similarly represented in the brain. As such, at this level of computation, individual attributes of the clothing items are no longer represented and instead a common value scale allows for comparison across items and contexts. **B)** We indeed find good generalization of participants' stimulus ratings (i.e., integrated value) across areas of prefrontal cortex. There are no significant between context differences in decoding accuracies and models trained on context show significant model performance on the two left-out contexts. Models trained on ski-context are shown as exemplars here (see supplemental figure x for other training contexts). **C)** In addition to a generalization decoding analysis, we performed integrated value (operationalized as stimulus ratings) decoding across the whole task. This decoding analysis was performed across sub-ROIs within the prefrontal cortex. We find universal significant representation of integrated value across frontal-cortex ROIs (***: $p < 0.001$). **D)** An encoding analysis confirms stimulus ratings are significantly represented in clusters that span areas of dorsomedial prefrontal cortex and ventral-lateral prefrontal cortex. We additionally find participants' stimulus ratings are encoded in lateral OFC. The encoding model was ran within LPFC (outlined in orange) and MPFC (outlined in light blue) ROIs and result are cluster corrected within the ROIs for FWE<0.05 with voxels surpassing $p < 0.001$.

Figure 2.5D). In the lateral prefrontal cortex, we find the strongest representation of participant value ratings in posterior parts of the ROI. Lastly, we find significant value rating encoding in lateral orbitofrontal cortex (OFC).

Overlapping and separable encoding of subjective value ratings and attributes in value space in prefrontal cortex.

We have found evidence that the brain represents both attributes in value space and subjective value ratings. Here we turn to address the open question of how the computation of these two value signals interact. We defined integrated value as the linear sum over attributes in value space (see Figure 2.1E and 3.1A and in line with previous research e.g. Iigaya et al., 2021). Given this close relationship between these two constructs (one being the sum of the others) we might predict that the neural representation for these two constructs is similar as well. However, if we cannot find separate neural correlates for these two value computations we might not feel confident they are indeed two separate steps in hierarchical multi-attribute value computation. We therefore tested the extent of the similarity between

integrated value ratings and attributes in value space. We employ a voxel-wise model in which the subjective stimulus ratings regressor directly competes with attributes in value space regressors to explain variance in the BOLD signal. In the most extreme case where attributes in value space are not computed independently from integrated value, we would expect to see the same voxels represent both variables. Alternatively, if we observe significant effects of attributes in value space or subjective integrated value ratings in this model these results would reflect each term's unique variance (see Methods for more details). Such independent neural representation would imply that there is variance unique to the subjective stimulus ratings provided by our participants that cannot be captured by the equal linear combination of the four included attributes in value space.

Using this combined model, we find substantial significant clusters of both subjective stimulus ratings and attributes in value space. Additionally, these clusters appear to be located in areas similar to those found in the models that focused on each construct separately (see Figure 2.6). We observed a similar ventral-dorsal separation in the prefrontal cortex ROI along the medial wall, with ventral correlates to attributes in value space. This ventral-dorsal pattern appears to be reversed within the lateral prefrontal cortex, in that the stimulus rating cluster occupies the posterior ventral LPFC. One major difference to note here is that in the combined analysis we no longer find significant stimulus rating encoding in lateral OFC but solely find evidence for attributes in value space across OFC regions. This suggests the lateral OFC representation of stimulus value rating (see Figure 2.5D) is now better explained by the attributes in value space regressors. This finding hints at the high similarity between the neural signal for integrated value ratings and attributes in value space in the region.

However, the extensive independent encoding of both attributes in value space and integrated value suggests these are separate neural computations. Part of this separability could originate from the great many raw attributes over which participant ratings (i.e., integrated value) might be computed, since we are necessarily only examining a small subset of such attributes. Thus, a linear combination over the attributes examined in the present study may not be sufficient to capture the full integrated value signal as it is represented in each individual. Accordingly, a difference between these representations may simply relate to this unexplained variance in integrated value arising from attributes outside the scope of our consideration.




Figure 2.6: **A)** Encoding results of a model including both integrated value and attributes in value space. Since these results are from one unified encoding model they can be directly compared. Additionally, only variance unique to each of these two constructs is modeled here. Interestingly there are a lot of clusters that remain similar compared to Figure 2.4F and Figure 2.5D hinting at substantial unique variance in the neural representation of these constructs. However, we also observe differences, with most noteworthy a lack of integrated value encoding in the lateral OFC in favor of attributes in value space encoding in the area. LPFC and MPFC ROIs are outlined in orange and light blue respectively. All result are cluster corrected within the ROIs for FWE<0.05 with voxels surpassing $p<0.001$.

Discussion

This study investigated the neural mechanisms underlying the computation of integrated value ratings for multi-attribute objects, a process central to flexible decision-making. We set out to adjudicate between two competing computational architectures: a direct integration model where goals merely adjust the weights applied to raw stimulus features, and a hierarchical model proposing the existence of “attributes in value space” as a critical intermediate computational step. While the first and final steps of valuation have been the topic of much previous work (e.g., Kahnt et al., 2011, Iigaya et al., 2020), the intermediate computations necessary to flexibly transform a visual attribute to an object-based value judgment had not been directly contrasted. Through a combination of behavioral experiments and fMRI, our findings reject the simpler direct integration hypothesis and provide significant support for a hierarchical model of value construction, demonstrating distinct neural representations for raw attributes, attributes in value space, and integrated value ratings.

Our behavioral results clearly demonstrate the dynamic nature of integrated value. Participants significantly modulated their liking ratings for the same clothing items based on the instructed goal-context (ski trip, tropical island, or job interview). Successfully demonstrating a goal-dependent shift in valuation is necessary to distinguish between stimulus-intrinsic attributes and their value. Employing in-

dividualized attribute models, we derived estimates of attributes in value space through goal-specific weighting of raw attributes. Crucially, we found an important distinction between low-level attributes and high-level, subjectively-rated attributes wherein only high-level attributes generalized well to left-out data. This finding both speaks to the predictive power of high-level attributes in capturing general valuation, and provides further validation for a hierarchical model of value construction.

Building directly on this behavioral evidence, our fMRI data revealed a distributed neural architecture that specifically supports this hierarchical computation rather than a direct goal-weighted integration of raw features. Raw attributes were robustly encoded in the visual cortex, consistent with their stimulus-intrinsic nature. If the brain relied solely on direct integration, we would not expect to find goal-modulated representations of individual attributes prior to the computation of overall object value. Consistent with our intermediate hypothesis, however, we found that regions within both lateral and medial prefrontal cortex demonstrated a significant shift in the representation of attributes across different goal-contexts. This context-dependent neural coding, whereby a model trained on an attribute under one goal performed poorly when tested on another, precisely matches the expected behavior of attributes in value space—signals that are dynamically modulated by current goals before final integration.

We then sought to provide further neural evidence supporting the novel computational step we put forth as part of multi-attribute valuation. We directly investigated the presence of attribute-encoding in value space, alongside the existence of an integrated value signal within the prefrontal cortex. We found evidence for separable encoding of these different value representations. Integrated value signals were predominantly found along the medial wall of the prefrontal cortex (stretching towards the anterior cingulate cortex), while attributes in value space showed a more ventral representation along the medial wall and a dorsal-to-anterior pattern in lateral prefrontal cortex. We additionally found significant encoding of attributes in value space in regions including the medial and lateral orbitofrontal cortex (OFC). This distinct neural topography for attributes in value space and integrated value, however supports the possibility that these are two separable computational components in the flexible valuation process.

As it was not our aim to be exhaustive in the set of attributes to decompose our stimuli into, this additional integrated value encoding might represent other attributes in value space not included in our set. Contrasting the neural representation of these

attributes in value space with integrated value additionally opens up future research questions about the composition of integrated value. We hypothesize the potential of a more distributed representation of individual attributes in value space (where individual attributes in value space are encoded by partially overlapping but distinct sub-regions of the ventral prefrontal cortex), with this information combined into an integrated value signal for decision making and comparison. The fact that integrated value signals generalize well across goal-contexts further enforces its separable and abstracted nature, serving as a unified decision variable.

While early theoretical work by Winterfeldt and Fisher, 1975 posited the potential for computing single attribute utility functions, direct neurobiological evidence for such value-like computation at the individual attribute level has remained elusive. Our findings, by demonstrating a neural correlate for attributes in value space, provide a critical missing piece. This work connects with previous research across various sub-domains of value computation; for instance, McNamee, Rangel, and O'Doherty, 2013 reported a ventral-to-dorsal separation along the medial wall of the prefrontal cortex in their study of category-dependent value signals. In their work, participants' subjective value for objects within specific categories (e.g., food, trinkets) was reflected in more ventral mPFC regions, while a more dorsal mPFC area held a category-independent value signal. Analogously, our study also reveals a gradient along the mPFC, where more ventral regions represent individual attribute values in a goal-dependent manner, and more dorsal regions encode the integrated value of overall objects. Both sets of results collectively underscore the brain's capacity for hierarchical abstraction, moving from specific, context-dependent object properties to a more generalized, comparable value signal to facilitate efficient decision-making.

Our findings are broadly consistent with a recent study by Castegnetti, Zurita, and DeMartino, 2021, which observed that activity in the ventromedial prefrontal cortex (vmPFC) correlated with an item's "usefulness" for the current goal, and that the neural code supporting this usefulness representation in the vmPFC was maintained and generalized across different goals. This finding resonates with our own report that integrated value signals, predominantly found along the medial wall of the prefrontal cortex, generalize well across goal-contexts, serving as a unified decision variable. In parallel, Castegnetti, Zurita, and DeMartino, 2021 also showed that while the dorsolateral prefrontal cortex (dlPFC) represented usefulness, its coding was specific to the current goal and did not generalize across contexts. The

present study builds on those previous findings by providing a mechanistic account of how a context-specific value code can be produced in the first place through the transformation of attributes from raw stimulus features through to context-dependent attribute values which ultimately yields an integrated context-specific value representation. Collectively these findings shed light on the sophisticated interplay within the prefrontal cortex, where some sub-regions maintain generalized value signals, while others provide highly specific, context-modulated attribute representations, collectively underpinning the brain's capacity to construct value on the fly and enable adaptive behavior in dynamically changing environments.

It is important to acknowledge several limitations of the present work. First of all, while our experimental paradigm successfully manipulated goal-context, the scope of attributes and goals investigated was constrained. Other attributes will undoubtedly vary with our goal manipulation and many more will vary when considering other goals. However, given the overlap found across stimulus types for similar valuation processes, we expect our findings to generalize to other attributes, object classes and goals. Follow-up research could explore the specific intricacies that might exist under other conditions/objects/goals. Additionally, our research does not address how goal-dependent (attribute) values are formed over time. We might hypothesize that experience with items consisting of similar attributes under various conditions/contexts could inform current attribute valuation. Probing the mechanisms of such a learning process on the attribute-level could be an important direction of future research. Lastly, our findings on the relationship between attributes in value space and integrated value suggest a combination of overlapping and unique representations. Probing the extent of this overlap at various levels of neural computation (i.e., voxels to individual neurons) will require further research.

This study provides a comprehensive account of the neural computations underlying flexible value construction. By proposing and providing evidence in support of the existence of "attributes in value space", we bridge the gap between static stimulus properties and dynamic, goal-dependent object valuation. Our findings highlight a hierarchical process, where sensory areas represent raw attributes, and prefrontal regions transform these into goal-sensitive attributes in value space, culminating in the computation of an integrated value signal that is robust and generalizable across contexts. This work advances our understanding of how the brain constructs flexible subjective value, a fundamental process for navigating a complex and ever-changing world.

Materials and Methods

fMRI experiment

Thirty-five participants (age 18 to 64 years, twenty female) participated in this experiment. Participants were recruited from the greater Pasadena (CA) area and were screened for MRI compatibility. Participants provided consent to participate in the research protocol approved by the California Institute of Technology Institutional Review Board (IRB number: IR20-1033) . Participants were paid \$55 on average. All participants completed the study successfully and are included in the analyses.

MRI task procedure Participants first answered a short demographics questionnaire after which they received the task instructions. These instructions included a more elaborate description of the goal-context to ensure a common understanding across participants. Following the task instructions, a set of ten practice trials was completed by each participant. During these practice trials participants were shown stimuli and a goal-context that were not used in the main experiment. This completed the pre-scanner portion of the experiment and participants were placed inside the scanner. When in the scanner participants were invited to browse through the instructions one more time after which the main portion of the experiment was started. Two sets of stimuli were created each consisting of various clothing items, one of which was female and the other male gendered. Stimuli are collected from the DeepFashion dataset and selected to not include any persons (see Liu et al., 2016). Participants saw the set of stimuli that best matched their gender identity. In addition to these clothing stimuli a symbol and text label identifying the current goal-context were presented. Trials were presented in blocks of fifteen, within which the goal remains constant. Participants were shown the clothing item stimulus and a reminder of the goal, after which participants reported how much they would like to wear the clothing item with the current goal in mind (see Figure 2.1). Three blocks of fifteen images within a context were shown consecutively within an fMRI run. Each such fMRI run took 7 minutes and a total of five runs was completed by each participant (one participant completed an additional sixth run). A circular rating screen was used on which the rating 1 through 7 was displayed. The participants could navigate around this screen by pressing a left and right button to move either counterclockwise or clockwise respectively. Once they had reached the number corresponding to their intended rating they pressed a third button to confirm their rating. While the order of numbers and the starting location to navigate from was kept consistent the numbers were rotated around the screen randomly every trial.

This setup helps dissociate between button pressing and the value rating inputted. For a traditional linear scale the ends of the scale (more button presses) are often associated with more extreme ratings. Additional correlation can arise since pressing left is often associated with the lower end of the scale while right is associated with higher ratings. Using the circular scale avoids these potentially confounding effects in the decoding analyses.

High-level attribute ratings After completing the in-scanner task, participants were asked to give high-level attribute ratings for the stimuli used. This rating was self-paced and took around 20 minutes to complete. The high-level attributes rated were ‘warmth’, ‘formalness’, ‘festiveness’, ‘youthfulness’ and ‘comfort’ (youthfulness attribute was excluded from analyses due to correlation with other attributes - see supplement Figure A.2F). Participants rated each stimulus they saw during the experiment on five high-level attributes on a rating scale from 1 to 10. One subject had no variance in their provided comfort ratings across stimuli, this subject is therefore excluded from any comfort-attribute based models.

Online experiment

The online sample consisted of 95 participants (57 female) and was collected through Prolific. Online participants were paid approximately \$5 on average.

Online task procedure Participants performed a rating task under four instructed goal-contexts (tropical island vacation, ski trip, job interview and summer wedding). Participants saw 100 unique, gender-matched stimuli total, resulting in 400 total trials (see Figure A.1 for online design). Participants indicated how much they would like to wear the item shown given the goal-context on a scale from 1 through 9 (1; don’t want to purchase at all, 9; want to purchase a lot). This rating was reported using the 1 through 9 keys at the top of a participants keyboard. Image order was fully randomized and trials under a certain goal-context were presented in a blocked manner (20 trials before switching). Trials were self-paced and participants moved on to the next stimulus-context combination when a rating was entered. The fourth goal condition included in the online experiment was not in the fMRI study due to time constraints in the scanner. This goal-condition is therefore excluded from further analyses.

Behavioral analyses

Note that integrated value and raw attribute ratings are independent of gender therefore data from both genders was analyzed jointly (see Supplemental Figure A.2B, D&E).

Low-level attribute extraction Low-level computer vision attributes were extracted using the same methods as those used in Iigaya et al., 2021. These attributes capture both segment-specific and overall stimulus characteristics. Please see Iigaya et al., 2021 and Li and Chen, 2009 for further details.

Mixed-effects modeling: attribute main effects For each attribute in each context independently we ran linear model predicting value ratings. This model included the attribute and a random intercept for each subject (value = attribute + (1|subjectID)). We then evaluated the significance of the attribute coefficient. This set up resulted in 4 attributes x 3 contexts models.

Beta weight estimation modeling High-level raw attribute beta weights (shown in Figure 3.1B&C) were fitted within subject for each context separately. The model used in this analysis included the four high-level attributes and was predicting the object value ratings of the participant and context in question. Since the online participants did not provide subjective attribute ratings, we instead used the stimulus' mean attribute rating provided by the MRI participants. Models were fit separately within subject and context to ensure individual and within goal differences were maximally captured. This is especially important since we would expect such inter-individual differences to also be reflected on a neural level. For example, for a participant whose value rating is highly dependent on the comfort of the item we would expect a stronger comfort representation than for a participant whose value ratings are independent of comfort.

Since the distribution of value ratings displayed a bimodal pattern we additionally repeated this analysis using a logistic regression model on binarized (median split) value ratings. The beta weights fit using this approach are shown to correlate strongly with those fit using the linear model (see Supplemental Figure A.2G).

Leave-one-out behavioral generalization modeling We used an elastic net model (alpha=0.2 and ratio=0.5) trained to predict value ratings on all but 1 left-out participant. The model was then used to predict value ratings for the left-out participant. The correlation between these predicted ratings and the true ratings was used to evaluate the model. This procedure was repeated for each participant in the sample.

Similarly, we trained an elastic net model to all participants within one gender-group and tested on one of the opposite-gender participants. This process was repeated until each subject had been left out once.

Neural analyses

Data preprocessing fMRI data were preprocessed using a custom pipeline in the Nipype (Gorgolewski et al., 2011) environment. Functional images were extracted using functions from the Nilearn (Abraham et al., 2014) toolbox and motion corrected using the MCFLIRT algorithm (Jenkinson et al., 2002) from the FMRIB Software Library (FSL; Jenkinson et al., 2012), which estimated six motion parameters to be included as covariates in the statistical model along with estimated motion outlier volumes. Other covariates were estimated using a principle components analysis on both the timeseries of white matter and CSF voxels, demarcated with the co-registered T1 structural image, as well as outlier timeseries using the CompCor algorithm (Behzadi et al., 2007). The functional data were unwarped to correct for distortion using FSL’s TOPUP algorithm (Andersson, Skare, and Ashburner, 2003) with reverse phase-encoded spin-echo EPI images. Both T1 and T2 structural images were corrected for field bias using the N4 algorithm (Tustison et al., 2010) and brain extracted using Advanced Normalization Tools (ANTs; Avants, Tustison, and Johnson, n.d.). Affine co-registration transformations were estimated between the functional image and the T2 as well as between the T2 and T1 structural images. A non-linear transformation was estimated to normalize the T1 structural to the MNI152 template (Mazziotta et al., 2001) using the symmetric normalization algorithm (Avants et al., 2008). All transformation matrices were estimated using ANTs, and applied to the functional data simultaneously to normalize to template space. Finally, the functional data were de-spiked and smoothed with a 6mm FWHM kernel (smoothing was applied for encoding analyses only).

Univariate encoding of raw attributes fMRI univariate encoding analyses were performed using FSL (Jenkinson et al., 2012). Custom onset files created regressors spanning the stimulus-on period. The raw attribute analysis included the three contexts as regressors in addition to the four high-level attribute regressors and four low-level attributes (average saturation of the 2nd segment, average brightness of the 2nd segment, mean saturation of the 2nd segment and entropy of the largest segment). We then constructed two independent F-tests over the four high-level and the four low-level attributes. F-test statistics were computed based on the group-

level results where voxels were thresholded at a z-score of 3.2 ($p \leq 0.001$) and cluster correction was performed ($p \leq 0.05$). The proportions of significant voxels across ventral and dorsal visual ROIs were computed as those passing the $p < 0.01$ threshold before cluster correction. ROIs were defined using the atlas from Wang et al., 2015.

ROI generalization decoding Single-trial β 's were extracted for each voxel from the un-smoothed neural data time-locked to the stimulus on period. These were then masked based on the ROI of interest. Visual ROI masks were created using Wang et al., 2015 atlas. Frontal ROI masks were created by merging regions specified in Tzourio-Mazoyer et al., 2020. More specifically, the IOFC ROI was created by merging the frontal mid orbital ROIs, frontal inferior orbital ROIs and the frontal superior orbital ROIs. The mOFC consisted of the frontal medial orbital ROIs and the rectus ROIs. The dmPFC ROI consisted of the frontal superior medial ROIs and the cingulum anterior ROIs. The dlPFC ROI was made up out of the frontal mid ROIs and the frontal superior ROIs. Lastly, the vlPFC consisted of the frontal inferior operculum ROIs and the frontal inferior triangularis ROIs. The larger LPFC and MPFC ROIs were created by merging the IOFC, dlPFC and vlPFC or mOFC and dmPFC respectively.

The generalization decoding was ran over each context combination, with each context considered the training context and each context considered as the test context (leading to a total of nine train-test combinations per variable of interest per ROI). Trials were split into training and testing sets using a 5-fold cross-validation scheme. Importantly, for the out-of-context decoding (training context is not equal to testing context) trials belonging to the training context were still subset using this approach to ensure equal training data for across and within context modeling. The decoding was performed using a ridge classification model. The dependent variable was median split to allow for binary classification. The average accuracy across folds was returned for model evaluation.

Test set group-level accuracy means were compared to permuted chance distributions. Chance distributions were computed through shuffling the dependent variable, retraining the classifier model and evaluating its performance using the same metrics as non-permuted. We computed 500 such shuffled model accuracies (across 5-folds each) for each train-test context pairing for each ROI and each attribute. Between context statistics were computed using T-tests for the differences for each train-test pair shown.

Additionally we computed subject-level decoding accuracy means across all train-test combinations for each ROI and attribute. We split these mean accuracies into those for within-context models and out-of-context models. The significance of the difference between within vs out-of-context models was computed using Wilcoxon signed rank tests.

ROI decoding of integrated value ROI-based decoding analyses were used to determine the presence of multi-variate integrated value representations across the whole task for the ROIs introduced above. For this analysis we again binarized the dependent variable (median split). We perform stratified 3-fold cross validation and trained a ridge classifier model for each fold. The mean accuracy across the cross-validation folds was saved for each ROI for each participant. Significance of the group-level mean was again compared to permuted null distribution. We again computed 500 permuted model accuracies for each individual (and 3-folds cross-validation within each shuffle). We then computed group level permuted accuracies for comparison.

Value-based encoding models Univariate integrated and attribute value regression models were set up as similarly as possible. All these models included three context regressors (ski, job and island context). In addition, the models of integrated value included a parametrically modulated regressor capturing either integrated value ratings reported by the participants (or behavioral model estimates for the integrated value - see supplemental Figure A.7).

For the attribute in value space model we use the same three context regressors in addition to 4 attribute in value space parametric regressors. These attribute in value space regressors were created by taking the subjective attribute ratings multiplied with subject specific β weights from the current-context's behavioral GLM. For example, if the trial stimulus is a jacket in the job interview context the β weights belonging to the attribute regressors in the job interview context GLM is used. Since the attributes in value space model included 4 regressors of in interest (one for each of our a priori high-level attributes) we additionally computed a contrast taking the mean across individual regressor effects. Similarly to the raw attribute encoding models, we thresholded individual voxels at $p < 0.001$ and cluster corrected to $p < 0.05$.

For the models including both integrated value and attributes in value space we combined the following regressors together: integrated value ratings, 4 attribute in value space regressors and three context regressors. The integrated value and the

4 attribute in value space regressors might be modeling overlapping variance. In extremes, two perfectly correlated constructs in one analysis should result in a null result when testing contrasts to identify one or other type of representation. As due to the extra sum of squares principle (Draper and Smith, 1998), all shared variance (i.e. which in such a case would be all of the variance) between these representations should have been partialled out.

References

- Abraham, A. et al. (2014). “Machine learning for neuroimaging with scikit-learn”. en. In: *Frontiers in Neuroinformatics* 8. ISSN: 1662-5196. DOI: 10.3389/fninf.2014.00014.
- Andersson, J. L., S. Skare, and J. Ashburner (Oct. 2003). “How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging”. en. In: *NeuroImage* 20.2, pp. 870–888. ISSN: 10538119. DOI: 10.1016/S1053-8119(03)00336-7.
- Avants, B. et al. (Feb. 2008). “Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain”. en. In: *Medical Image Analysis* 12.1, pp. 26–41. ISSN: 13618415. DOI: 10.1016/j.media.2007.06.004.
- Avants, B. B., N. Tustison, and H. Johnson (n.d.). “Advanced Normalization Tools (ANTS)”. en. In: ().
- Behzadi, Y. et al. (Aug. 2007). “A component based noise correction method (CompCor) for BOLD and perfusion based fMRI”. en. In: *NeuroImage* 37.1, pp. 90–101. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2007.04.042.
- Castegnetti, G., M. Zurita, and B. DeMartino (2021). “How usefulness shapes neural representations during goal-directed behavior”. en. In: *Science Advances* 7.15. DOI: 10.1126/sciadv.abd5363.
- Draper, N. and H. Smith (1998). *Applied Regression Analysis*. 3rd ed. Wiley.
- Farashahi, S. et al. (Sept. 2019). “Flexible combination of reward information across primates”. en. In: *Nature Human Behaviour* 3.11, pp. 1215–1224. ISSN: 2397-3374. DOI: 10.1038/s41562-019-0714-3.
- Gorgolewski, K. et al. (2011). “Nipype: A Flexible, Lightweight and Extensible Neuroimaging Data Processing Framework in Python”. en. In: *Frontiers in Neuroinformatics* 5. ISSN: 1662-5196. DOI: 10.3389/fninf.2011.00013.
- Iigaya, K. et al. (Feb. 2020). *Aesthetic preference for art emerges from a weighted integration over hierarchically structured visual features in the brain*. en. preprint. Neuroscience. DOI: 10.1101/2020.02.09.940353.

- Iigaya, K. et al. (June 2021). “Aesthetic preference for art can be predicted from a mixture of low- and high-level visual features”. en. In: *Nature Human Behaviour* 5.6, pp. 743–755. ISSN: 2397-3374. DOI: 10.1038/s41562-021-01124-6.
- Iigaya, K. et al. (Jan. 2023). “Neural mechanisms underlying the hierarchical construction of perceived aesthetic value”. en. In: *Nature Communications* 14.1, p. 127. ISSN: 2041-1723. DOI: 10.1038/s41467-022-35654-y.
- Jenkinson, M. et al. (Oct. 2002). “Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images”. en. In: *NeuroImage* 17.2, pp. 825–841. ISSN: 10538119. DOI: 10.1006/nimg.2002.1132.
- Jenkinson, M. et al. (Aug. 2012). “FSL”. en. In: *NeuroImage* 62.2, pp. 782–790. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2011.09.015.
- Kahnt, T. et al. (May 2011). “Decoding different roles for vmPFC and dlPFC in multi-attribute decision making”. en. In: *NeuroImage* 56.2, pp. 709–715. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2010.05.058.
- Li, C. and T. Chen (Apr. 2009). “Aesthetic Visual Quality Assessment of Paintings”. en. In: *IEEE Journal of Selected Topics in Signal Processing* 3.2, pp. 236–252. ISSN: 1932-4553. DOI: 10.1109/JSTSP.2009.2015077.
- Lim, S.-L., J. P. O’Doherty, and A. Rangel (May 2013). “Stimulus Value Signals in Ventromedial PFC Reflect the Integration of Attribute Value Signals Computed in Fusiform Gyrus and Posterior Superior Temporal Gyrus”. en. In: *The Journal of Neuroscience* 33.20, pp. 8729–8741. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.4809-12.2013.
- Liu, Z. et al. (June 2016). “DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations”. en. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Las Vegas, NV, USA: IEEE, pp. 1096–1104. ISBN: 978-1-4673-8851-1. DOI: 10.1109/CVPR.2016.124.
- Mazziotta, J. et al. (Sept. 2001). “A Four-Dimensional Probabilistic Atlas of the Human Brain”. en. In: *Journal of the American Medical Informatics Association* 8.5, pp. 401–430. ISSN: 1067-5027, 1527-974X. DOI: 10.1136/jamia.2001.0080401.
- McNamee, D., A. Rangel, and J. P. O’Doherty (2013). “Category-dependent and category-independent goal-value codes in human ventromedial prefrontal cortex”. en. In: *Nature Neuroscience* 16.4, pp. 479–485. DOI: 10.1038/nn.3337.
- O’Doherty, J. P., U. Rutishauser, and K. Iigaya (Oct. 2021). “The hierarchical construction of value”. en. In: *Current Opinion in Behavioral Sciences* 41, pp. 71–77. ISSN: 23521546. DOI: 10.1016/j.cobeha.2021.03.027.
- Platt, M. L. and P. W. Glimcher (July 1999). “Neural correlates of decision variables in parietal cortex”. en. In: *Nature* 400.6741, pp. 233–238. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/22268.

- Suzuki, S., L. Cross, and J. P. O'Doherty (2017). "Elucidating the underlying components of food valuation in the human orbitofrontal cortex". en. In: *Nature Neuroscience* 20.12, pp. 1780–1786. DOI: 10.1038/s41593-017-0008-x.
- Tustison, N. J. et al. (June 2010). "N4ITK: Improved N3 Bias Correction". en. In: *IEEE Transactions on Medical Imaging* 29.6, pp. 1310–1320. ISSN: 0278-0062, 1558-254X. DOI: 10.1109/TMI.2010.2046908.
- Tzourio-Mazoyer, N. et al. (Feb. 2020). *Automated Anatomical Labeling (AAL1) atlas*. DOI: 10.25493/WJHA-2NY.
- Wang, L. et al. (Oct. 2015). "Probabilistic Maps of Visual Topography in Human Cortex". en. In: *Cerebral Cortex* 25.10, pp. 3911–3931. ISSN: 1047-3211, 1460-2199. DOI: 10.1093/cercor/bhu277.
- Winterfeldt, D. von and G. Fisher (1975). "Multi-attribute utility theory: models and assessment procedures". In: *Wendt, D., Vlek, C.A.J. (Eds.), Utility, probability, and human decision making*. Dordrecht, The Netherlands: Reidel, pp. 47–86.
- Zysset, S. et al. (July 2006). "The neural implementation of multi-attribute decision making: A parametric fMRI study with human subjects". en. In: *NeuroImage* 31.3, pp. 1380–1388. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2006.01.017.

Chapter 3

CHOICE-DRIVEN REMAPPING OF ACTION- AND STIMULUS-ANCHORED VALUE IN HUMAN SINGLE NEURONS

Fransen, A. et al. (2026). “Choice-driven remapping of action- and stimulus-anchored value in human single neurons”.

Abstract

Whether stimulus- and action-based choices rely on a shared neural architecture remains intensely debated. While lesion and fMRI studies show conflicting regional dissociation and overlap, the underlying single neuron substrates remain unknown. We recorded single neuron activity from the human ventromedial prefrontal cortex (vmPFC), anterior cingulate cortex (ACC), pre-supplementary motor area (preSMA), and amygdala during distinct action- and stimulus-based reward tasks. Prior to decision-making, neurons across all regions broadly encode value for both choice domains. However, choice commitment triggers domain segregation across the prefrontal cortex. While ACC and vmPFC neurons flexibly track both chosen actions and stimuli, they exhibit a double dissociation for chosen valuation: ACC neurons selectively maintain action-based values, whereas vmPFC neurons selectively maintain stimulus-based values. Conversely, preSMA neurons invariant to task domain track chosen action identity and value. These findings reveal how prefrontal networks transition from domain-general evaluation to domain-segregated choice execution.

Introduction

Value-based choices rely on the computation of expected values for potential candidates, whether those candidates are visual stimuli or motor actions (Rangel, Camerer, and Montague, 2008). A fundamental debate in the study of value-based decision-making is whether these value computations are segregated by the choice-domain (i.e. stimulus versus action-based choices) or instead rely on a shared neural substrate (see ‘common currency’ literature e.g., Levy and Glimcher, 2012). Historically, specific neural circuits have been theoretically partitioned based on these

choice-domains. For instance, evidence from lesion studies in both monkeys and humans have demonstrated a macro-level functional dissociation, revealing specific action or stimulus selection deficits following damage to the anterior cingulate cortex (ACC) and orbitofrontal cortex (OFC), respectively (see Rudebeck et al., 2008; Camille, Tsuchida, and Fellows, 2011). While these findings suggest a broad regional divide in necessary neural infrastructure, it remains unclear if such distinct separation appears at the cellular level. Additionally, these studies leave open the question of how these action- and stimulus-based value signals are dynamically computed over the stages of the decision process.

To characterize the encoding of stimulus and action values, previous studies have employed various recording modalities, though critical gaps remain. Single-neuron studies in non-human primates have been instrumental in mapping valuation circuitry, frequently identifying regions such as the striatum, premotor cortex, and anterior cingulate cortex as key nodes for action-based decision-making (e.g., Samejima et al., 2005; Scangos and Stuphorn, 2010; Hayden and Platt, 2010; Seo and Lee, 2007; Kennerley et al., 2006). However, a pervasive methodological limitation in much of this literature is the inherent conflation of motor actions with spatial targets. In some experimental paradigms, animals are required to indicate their choices by executing a movement (such as a saccade or a reach) to a specific spatial coordinate (e.g., the left or right side of a computer screen). As such, the motor actions performed only varied with respect to the spatial location they were directed towards. Consequently, neural signals that have been interpreted as encoding ‘action value’ might simultaneously reflect the value of the spatial target, the spatial attention directed toward it, or the directional movement vector itself (see for discussion; Padoa-Schioppa, 2011). This confound makes it exceedingly difficult to isolate ‘pure’ motor action value. Here we therefore aimed to define an action space that spanned unique physical effectors (e.g., selecting a left-hand versus a right-foot movement) and was as dissociated from the visual stimuli/attention as possible, in an effort to more directly pair value to distinct motor actions.

In parallel to investigations of action value, a robust line of single-neuron research has sought to characterize the encoding of stimulus value, primarily focusing on the orbitofrontal cortex (OFC) and ventromedial prefrontal cortex (vmPFC). Studies in non-human primates have demonstrated that neurons in these regions encode the economic value of offered and chosen visual stimuli, often doing so independently of the sensorimotor contingencies required to obtain them (Padoa-Schioppa and

Assad, 2006; Strait, Blanchard, and Hayden, 2014). This has led to the prominent hypothesis that the OFC/vmPFC computes a purely abstract, domain-general “common currency” for economic choice, evaluating choice options independent to any motor planning (Padoa-Schioppa and Assad, 2006; Levy and Glimcher, 2012). However, testing whether these stimulus value signals are truly distinct from, or integrated with, action value computations requires directly recording from both domains simultaneously within the same neural populations.

In the human brain, functional magnetic resonance imaging (fMRI) has been deployed to investigate whether these valuation domains are represented by distinct networks. Such studies have yielded evidence for regional overlap, implicating the ventromedial prefrontal cortex (vmPFC) in encoding value-related signals supporting choices over stimuli and actions (Gläscher, Hampton, and O’Doherty, 2009; Wunderlich, Rangel, and O’Doherty, 2010). Additional studies have set out to map the process of stimulus to action mapping in stimulus-based decisions and found a network involving vmPFC, dorsomedial PFC and motor cortex to support such stimulus-to-action decisions (Hare et al., 2011). However, while fMRI has successfully mapped relevant valuation circuitry, its reliance on blood-oxygen-level-dependent (BOLD) signals inherently limits both spatial- and temporal resolution. A single functional voxel exhibiting overlapping representations of stimulus and action value could reflect a genuinely integrated, mixed-selective neural population, but it could equally arise from entirely segregated (but anatomically interleaved) populations of domain-specific neurons. Similarly, the poor temporal resolution of the hemodynamic response function makes differentiating between distinct facets of the decision process difficult. fMRI activity related to value signals within a trial could reflect neuronal firing pre-choice, post-choice or a mixture (or compounding) of both. Temporally segregating the neural response to each of these decision stages is necessary to gain a deeper mechanistic understanding of these stimulus- and/or action-based value computations. Direct electrophysiological recordings in humans are uniquely positioned to bridge these gaps and resolve the question of true single-neuron separability.

Previous research utilizing human single-unit recordings has successfully identified neurons that represent the value of available choice candidates, regardless of whether they are chosen (here also referred to as “decision value”), in pre-supplementary motor area (preSMA) during value-based choice (Aquino et al., 2023). Yet, because the behavioral task employed in this study was not designed to decouple the visual

stimuli used from the physical motor effectors required to select them, it remains unclear whether these human value signals are stimulus- or action-based in their format, or part of a domain-general computation.

To address this open question, we set out to differentiate the encoding of stimulus value and action value at the single-neuron level in the human brain. We designed an experiment to tackle the two limitations discussed above. We included a stable set of stimuli and actions, allowing us to track the expected value attached to each of these choice options throughout the experiment. Randomization of stimulus-action pairing additionally orthogonalized stimulus- from action values. Lastly, we employed clearly distinct motor actions and tied reward probabilities to these motor responses in the action portion of the experiment. This set of hand- and foot-operated buttons helped isolate action value as it is tied to different motor effectors. We recorded single-unit activity from 902 neurons across the vmPFC, ACC, preSMA, and amygdala to investigate the neural representation of stimulus and action value. We leverage the temporal resolution of these recordings to trace the computation of value as the decision process evolves from the pre-decision representation of choice alternatives to choice commitment and selection.

Results

Our primary objective was to characterize and compare the single-neuron computations underlying value-based decision-making when choices are framed over motor actions versus visual stimuli. To address this, we recorded 230 vmPFC, 299 preSMA, 157 dACC and 216 amygdala single neurons in 20 sessions across 16 patients implanted with hybrid macro- and micro-electrodes for epilepsy monitoring. Participants performed an experiment specifically designed to differentiate the encoding of stimulus- and action-based decisions. As part of this experiment, participants performed two structurally parallel reinforcement learning tasks in counterbalanced order (see Methods for details). In the “Action Task,” reward probabilities were tied to specific motor responses. Given our questions’ emphasis on distinct action value encoding, we chose to enrich the action-space beyond the typical 2 hand-button press used in previous choice experiments. As such, our motor actions included a left-hand button press, a right-hand button press, a left-foot pedal press, and a right-foot pedal press. In contrast, in the “Stimulus Task,” reward probabilities were associated with visual stimuli. The Stimulus Task took advantage of the same four motor actions for choice selection. Importantly, we ensured the value of visual stimuli was independent of the physical action required to select

them through trial-wise pseudo-random stimulus-action pairing (see Figure 3.1A, B and Methods). Participants completed over five blocks of 36 choice trials on each task on average (see Methods).

Behavioral Results




Figure 3.1: Task and behavioral results. See below for continued caption.

We started our investigation of participants' behavior across these two tasks by determining if participants were able to learn to select more rewarding choice candidates (i.e., actions in the Action Task and stimuli in the Stimulus Task, note that *options* is a more typical word used for decision candidates but we refrain from doing so as to avoid confusion with the notion of an *option* from reinforcement-learning (Sutton, Precup, and Singh, 1999)). Participants could learn to select candidates with a higher probability of reward through experiencing the win or no-win outcomes resulting from the selection. As such, we would expect to see an increase in the selection of the most rewarding choice candidate over the course of a block. Indeed participants demonstrated significant learning over the course of a 36-trial block in both tasks. In the Action Task, the probability of selecting the candidate with the highest expected value increased significantly over time (Wilcoxon test over

Figure 3.1: (continued) **A)** Trial design of action task. Fixation cross is followed by the “Go” signal to indicate one of three instructed actions can be chosen and executed. Actions include; left hand press, left foot press, right hand press and right foot press. The check mark confirms action selection independent of the specific action chosen. The trial ends with either a reward (yellow coin) or no reward (red coin) in accordance with the reward probability attached to the chosen action. Participants complete 36 such trials over the course of a block. See Methods for more details on task design. **B)** The stimulus trial design mirrors the action task with the exception that stimuli are displayed in the “Go” window. There are four total stimuli; tree, house, dog and camera. Participants make a choice between the 2/3 pseudo-randomly drawn stimuli displayed on the screen. The stimulus location determines the action required to make the selection. Stimuli displayed on the top of the screen can be selected using the hand buttons (i.e. left hand to select the tree in example) while stimuli on the bottom of the screen can be selected using their left or right foot (i.e. left foot to select camera in example). In this task the probability of winning (i.e. reward) is dependent on the stimulus selected (and independent of the action executed). **C)** Probability of selecting the candidate with the highest expected value. Solid line indicates average performance across participants and blocks in the action task (with the shaded band representing the 95% confidence interval). Individual session’s performance is shown in the thinner background lines. **D)** Average performance similarly increases over the course of blocks in the stimulus task. Please note chance here is $P=1/2$ since participants select one out of two available candidates in each trial. **E)** Probability of selecting the correct candidate (i.e. highest expected value) in the stimulus task increases with the difference in win probability between the available candidates. In other words, participants perform better on trials that are easier. Whiskers indicate standard error around the mean (indicated with the central dot). These effects are significant (Mann-Whitney U test: ** : $p < 0.01$ and *** : $p < 0.001$) **F)** Participants are more likely to repeat an choice that was rewarded on the previous trial (t-1) across both tasks. We see this effect quickly decays further back into the past where at two trials (t-2) and three trials back (t-3) this effect is no longer significant. Whiskers indicate standard error around the mean. Estimates from a logistic regression model (win vs no win, see Methods) transformed to log odds probabilities. **G)** The probability of staying after a win on the previous trial (darker colored bars) or a no-win on the previous trial (lighter colored bars) in the action task. Bars indicate the mean across sessions with the whiskers representing the standard error. Leftmost bars display real participant stay tendencies. The middle bars are simulated using the RL model (see Methods), while rightmost bars indicate choice kernel model simulations (see Methods). **H)** Model simulated stay behavior in the stimulus task. Plot follows the same conventions as G.

the last 10 trials: $p < 0.001$; Figure 3.1C). At the start of a block, when the reward probabilities associated with each action are unknown, participants' performance was at chance level. However, this performance improved over the course of the block indicating participants learned to associate the available actions with their reward probabilities (see Methods). Similarly, in the Stimulus Task, participants selected the more rewarding stimulus more than expected by chance at the end of a block (Wilcoxon test over the last 10 trials: $p < 0.001$; Figure 3.1D).

Furthermore, choice accuracy in the Stimulus Task was sensitive to the relative value of the available stimuli. The probability of selecting the most rewarding choice candidate available increased monotonically with the difference in win probability ($\Delta p(\text{win})$) between the two presented stimuli (Figure 3.1E). This indicates that participants were more effective at identifying the optimal choice when the expected value difference was larger (i.e., when the choice was easier).

Temporal Dynamics of Reinforcement

To investigate how recent reward history shaped choice behavior, we performed a session-level logistic regression analysis to determine the probability participants repeat rewarded choices (Figure 3.1F and see Methods). We hypothesized that if participants were utilizing reinforcement to guide their behavior, they would be more likely to repeat a choice that resulted in a win. Our results confirmed this predictions across both tasks. We observe participants are approximately 15 to 30% more likely to repeat the choice for which they received a reward on the immediately preceding trial (for action and stimulus tasks respectively; $p = 0.009$ for action and $p < 0.001$ for stimulus task). Notably, this effect is highly similar across both tasks, suggesting that the integration of rewards follows a similar temporal scale regardless of whether the reward is associated with a motor action or a visual stimulus.

Computational Modeling of Choice Behavior

Finally, we sought to formalize participants' choice behavior by fitting a computational model. We fit a reinforcement learning (RL4) model and a choice kernel (CK Only) model to the participants' choice data (see Methods for model details). The choice kernel model captures participants tendency to repeat previously chosen actions (see Figure 3.1G, H). However, this model is unable to replicate the observed difference in selection repetition after a win versus no-win (see Supplemental Figure B.2 for quantitative model comparison results). We therefore utilized

a four-parameter reinforcement learning model that tracks both expected values (Q -values) and choice history (choice kernel), allowing us to capture both value-based updates and choice perseveration (see Figure 3.1G,H and Supplemental Table B.2 and Figure B.2).

To connect our behavioral findings with the hypothesized single-neuron computations, we utilized the selected reinforcement learning model to extract latent, trial-by-trial decision variables for each session. While our behavioral analysis demonstrated that this model successfully captures the dynamics of choice, its primary utility for our neural analysis lies in generating precise, session-specific estimates of expected values. This process allows us to isolate these internal computational signals as they evolve over the course of a block through experience with the outcomes (wins or no-wins). We can then use these latent variables as direct trial-by-trial regressors to elucidate the associated neural computations. As a result, we can map exactly where and when individual neurons across our regions of interest dynamically track these specific decision signals during action- and stimulus-based choices.

Neural representations of Action and Stimulus Decision Value

Having established that participants utilize a reinforcement-based strategy in both tasks, we turn to investigating the underlying neural computations. Our primary goal was to determine whether single neurons in our areas of interest (vmPFC, preSMA, ACC and amygdala) represent value across choice domains (i.e., coding for the value of both stimuli and actions) or if these representations are domain-specific. There are various points in a trial where we might hypothesize different types of value signals are computed. We start by analyzing the representation of the value of choice options that are available for selection. These value signals will serve as an input to the comparison/decision process. Here we will refer to these values that proceed choice as ‘decision value’ (note that these types of value signals are sometimes referred to as ‘offer value’ (e.g., Shi, Ballesta, and Padoa-Schioppa, 2022; Fine et al., 2023)).

Based on previous literature (Camille, Tsuchida, and Fellows, 2011; Kennerley and Walton, 2011), we might hypothesize domain-specific value computations, with regionally clustered neurons in the prefrontal cortex specializing in either action- or stimulus-based decision value representation. To test this hypothesis, we identified neurons whose firing rates significantly correlated with the decision value of actions or stimuli available for selection, as derived from our RL model, after stimulus onset.

In the action task, we surprisingly found a significant population of neurons across vmPFC ($n = 46$, 20%, $p < 0.001$), preSMA ($n = 46$, 15.4%, $p < 0.001$), ACC ($n = 35$, 22.3%, $p < 0.001$) and amygdala ($n = 30$, 13.9%, $p < 0.001$) that tracked the value of available actions (Figure 3.2B; proportions evaluated via binomial tests on permuted neuron significance, see Methods). A temporally resolved encoding analysis (see Methods) additionally highlighted that neurons across these areas encode the expected value (Q -value) of available actions early, even before the ‘Go’-screen prompting choice initiation.

Similarly, in the stimulus task, a significant population of neurons across the regions of interest tracked the value of the available stimuli (vmPFC: $n = 32$, 14%, $p < 0.001$; preSMA: $n = 41$, 13.7%, $p < 0.001$; ACC: $n = 39$, 24.8%, $p < 0.001$; amygdala: $n = 27$, 12.5%, $p < 0.001$)(see Figure 3.2E). In elucidating the temporal dynamics of these signals, we observed a clear stimulus-onset-induced ramping of encoding (Figure 3.2F). In the vmPFC, specifically, we observed a steady increase in the number of neurons that code for stimulus-based decision value as time progresses post stimulus onset.

Lastly, we evaluated the representation of available action values during the stimulus task. Although reward probabilities in this task were explicitly tied to the visual stimuli, the brain must ultimately translate these expected values into specific motor plans to execute a decision. To test whether the brain represents pre-decision action values in a stimulus-based task, we mapped the expected value of each presented stimulus to the specific motor effector required to select it on that given trial (e.g., if a stimulus yields a Q -value of 0.8 and its current paired action is the left-hand button press, the available left-hand action value is set to 0.8). This operationalization, which is standard for computing action values in spatial stimulus-based bandit tasks (e.g., Aquino et al. (2023)), enables us to test whether the brain automatically evaluates potential motor actions even when the decision is primarily stimulus-driven. This analysis revealed significant action decision value encoding solely in the preSMA (preSMA: $n = 29$, 9.7%, $p < 0.001$; vmPFC: $n = 13$, 5.7%, $p = 0.37$; ACC: $n = 13$, 8.3%, $p = 0.052$; amygdala: $n = 12$, 5.6%, $p = 0.40$ - see Figure 3.2H).

Together, these results reveal that, contrary to our hypothesis, there is substantial overlap in the regions that encode for stimulus decision value in the stimulus task and action decision value in the action task. However, this does not yet determine whether the same neurons underpin these computations across domains. We tested

this domain generality at the single-neuron level by computing the stability of model statistics of single neurons across tasks (Figure 3.2J; see Methods). Here we observe a significant correlation between action value and stimulus value coding strength across tasks in all ROIs (vmPFC: $\rho = 0.26$, $p < 0.001$, preSMA: $\rho = 0.27$, $p < 0.001$, ACC: $\rho = 0.36$, $p < 0.001$, amygdala: $\rho = 0.20$, $p = 0.003$, see Figure 3.2J in purple). These correlations indicate that neurons strongly encoding the value of actions in the action task were significantly more likely to also encode the value of stimuli in the stimulus task. Conversely, neurons that contributed to action value coding in the stimulus task were not more likely to code for action value in the action task. This pattern appears in all areas except preSMA ($\rho = 0.26$, $p < 0.001$, see 3.2J in orange) and is particularly striking in the ACC (ACC: $\rho = 0.10$, $p = 0.22$, vmPFC: $\rho = 0.14$, $p = 0.04$, amygdala: $\rho = 0.12$, $p = 0.08$). We thus conclude that the neural representation of decision values is supported by a more overlapping set of neurons across choice-domains than the representation of action values per se. In other words, neurons are more likely to encode the decision value in whichever domain is relevant to the current choice than they are to maintain a domain-specific encoding regardless of task relevance (i.e., always code for action value).

In the ACC we can clearly observe this dichotomy when plotting the individual neurons' model statistics (colored by significance; Figure 3.2K, L). Among ACC neurons, we can observe four distinct clusters: a subpopulation coding for both stimulus and action value (purple), alongside strictly action-selective (red), strictly stimulus-selective (blue), and non-value-coding (gray) populations. This is in stark contrast to the cross-task action value comparison (Figure 3.2L). While we observe neurons that code for action in each task there is a complete lack of overlap in these populations.

When further investigating the coding scheme employed in the subset of neurons that code for both stimulus values and action values (e.g., those in purple in Figure 3.2K) we find a remarkably consistent flip in coding directionality in ACC neurons (see Figure B.3). Nearly all neurons in the ACC found to encode stimulus value do so using a positive coding scheme (where the firing rate increases with value). In contrast, these same neurons in the ACC largely negatively code for action value in the action task. While the number of neurons in this analysis is small this result might point to the ACC using anti-correlated subspaces to encode decision values across these different choice domains. Together, these results demonstrate a common neural substrate for domain-relevant decision values (while relying on




Figure 3.2: Decision value coding. See below for continued caption.

different subspaces within the ACC particularly).

Representations of Chosen Value

Having detailed the neural encoding of decision values, we next investigated whether these regions also reflect the value of the final decision (i.e., the chosen value). More specifically, we sought to determine if the expected value of the selected choice candidate (chosen Q -value) is represented across our ROIs and whether this encoding differs depending on the task domain. Again, previous research might lead us to hypothesize a regional specialization in either stimulus- or action-value. Alternatively, there could be a continued domain-flexibility as the choice progresses from decision value to chosen value.

To differentiate between these possible hypotheses, we evaluated chosen value encoding during the Action Task. In the period leading up to the response, we found that chosen action values were selectively represented in regions heavily implicated in motor planning and action valuation. Specifically, a significant proportion of neurons in the preSMA ($n = 31$, 10.4%, $p < 0.001$) and the ACC ($n = 17$, 10.8%, $p = 0.002$) tracked the value of the action ultimately executed by the participant

Figure 3.2: **A)** Action value representation in an example Anterior Cingulate Cortex (ACC) neuron. Top: Raster plot of individual trials aligned to stimulus onset (dashed vertical line at 0 ms). Trials are sorted and color-coded by action value tertile (blue: high, orange: mid, green: low). Bottom: Peristimulus time histogram (PSTH) showing the average firing rate for each tertile. Shaded regions represent the standard error. **B)** Proportion of significant neurons encoding action value during the action task across four recorded brain regions: ventromedial prefrontal cortex (vmPFC), pre-supplementary motor area (preSMA), ACC, and Amygdala. The generalized linear model used for estimation is displayed above, highlighting the four available action options. Dark gray bars represent the percentage of significant units, while white bars represent the remaining non-significant population. Dashed lines indicate the chance level. **C)** Time-resolved proportion of significant units encoding action value relative to stimulus onset, calculated using a sliding window approach (see Methods). Solid lines denote the proportion for each brain region at $p < 0.05$ uncorrected. **D)** Stimulus value representation in the exact same ACC neuron shown in panel A, demonstrating mixed selectivity. The raster plot and PSTH follow the same conventions as in panel A, but trials are sorted by stimulus value tertiles during the stimulus task. **E)** Proportion of significant neurons encoding stimulus value across the four brain regions. The corresponding stimulus value model is shown above, highlighting the distinct visual stimuli. **F)** Time-resolved proportion of significant units encoding stimulus value relative to stimulus onset, following the same conventions as panel C. **G)** Action value representation in an example preSMA neuron, following the same conventions as panel A. **H)** Proportion of significant neurons encoding action value during the stimulus task, following the conventions of panel B. **I)** Time-resolved proportion of significant units encoding action value during the stimulus task relative to stimulus onset. **J)** Correlation in model statistics across tasks and value domains. Bar plot displays the Spearman rank correlation (ρ) of likelihood ratio statistics for each brain region. Dark orange bars compare action value coding in the stimulus task versus the action task. Dark purple bars compare action value coding versus stimulus value coding across the respective tasks. Filled bars indicate the correlation is significant at $p < 0.01$. **K)** Scatter plot comparing the model statistic for stimulus value in the stimulus task (y-axis) against action value in the action task (x-axis) for individual units in the ACC. Units are colored based on significance across both tasks. We then label neurons red when they are significant only in the action task, blue if they solely contribute significantly in the stimulus task. Units labeled in purple if they contribute significantly across both tasks/variables. Dashed line added to the 45-degree line for visual aid in comparing units to 'equal contribution'. **L)** Scatter plot comparing the model statistic for action value in the stimulus task (y-axis) against action value in the action task (x-axis) in the ACC. Conventions follow panel K. No units are found to be significant across both tasks.

(see Figure 3.3A, B). In contrast, we did not find significant encoding of chosen action value at our predefined threshold ($p < 0.01$) in the vmPFC ($n = 17$, 7.4%, $p = 0.071$) or the amygdala ($n = 18$, 8.3%, $p = 0.024$).

The same analysis run on the Stimulus Task revealed a dissociation in chosen value coding. Rather than utilizing a static, domain-general network across both decision domains, we observed a partial shift in the regions recruited to represent chosen stimulus value. In the stimulus domain, the vmPFC reliably encoded the value of the chosen option ($n = 24$, 10.4%, $p < 0.001$), alongside the preSMA ($n = 30$, 10.0%, $p < 0.001$). Conversely, the ACC ($n = 5$, 3.2%, $p = 0.897$), which had prominently tracked chosen value in the action task, failed to exhibit significant coding for chosen stimulus value, as did the amygdala ($n = 14$, 6.5%, $p = 0.195$) (see Figure 3.3D, E). To formally test this double dissociation between the vmPFC and ACC, we utilized a logistic regression to evaluate the likelihood of a neuron significantly encoding chosen value as a function of region and task (see Methods). This analysis revealed a highly significant region $\times$ task interaction ($\beta = 1.68$, $z = 2.72$, $p = 0.006$), confirming a double dissociation.

Together, these results highlight a more specialized, domain-specific set of neurons track the value of the chosen option in stimulus- versus action-framed choices. Unlike in the encoding of pre-choice decision values, at the region-level we already observed a domain-specific specialization. A significant proportion of neurons in the ACC track the chosen value exclusively when decisions are made over motor actions. Conversely, the vmPFC demonstrates the opposite specialization, being selectively recruited to track chosen value only during stimulus-based decisions.

Bridging this domain-specific dichotomy is the preSMA, which encodes chosen value regardless of the choice domain. To further characterize the encoding of chosen value at the level of individual neurons, we tested whether a neuron more strongly encoding chosen value in the stimulus-domain would equally encode chosen value when the choice was action-driven (correlation of model statistics - see Methods). This analysis revealed that the encoding of chosen value in stimulus- and action-domains relies on separated sub-populations (vmPFC: $\rho = 0.03$, $p = 0.61$, preSMA: $\rho = 0.02$, $p = 0.67$, ACC: $\rho = 0.13$, $p = 0.11$, amygdala: $\rho = 0.13$, $p = 0.06$, see Figure 3.3G). This separation of value coding appears even in the preSMA, where we did observe both stimulus and action chosen value at the region level (see Figure 3.3A, D, G, H). From this analysis we conclude that different subsets of neurons within the preSMA facilitate the representation of domain-specific

chosen value codes. Maintaining such information would be relevant for mapping to motor plans and credit-assignment (see Discussion). Lastly, despite being driven by separate sub-populations in the preSMA, an analysis of the latency of chosen value coding revealed no latency difference between the encoding of chosen values in the stimulus and action task (see Figure 3.3I). As such, while these chosen values are encoded in regionally segregated neurons, the similar temporal evolution of this computation might hint at its shared function across choice-domains.




Figure 3.3: Chosen value encoding. See below for continued caption.

Representations of Chosen Identity

Having established that distinct regional populations specialize in tracking the value of the chosen candidate depending on task domain, we next investigated whether these regions also represent the identity of the chosen candidate. Note that an isolated value signal is insufficient for goal-directed behavior; the brain must explicitly encode the identity of the chosen candidate to successfully orchestrate its selection and to properly assign credit during subsequent learning updates. We would therefore hypothesize that the regions found to represent chosen value would also represent the chosen identity.

We first evaluated the chosen action identity (the specific button press) encoding across our regions of interest during the action task. Remarkably, we observed robust representation of the selected action across prefrontal regions. Specifically, a highly significant proportion of neurons in the preSMA ($n = 56$, 18.7%, $p < 0.001$), the ACC ($n = 41$, 26.1%, $p < 0.001$), but also the vmPFC ($n = 24$, 10.4%,

Figure 3.3: (continued) **A)** Proportion of significant neurons encoding chosen value during the action task across four recorded brain regions: ventromedial prefrontal cortex (vmPFC), pre-supplementary motor area (preSMA), anterior cingulate cortex (ACC), and Amygdala. Dark gray bars indicate regions where the percentage of significant units exceeds chance (at $p < 0.01$), while white bars represent non-significant populations. Dashed horizontal lines indicate the specific chance level for each region (see Methods). **B)** Time-resolved proportion of significant units encoding chosen value relative to response execution (dashed vertical line at 0 ms) in the action task. Solid lines denote the proportion for each brain region computed using a sliding window approach (see Methods). **C)** Chosen value representation in an example ACC neuron during the action task. Top: Raster plot of individual trials aligned to response onset. Trials are sorted and color-coded by chosen value tertile (blue: high, orange: mid, green: low). Bottom: Peristimulus time histogram showing the average firing rate for each tertile. Shaded regions represent the standard error. **D)** Proportion of significant neurons encoding chosen value during the stimulus task. Conventions are identical to panel A. **E)** Time-resolved proportion of significant units encoding chosen value relative to the response during the stimulus task. Conventions follow panel B. **F)** Chosen value representation in an example preSMA neuron during the stimulus task. Conventions follow panel C. **G)** Spearman correlation for chosen value across tasks. The bar plot displays the correlation (ρ) of chosen value model statistics between the stimulus task and the action task for each brain region. None of the correlations are significant. **H)** Layout as in Figure 3.2K. Scatter plot comparing the model statistic for chosen value in the stimulus task (y-axis) against chosen value in the action task (x-axis) for individual units in the preSMA. Units are colored based on significance across both tasks (red for significance in action task; blue for the stimulus task; purple for both; gray for neither). Dashed line added to the 45-degree line for visual aid in comparing units to ‘equal contribution’. **I)** Latency analysis of the onset of firing rate change in significant units. Significance determined based on comparison with the expected Poisson baseline firing rate (see Methods). Analysis aligned to response (indicated by time 0).

$p < 0.001$) tracked the identity of the chosen action (see Figure 3.4A). A temporally resolved analysis further revealed that the proportion of engaged neurons in both the ACC and preSMA exhibits a ramping trajectory that progressively builds as the moment of choice execution approaches (see Figure 3.4B). Such ramping would be consistent with the hypothesis that these signals might contribute to the motor planning leading up to selection. The amygdala, however, did not exhibit significant encoding of chosen action identity in the action task ($n = 15$, 6.9%, $p = 0.126$).

We next examined chosen action identity during the stimulus task. In this task, the motor output is merely a vehicle to obtain a desired visual stimulus. Crucially, under these conditions, the comprehensive representation of the to-be-executed action collapsed in most regions. At our predefined threshold ($p < 0.01$), the preSMA was the only region that maintained a significant representation of the motor response ($n = 44$, 14.7%, $p < 0.001$). Strikingly, the ACC ($n = 10$, 6.4%, $p = 0.261$) and the vmPFC ($n = 18$, 7.8%, $p = 0.042$) ceased to reliably track the identity of the action, as did the amygdala ($n = 17$, 7.9%, $p = 0.045$). This result replicates findings from previous research (e.g., Aquino et al., 2023) that found the preSMA to encode chosen action identity in a stimulus-based choice task. Additionally, the clear ramping pattern observed in preSMA (see Figure 3.4F), further reinforces the interpretation that this area might translate value information to motor commands to relay to neighboring motor cortex. When comparing this encoding of chosen action identity across tasks, we find a subset of preSMA neurons significantly codes for chosen action identity across both tasks (see Figure 3.4G, I, preSMA: $\rho = 0.22$, $p < 0.001$, vmPFC: $\rho = 0.20$, $p = 0.003$, ACC: $\rho = 0.24$, $p = 0.002$, amygdala: $\rho = 0.24$, $p < 0.001$). Additionally, a generalization analysis in this area confirms this motor action representation is indeed task invariant in the preSMA (see Figure 3.1H and Methods).

In contrast to this specialized preSMA encoding of chosen action in the stimulus task, the encoding of chosen stimulus identity appeared more widespread. During the response period of the stimulus task, we found widespread, significant encoding of the chosen stimulus identity across all of our recorded regions of interest. The vmPFC ($n = 31$, 13.5%, $p < 0.001$), ACC ($n = 21$, 13.4%, $p < 0.001$), amygdala ($n = 23$, 10.6%, $p < 0.001$), and preSMA ($n = 25$, 8.4%, $p = 0.009$) all possessed significantly large neural populations dedicated to signaling which visual stimulus the participant had committed to (see Figure 3.4C). While we do also observe significant encoding of chosen stimulus in the preSMA, this code depends on a

distinct sub-populations of neurons than that for chosen action in the action task (see Figure 3.4J & G). An analysis of the onset of these chosen identity codes in the stimulus task shows the chosen stimulus encoding proceeds the chosen action encoding in the preSMA (see Figure 3.4K). This finding supports the idea that the preSMA might translate stimulus information into action selection.

Lastly, while there is a significant correlation between single neuron encoding strength across choice domains in the ACC ($\rho = 0.29$, $p < 0.001$), the encoding of chosen stimulus identity and action identity appear to rely on independent sub-groups of neurons within vmPFC, preSMA and amygdala (vmPFC: $\rho = 0.14$, $p = 0.03$, preSMA: $\rho = 0.07$, $p = 0.22$, amygdala: $\rho = 0.01$, $p = 0.84$, see Figure 3.4G in purple). We can thus conclude that while there is regional overlap in the encoding of chosen identity across task-domains, this encoding is distinctly more segregated than we observed for decision value encoding (see Figure 3.2J).

The ubiquity of these chosen identity signals might hint at their function in maintaining the code necessary to ensure the new value assigned to the chosen candidate after outcome gets appropriately assigned. However, our limited set of recording locations prevents characterizing these processes in, for example, the basal ganglia. The temporal resolution of these recordings does however offer us a potential for a post-hoc test to falsify the role of these chosen identity signals in such an outcome-driven update tagging. If these representations are indeed used in value-assignment, we might expect to see persistent representation continuing into the outcome window. Alternatively, if the encoding of chosen identity is temporally limited to the response period this would exclude the possibility of their use for tagging an outcome-driven value update.

We thus extended our analysis of chosen identity to the outcome window (see Supplemental Figure B.4). This analysis revealed robustly maintained encoding of chosen actions in the Action Task and chosen stimuli in the Stimulus Task across prefrontal ROIs (see Supplemental Figure B.4A, C). When plotting the time course of these signals we find strong persistent encoding of the chosen action into the outcome window particularly for the ACC and preSMA in the action task (see Supplemental Figure B.4B). For the stimulus task this temporally resolved analysis shows maintained encoding (particularly) in vmPFC and what appears as re-instantiation in the ACC and preSMA (see Supplemental Figure B.4D). This finding stands in contrast to the encoding of chosen action in the stimulus task, for which we observe a decline in encoding prevalence post choice execution (see




Figure 3.4: Chosen identity coding. See caption below.

Supplemental Figure B.4E, F). This temporally restricted encoding to the response-window confirms that chosen action coding in the stimulus task is being used exclusively for choice execution. In contrast, the observed pattern in domain-relevant chosen identity coding is consistent with what we would hypothesize if these signals fulfilled a role in updated-value assignment (see Discussion).

Discussion

An extensive debate has revolved around whether stimulus- and action-based choices rely on a shared neural architecture, with lesion studies and human fMRI work appearing to find opposing evidence (e.g., Rudebeck et al., 2008; Gläscher, Hampton, and O'Doherty, 2009). By recording from single neurons across a distributed pre-

Figure 3.4: (continued) **A)** Proportion of significant neurons encoding the chosen action ($Chosen_{action}$) during the action task across four regions of interest: ventromedial prefrontal cortex (vmPFC), pre-supplementary motor area (preSMA), anterior cingulate cortex (ACC), and Amygdala. Dark gray bars indicate regions where the percentage of significant units exceeds chance (at $p < 0.01$); white bars represent non-significant populations. Dashed horizontal lines indicate the region-specific chance levels. **B)** Time-resolved proportion of significant units encoding the chosen action relative to response execution (dashed vertical line at 0 ms) in the action task. Solid lines denote the proportion for each brain region computed using a sliding window approach (see Methods). **C)** Proportion of significant neurons encoding the chosen stimulus ($Chosen_{stimulus}$) during the stimulus task. Conventions follow panel A. **D)** Time-resolved proportion of significant units encoding the chosen stimulus relative to the response during the stimulus task. Conventions follow panel B. **E)** Proportion of significant neurons encoding the chosen action during the stimulus task. Conventions follow panel A. **F)** Time-resolved proportion of significant units encoding the chosen action relative to the response during the stimulus task. Conventions follow panel B. **G)** Correlation between model statistics for chosen identity across tasks. The bar plot displays the Spearman rank correlation (ρ) comparing unit model statistics for chosen action versus chosen stimulus within the stimulus task (purple), and chosen action in the stimulus task versus the action task (orange) across brain regions. Filled bars indicated a significant $p < 0.01$ correlation while outlined bars are non-significant correlations. **H)** Generalization performance of chosen identity decoding across brain areas: vmPFC, preSMA, ACC, and Amygdala. Box plots show cross-condition decoding accuracy for different train-test combinations. Lighter red boxes show decoding performance for models trained on action identity in the chosen action task and tested on the chosen action identity in the stimulus task. Darker red boxes show the reverse: models trained on the chosen action identity in the stimulus task and test on the action task. Asterisks denote decoding performance significantly above the chance level ($*** : p < 0.001$). The chance threshold ($p < 0.05$) is indicated with a black line while the mean of the permuted chance distribution (see Methods). **I)** Layout as in Figure 3.2K. Likelihood ratio statistic for individual units in the preSMA for either chosen action in the stimulus task (on y-axis, with significance in blue) or chosen action in the action task (on the x-axis, with significance in red). There are a few neurons that code for chosen action significantly across both tasks (labeled in purple). **J)** Likelihood ratio statistics of chosen stimulus (in stimulus task, in blue along the y-axis) and chosen action (as in Figure 3.4I). **K)** Latency analysis for preSMA units significantly coding for chosen stimulus identity or chosen action identity in the stimulus task. Latency distributions compared using two-sided Wilcoxon rank-sum test ($*** : p < 0.001$). Analysis run aligned to response (indicated by time 0).

frontal network during both action- and stimulus-based decisions, our results unify these findings by suggesting that the brain relies on a dynamic architecture within which the representations of action and stimulus value reorganize as the decision unfolds. Prior to choice commitment, valuation recruits broad, partially overlapping populations of neurons across the prefrontal cortex across both choice-domains. However, upon commitment to a choice, these populations undergo a remarkable functional bifurcation, separating into highly specialized, modality-specific codes.

In the pre-decision valuation phase neurons across the prefrontal cortex (including both the vmPFC and ACC) code the decision value for both actions and stimuli. These representations of action- and stimulus-based value use a shared neural substrate, where the a neuron coding for one value-domain is more likely to also code for the other. We want to briefly note however, that this does not prove the existence of a true ‘common currency’ across these choice domains (Padoa-Schioppa, 2011; Levy and Glimcher, 2012). While we do find a small subset of neurons that similarly encode action- and stimulus value, there are also preliminary signs of a domain-specific value coding scheme within neurons in the ACC. Future research will have to probe the population-level geometry of these value representations to determine if these value computations indeed rely on (semi-) orthogonal subspaces (as suggested (although strictly at the population-level) in Johnston et al., 2024). These findings resonate with human fMRI literature reporting overlapping BOLD activations for both stimulus and action values in prefrontal cortex regions (Gläscher, Hampton, and O’Doherty, 2009). Crucially, our temporally resolved single-unit data reveals that any overlapping encoding is a strictly pre-decision phenomenon.

As the decision process transitions to the commitment phase we observe a more segregated, modality-specific architecture. This suggests the prefrontal value encoding network is actively reshaping its representations to separately encode the value of chosen actions and stimuli. Such a regional specialization across choice-domains could facilitate the correct encoding of the “value-relevance” of either the action and stimulus. Inappropriately assigning the outcome to, for example, the action chosen in a stimulus-based decision could be detrimental for value learning. Indeed, segregating the representation of chosen action values from chosen stimulus values would allow the brain to directly link selected candidate to the relevant domain (as might be necessary for credit assignment see below). Reflecting this need for compartmentalization, we observe that domain-overlap ends during the post-choice valuation phase, revealing a stark double dissociation between vmPFC and ACC. The ACC

specializes by selectively maintaining action-based chosen value representations. This finding extends classical non-human primate literature (e.g., Hayden and Platt, 2010) by demonstrating that this specialization applies even to more clearly separated motor actions. Conversely, the vmPFC selectively maintains stimulus-based chosen value representations, supporting its role as a post-choice evaluator of stimuli (Wunderlich, Rangel, and O’Doherty, 2010; Hare et al., 2011). This network-level dichotomy provides a compelling neuronal account for macro-level lesion studies that reported specific action or stimulus selection deficits following damage to the ACC and OFC, respectively (Rudebeck et al., 2008; Camille, Tsuchida, and Fellows, 2011).

Additionally, in the commitment phase, we find the prefrontal cortex broadly represents the identity of the chosen option across both action- and stimulus-based choices. For example, both the vmPFC and ACC maintain this cross-task encoding for choice identity (dynamically tracking the chosen action or chosen stimulus) but they do so utilizing largely independent sub-populations within these regions. Beyond immediate choice execution, this post-decision representation of chosen identity and value within the relevant choice domain is necessary for effective credit assignment in value-based learning. When an outcome (i.e., win or no-win) is received after choice, the brain must determine whether to attribute that outcome to the physical action taken or the visual stimulus chosen. We hypothesize that by selectively isolating the chosen action value within the ACC and the chosen stimulus value within the vmPFC, the brain establishes dedicated structural channels as it prepares the input to an outcome-based value update. Such an architecture would allow prediction error signals to selectively update the value of the relevant choice dimension.

We are, however, limited in our observation of this update process due to our clinically defined recording locations. Direct signals of such value updates have been found in lateral PFC, basal-ganglia and the midbrain (e.g., Valentin and O’Doherty, 2009; Asaad and Eskandar, 2011; Lak, Stauffer, and Schultz, 2014; Man et al., 2024). Despite this limitation, we do find a widespread and persistent chosen identity coding that would be consistent with these signals fulfilling a secondary function beyond correct choice execution. The chosen identity’s function in selecting the relevant option would only require it to be presented until that selection is complete. Indeed, while we observe this post-response ramp down for chosen action identity in the stimulus task, the task-relevant chosen identity persists into the outcome

window. Such a persistent encoding would be expected if the chosen identity additionally facilitates a tagging mechanism linking an updated value to the correct item. Future research should aim to further elucidate the updating of both action- and stimulus-based value signals. Ideally these studies will be able to directly record from areas of basal-ganglia in conjunction with a wide set of prefrontal regions (including lateral PFC). A combined recording across these regions will be able to expand our analysis of value signals across choice-domains to additionally include reward-based updates.

Outside of the various valuation signals found in the vmPFC and ACC, our results position the preSMA as a highly dynamic encoder of both stimulus and action information. Rather than acting as a simple final motor output region, the preSMA exhibits highly flexible encoding of pre-decision value, chosen value and chosen identity regardless of the choice domain. Furthermore, we observed a distinct temporal cascade of single-neuron activity within the preSMA during stimulus-based decisions. Specifically, the encoding of the chosen stimulus identity significantly precedes the encoding of the chosen action. This sequential encoding mirrors observations in the supplementary eye field (Chen and Stuphorn, 2015) and strongly aligns with an active 'good-to-action' transformation (Cai and Padoa-Schioppa, 2014).

Consequently, the preSMA's role extends far beyond a passive motor relay, especially relevant to stimulus-to-action choices, it acts as an active across-domain transformer. In computational terms, while the vmPFC resolves the value of a stimulus, it lacks the mapping of this stimulus to the sensorimotor contingencies required to interact with it. The robust, sequential encoding of stimulus identity followed by action mapping within the preSMA suggests this region houses the transformational logic required to convert a chosen stimulus-value into a concrete behavioral policy. By holding both the stimulus identity and the motor effector, the preSMA dynamically routes the output of the prefrontal valuation network into the specific motor effectors required for goal acquisition, effectively acting as the central hub where value-based intention is translated into physical consequence.

Ultimately, our results provide a framework for understanding how the human frontal cortex bridges the gap between abstract valuation and concrete choice execution. We demonstrate that translating value into choice is not achieved through a single, static pathway. Instead, the brain leverages highly dynamic representations, relying on a shared set of neurons for the encoding of pre-decision values, before rapidly shifting into specialized, modality-specific roles to value and execute the finalized

choice.

Acknowledgments

We would like to thank all members of the O'Doherty and Rutishauser laboratories for their feedback. We thank the participants and their families for their participation, and nurses and medical staff for their work. We would like to acknowledge funding from the following grants; NIH MH133729 and NSF 2318899 to J.P.D. and U.R., as well as NIH U01NS117839 to U.R.

Materials and Methods

Participants

Sixteen adult patients (see Table B.1 for demographics) with drug-resistant epilepsy participated in this study. All patients were undergoing stereotactic depth electrode implantation for seizure localization and potential surgical treatment at Cedars- Sinai Medical Center. Electrode placement was determined solely on clinical grounds by the clinical care team. Each patient provided written informed consent (Feinsinger et al., 2022) to participate in research protocols approved by the Institutional Review Boards of both the California Institute of Technology (IR23-1313) and Cedars-Sinai Medical Center (STUDY00000572).

Task Design

General Experimental Structure

Participants (see Table B.1 for demographics) performed two variants of a probabilistic reinforcement learning task: an Action Task and a Stimulus Task. Each trial followed a consistent temporal structure: a fixation cross (500–1500 ms), a response window (max 3000 ms), a selection confirmation (750 ms), and a final outcome phase (1250 ms). Rewards were represented by a yellow coin (value = 1), while non-rewarded trials were indicated by a red coin (value = 0). The order of these tasks was counterbalanced across sessions (Stimulus Task first: N=13 sessions). Reward probabilities were (re-)set at the start of each block (36 trials) and were pseudo-randomly drawn from the following set [$P = 0.8$, $P = 0.6$, $P = 0.4$, $P = 0.2$] such that each choice option was matched with a unique reward probability.

Action-Based Learning Task

In the action task, reward probabilities were mapped directly to specific motor actions (left hand, right hand, left foot, or right foot press). In each block, a subset of three actions was available for selection. Participants had to learn which action yielded the highest probability of reward over 36 trials. Participants completed between 4-12 blocks of the action task (5.7 blocks on average across sessions).

Stimulus-Based Learning Task

The stimulus task presented a binary choice between two stimuli drawn from a set of three available within a block (out of the total set of four; tree, house, dog, and camera). To decouple reward learning from motor responses, stimulus locations

(top or bottom and left or right) were pseudo-randomized. Top locations required a hand press, while bottom locations required a foot press. Reward probability was dependent solely on the stimulus identity. Participants completed between 3-9 blocks of the stimulus task (5.25 blocks on average).

Behavioral Data Analysis

Learning Curves and Choice Probability

To visualize behavioral adaptation over the course of the block, we generated per-subject learning curves by plotting the probability of selecting the highest-value option across trial indices. A linear regression was fitted to each individual subject's choice data across the 36 trials to capture individual learning trajectories.

To determine if participants successfully learned the task parameters, we isolated behavioral performance during the final 10 trials (trials 27–36) of each block. Overall group performance was tested using a one-sided Wilcoxon signed-rank test, comparing the median accuracies of all subjects in the final 10 trials against chance.

Reward History and Choice Continuity (GLMs)

To quantify how recent reward history influenced choice perseveration, we implemented a combined-history logistic regression model. The model predicted the binary outcome of choice continuity (staying with the previous choice vs. switching) as a function of the rewards received on the preceding three trials ($t - 1$, $t - 2$, and $t - 3$).

The logistic regression was specified as:

$$P(\text{Stay}_t) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 R_{t-1} + \beta_2 R_{t-2} + \beta_3 R_{t-3})}}$$

where R_{t-k} represents the binary reward outcome at lag k . From the fitted model, we extracted the logit coefficients (β) to represent the log-odds of staying, as well as the marginal probability increase. The probability increase was calculated by taking the difference between the baseline probability of staying (derived from the intercept β_0) and the probability of staying given a reward at a specific lag. Group-level significance for the effect of each lag was assessed by testing the subject-level coefficients against zero using a one-sample t -test.

Computational Modeling

Reinforcement learning - 2 parameter

We modeled participant choice behavior using a standard reinforcement learning framework that relied solely on expected values (Q).

Choice Rule

The probability P of choosing option i on trial t was determined by a softmax choice rule:

$$P(choice_t = i) = \frac{\exp(\beta \cdot Q_{i,t})}{\sum_{j \in available} \exp(\beta \cdot Q_{j,t})} \quad (3.1)$$

where β is the inverse temperature parameter governing choice stochasticity.

Update Rules

Q-values were updated based on the reward prediction error:

$$Q_{chosen,t+1} = Q_{chosen,t} + \alpha(R_t - Q_{chosen,t}) \quad (3.2)$$

where R_t is the reward and α is the learning rate.

Model Fitting

Parameters (α, β) were estimated by minimizing the regularized negative log-likelihood:

$$-\log(L) = -\sum \log(P_{chosen}) - \sum \log(\mathcal{N}(params | \mu = 0, \sigma = 2)) \quad (3.3)$$

We utilized a normal distribution prior ($\mu = 0, \sigma = 2$) to regularize the parameters during the optimization process.

Choice kernel only model

We modeled participant choice behavior using a framework that relied entirely on choice history (CK), completely omitting reward-based value learning.

Choice Rule

The probability P of choosing option i on trial t was determined by a choice-history softmax rule:

$$P(choice_t = i) = \frac{\exp(\beta_{CK} \cdot CK_{i,t})}{\sum_{j \in available} \exp(\beta_{CK} \cdot CK_{j,t})} \quad (3.4)$$

where β_{CK} represents the strength of choice perseveration (or alternation, if negative).

Update Rules

The choice kernel was updated exclusively to track past choices:

$$CK_{i,t+1} = CK_{i,t} + \alpha_{CK}(D_{i,t} - CK_{i,t}) \quad (3.5)$$

where $D_{i,t}$ is a dummy variable (1 if chosen, 0 otherwise) and α_{CK} is the kernel learning rate (decay).

Model Fitting

Parameters $(\alpha_{CK}, \beta_{CK})$ were estimated using the regularized negative log-likelihood optimization described in Equation 3.3.

Reinforcement learning - 4 parameter

We modeled participant choice behavior using a hybrid reinforcement learning framework that integrated both expected values (Q) and choice history (CK).

Choice Rule

The probability P of choosing option i on trial t was determined by a combined softmax choice rule:

$$P(choice_t = i) = \frac{\exp(\beta \cdot Q_{i,t} + \beta_{CK} \cdot CK_{i,t})}{\sum_{j \in available} \exp(\beta \cdot Q_{j,t} + \beta_{CK} \cdot CK_{j,t})} \quad (3.6)$$

where β is the inverse temperature parameter and β_{CK} represents the strength of choice perseveration.

Update Rules

Both systems operated in parallel: Q-values were updated based on reward prediction errors (Equation 3.2), and the choice kernel was simultaneously updated to track choice history (Equation 3.5).

Model Fitting

The full parameter set $(\alpha, \beta, \alpha_{CK}, \beta_{CK})$ was estimated utilizing the identical regularized optimization procedure established in Equation 3.3.

Posterior Predictive Checks

Simulation Procedure

To evaluate how well our fitted models captured the empirical behavioral patterns, we conducted posterior predictive checks. For each participant, we generated synthetic data using their best-fitting parameters for three models: a standard Reinforcement Learning model (RL2; α, β), a Choice Kernel model (CK; α_{CK}, β_{CK}), and a hybrid model combining both (RL4; $\alpha, \beta, \alpha_{CK}, \beta_{CK}$).

To account for the stochastic nature of the models' decision rules, we simulated 100 independent experimental runs per participant, per model, for both the Action and Stimulus tasks. Each simulation strictly mirrored the empirical task structure experienced by that specific participant, including the exact number of blocks, trials, and the specific trial-by-trial bandit configurations. To ensure reproducibility, simulation stochasticity was controlled using unique random seeds derived from participant identifiers.

Behavioral Metrics

We assessed model performance by comparing the synthetic data to the empirical data across key behavioral signatures of reinforcement learning:

- **Win-Stay / Lose-Shift Probabilities:** We calculated the raw probability of a participant repeating their previous choice, conditioned separately on whether the previous trial resulted in a reward or no reward. Trials where the previously chosen option was no longer available were excluded from this analysis.
- **Lagged Reward Effects:** To quantify the extended impact of past rewards on current choices, we fitted trial-by-trial logistic regression models predicting

the probability of staying with the same choice. The model included the outcomes of the three most recent trials (lags $t - 1$, $t - 2$, and $t - 3$) as predictors.

Simulation Aggregation

For the simulated datasets, the behavioral metrics (stay probabilities and logistic regression estimates) were calculated independently for each of the 100 simulated runs. These values were then averaged within each participant to produce a single, stable model prediction per metric. These averaged predictions were directly compared against the metrics derived from the participants' actual empirical data to evaluate model fit.

Neural data

Electrophysiology and Recording

We used Behnke-Fried hybrid depth electrodes (AdTech Medical), positioned exclusively according to clinical criteria (Carlson, Rutishauser, and Mamelak, 2018)(see Supplementary Figure B.1 for electrode locations). Broadband extracellular recordings were performed with a sampling rate of 32 kHz and a bandpass of 0.1–9,000 Hz (ATLAS System, Neuralynx). The dataset of 902 single units reported in this study was obtained bilaterally from the ventromedial prefrontal cortex (vmPFC), anterior cingulate cortex (ACC), pre-supplementary motor area (preSMA), and amygdala, with typically one macroelectrode targeted to each region per hemisphere. Each macroelectrode contained a bundle of eight 40- μm microelectrodes extending past the macroelectrode tip. Recordings were bipolar, utilizing one microelectrode in each bundle as a local reference to isolate localized single-unit activity.

Spike Sorting and Unit Isolation

Spike sorting and unit isolation were performed offline using the open-source OSort algorithm (Rutishauser, Schuman, and Mamelak, 2006). Spike detection was performed using an amplitude threshold set to five standard deviations above the median noise level of the filtered signal.

Following detection, the OSort algorithm automatically clustered the extracted waveforms. This automated step grouped spikes into distinct putative units based on a distance metric utilizing their waveform features. After automatic clustering, all putative units underwent rigorous manual curation to ensure high data quality. We

evaluated each cluster based on waveform shape consistency, temporal stability across the recording session, and the distribution of inter-spike intervals (ISIs).

To ensure the purity of our single-unit isolations, we strictly enforced a physiological refractory period. Any cluster exhibiting more than a 3% violation rate within a 3 ms ISI period was classified as multi-unit activity or noise and excluded from the single-unit analyses. Furthermore, automated clusters representing mechanical or electrical artifacts were manually identified and discarded. This combination of automated sorting and rigorous manual refinement yielded the final dataset of 902 high-quality single units utilized for all subsequent analyses.

Single-unit encoding analysis

To characterize how individual neurons represent value and choice identity, we modeled single-unit spike counts using Poisson generalized linear models (GLMs). For our primary epoch-based analyses, spike counts were extracted during fixed task windows, such as a 1000 ms window immediately following stimulus onset (for option value results) or a 1000 ms window preceding the execution of the motor response (for chosen value and chosen identity results).

GLM's were constructed to include the variable of interest (e.g., chosen value, chosen identity) alongside relevant nuisance regressors (such as the availability of specific options on a given trial). Models were fit using Iteratively Reweighted Least Squares (IRLS); if convergence failed or algorithmic instability was detected, the optimizer automatically fell back to the L-BFGS algorithm.

The explanatory power of each variable was evaluated using a likelihood ratio test (LRT), comparing the log-likelihood of the full model to a reduced model omitting the specific term of interest (e.g. option values across the 4 stimuli). To robustly establish statistical significance and account for non-normal spike distributions, we computed empirical p -values via permutation testing. For each neuron, the variable of interest was randomly shuffled across trials 500 times. The empirical p -value was calculated as the proportion of permuted likelihood ratio statistics that met or exceeded the true, unpermuted statistic.

Proportion of significant units

To quantify whether a given brain region represented a variable of interest at the population level, we calculated the proportion of individually significant units within that region. A single unit was classified as significantly encoding the variable (e.g.,

chosen Q -value) if the p -value for the corresponding regressor in the first-level generalized linear model (GLM) fell below our predefined alpha level ($\alpha = 0.05$). The total number of units included in the GLM for a specific region and task served as the denominator. To statistically determine if the observed proportion of significant units in a given region was greater than what would be expected by random chance, we performed a one-sided binomial test. The expected chance probability was mathematically tied to the single-unit significance threshold ($p = 0.05$). The resulting binomial p -values were evaluated at a significance threshold of $p < 0.01$.

Double-dissociation analysis

To formally test for functional specialization between regions across task domains (i.e., a double dissociation), we evaluated whether the likelihood of a neuron encoding chosen value was dependent on the interaction between brain region and task domain. We extracted the single-unit significance status (1 = significant at $\alpha = 0.05$, 0 = not significant) for all units recorded in the predefined regions of interest (vmPFC and ACC) during both the stimulus and action tasks. Units included in the GLM for only one task were retained in the denominator for that specific task to ensure accurate, independent population proportions. We then fit a logistic regression model to predict the binary significance status of each unit. The model included categorical predictors for Brain Region (vmPFC vs. ACC), Task Domain (Stimulus vs. Action), and their interaction term (Region $\times$ Task). The formal test for the double dissociation was evaluated based on the statistical significance of the Region $\times$ Task interaction coefficient.

Sliding-window encoding analysis

To capture the fine-grained temporal dynamics of value and identity representations, we extended our GLM framework using a sliding-window approach. For each trial, spike counts were extracted within 500 ms rolling windows, advancing in 50 ms increments. These sliding windows were aligned to either the onset of the stimulus (spanning -500 to 1000 ms) or the execution of the motor response (spanning -1000 to 500 ms). The exact same joint Poisson GLM, optimization, and permutation testing procedures described above were applied independently to each time window. Units that survived the permutation at $p < 0.05$ are included in the proportion.

Latency analysis

To determine the precise timing of neural responses relative to task events, we computed response latencies for individual single units. This analysis was restricted to the subpopulation of units that demonstrated significant task-related encoding ($p <$

0.05, determined via the permutation-tested generalized linear models described above) to ensure latencies were only calculated for functionally responsive neurons.

Latency was defined as the first point in time at which a neuron's spiking rate significantly deviated from its expected baseline firing rate. We estimated this baseline rate using a 500 ms window immediately preceding stimulus onset. To maximize robustness, particularly for neurons with sparse baseline firing, we calculated a session-wide baseline firing rate for each neuron by aggregating all pre-stimulus baseline spikes across the entire recording session.

For each trial, we iteratively scanned the spike train relative to the event of interest (e.g., motor response). Moving temporally away from the event, we tallied the cumulative number of spikes and calculated the exact time elapsed for each incoming spike. At each spike time, we computed the expected number of spikes based on the session-wide baseline Poisson rate. We then utilized the Poisson cumulative distribution function (CDF) and survival function (SF) to test whether the observed cumulative spike count represented a statistically significant deviation ($p < 0.05$) from the baseline expectation.

A significant increase in firing (a burst) was identified if the probability of observing the tallied number of spikes (or more) was less than 0.05. Conversely, a significant decrease in firing (a suppression) was identified if the probability of observing the tallied number of spikes (or fewer) fell below 0.05. The exact time elapsed to this first significant deviation was recorded as the trial's response latency.

To compare the temporal dynamics of these responses across brain regions and task domains, we aggregated the significant trial-level latencies for each region. Differences in the median latency distributions between specific regions, as well as between the stimulus-based and action-based tasks within a given region, were evaluated for statistical significance using two-sided non-parametric Wilcoxon rank-sum tests.

Decoding generalization analysis

To determine whether representations of choice identity generalize across decision domains, we performed a cross-task population decoding analysis. For each brain region, we constructed pseudo-populations by pooling independently recorded single units across all participants. Spike counts for each unit were extracted during the window of interest and Z-scored. To prevent class imbalances from biasing the

decoder, we randomly subsampled trials to ensure an equal number of observations per condition. Pseudo-population feature matrices were standardized prior to model fitting.

Since chosen option identity is a categorical variables, we utilized a Linear Support Vector Classifier (SVC) with L2 regularization. To explicitly test for representational generalization, each model was trained exclusively on data from one task domain (e.g., the action-based task) and tested on unseen data from the alternative domain (e.g., the stimulus-based task). Model performance was quantified using classification accuracy.

Statistical significance was established using a rigorous permutation and bootstrapping approach. First, to account for variance introduced by the random trial subsampling, the entire pseudo-population generation and cross-task testing pipeline was repeated across 50 independent random seeds. Within each seed, we generated a local null distribution by shuffling the test labels 100 times. To compute the final overarching significance, we employed a “null of means” bootstrapping procedure: we drew one permuted score from each of the 50 seeds to calculate a null mean, repeating this process 5,000 times. The empirical p -value was then calculated as the proportion of these bootstrapped null means that performed better than or equal to the true observed mean across the 50 seeds (i.e., higher accuracy for categorical variables).

References

- Aquino, T. G. et al. (Mar. 2023). “Neurons in human pre-supplementary motor area encode key computations for value-based choice”. en. In: *Nature Human Behaviour* 7.6, pp. 970–985. ISSN: 2397-3374. DOI: 10.1038/s41562-023-01548-2.
- Asaad, W. F. and E. N. Eskandar (Dec. 2011). “Encoding of Both Positive and Negative Reward Prediction Errors by Neurons of the Primate Lateral Prefrontal Cortex and Caudate Nucleus”. en. In: *The Journal of Neuroscience* 31.49, pp. 17772–17787. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3793-11.2011.
- Cai, X. and C. Padoa-Schioppa (Mar. 2014). “Contributions of Orbitofrontal and Lateral Prefrontal Cortices to Economic Choice and the Good-to-Action Transformation”. en. In: *Neuron* 81.5, pp. 1140–1151. ISSN: 08966273. DOI: 10.1016/j.neuron.2014.01.008.
- Camille, N., A. Tsuchida, and L. K. Fellows (Oct. 2011). “Double Dissociation of Stimulus-Value and Action-Value Learning in Humans with Orbitofrontal

- or Anterior Cingulate Cortex Damage”. en. In: *Journal of Neuroscience* 31.42, pp. 15048–15052. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3164-11.2011.
- Carlson, A. A., U. Rutishauser, and A. N. Mamelak (2018). “Safety and Utility of Hybrid Depth Electrodes for Seizure Localization and Single-Unit Neuronal Recording”. en. In: *Stereotactic and Functional Neurosurgery* 96.5, pp. 311–319. ISSN: 1011-6125, 1423-0372. DOI: 10.1159/000493548.
- Chen, X. and V. Stuphorn (Nov. 2015). “Sequential selection of economic good and action in medial frontal cortex of macaques during value-based decisions”. en. In: *eLife* 4, e09418. ISSN: 2050-084X. DOI: 10.7554/eLife.09418.
- Feinsinger, A. et al. (Jan. 2022). “Ethical commitments, principles, and practices guiding intracranial neuroscientific research in humans”. en. In: *Neuron* 110.2, pp. 188–194. ISSN: 08966273. DOI: 10.1016/j.neuron.2021.11.011.
- Fine, J. M. et al. (June 2023). “Abstract Value Encoding in Neural Populations But Not Single Neurons”. en. In: *The Journal of Neuroscience* 43.25, pp. 4650–4663. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.1954-22.2023.
- Gläscher, J., A. N. Hampton, and J. P. O’Doherty (Feb. 2009). “Determining a Role for Ventromedial Prefrontal Cortex in Encoding Action-Based Value Signals During Reward-Related Decision Making”. en. In: *Cerebral Cortex* 19.2, pp. 483–495. ISSN: 1460-2199, 1047-3211. DOI: 10.1093/cercor/bhn098.
- Hare, T. A. et al. (Nov. 2011). “Transformation of stimulus value signals into motor commands during simple choice”. en. In: *Proceedings of the National Academy of Sciences* 108.44, pp. 18120–18125. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1109322108.
- Hayden, B. Y. and M. L. Platt (Mar. 2010). “Neurons in Anterior Cingulate Cortex Multiplex Information about Reward and Action”. en. In: *The Journal of Neuroscience* 30.9, pp. 3339–3346. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.4874-09.2010.
- Johnston, W. J. et al. (Nov. 2024). “Semi-orthogonal subspaces for value mediate a binding and generalization trade-off”. en. In: *Nature Neuroscience* 27.11, pp. 2218–2230. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/s41593-024-01758-5.
- Kennerley, S. W. et al. (July 2006). “Optimal decision making and the anterior cingulate cortex”. en. In: *Nature Neuroscience* 9.7, pp. 940–947. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn1724.
- Kennerley, S. W. and M. E. Walton (June 2011). “Decision making and reward in frontal cortex: Complementary evidence from neurophysiological and neuropsychological studies.” en. In: *Behavioral Neuroscience* 125.3, pp. 297–317. ISSN: 1939-0084, 0735-7044. DOI: 10.1037/a0023575.

- Lak, A., W. R. Stauffer, and W. Schultz (2014). “Dopamine prediction error responses integrate subjective value from different reward dimensions”. en. In: *Proceedings of the National Academy of Sciences* 111.6, pp. 2343–2348. DOI: 10.1073/pnas.1321596111.
- Levy, D. J. and P. W. Glimcher (2012). “The root of all value: a neural common currency for choice”. en. In: *Current Opinion in Neurobiology* 22.6, pp. 1027–1038. DOI: 10.1016/j.conb.2012.06.001.
- Man, V. et al. (Mar. 2024). “Temporally organized representations of reward and risk in the human brain”. en. In: *Nature Communications* 15.1, p. 2162. ISSN: 2041-1723. DOI: 10.1038/s41467-024-46094-1.
- Padoa-Schioppa, C. (July 2011). “Neurobiology of Economic Choice: A Good-Based Model”. en. In: *Annual Review of Neuroscience* 34.1, pp. 333–359. ISSN: 0147-006X, 1545-4126. DOI: 10.1146/annurev-neuro-061010-113648.
- Padoa-Schioppa, C. and J. A. Assad (May 2006). “Neurons in the orbitofrontal cortex encode economic value”. en. In: *Nature* 441.7090, pp. 223–226. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature04676.
- Rangel, A., C. Camerer, and P. R. Montague (July 2008). “A framework for studying the neurobiology of value-based decision making”. en. In: *Nature Reviews Neuroscience* 9.7, pp. 545–556. ISSN: 1471-003X, 1471-0048. DOI: 10.1038/nrn2357.
- Rudebeck, P. H. et al. (Dec. 2008). “Frontal Cortex Subregions Play Distinct Roles in Choices between Actions and Stimuli”. en. In: *The Journal of Neuroscience* 28.51, pp. 13775–13785. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3541-08.2008.
- Rutishauser, U., E. M. Schuman, and A. N. Mamelak (June 2006). “Online detection and sorting of extracellularly recorded action potentials in human medial temporal lobe recordings, in vivo”. en. In: *Journal of Neuroscience Methods* 154.1-2, pp. 204–224. ISSN: 01650270. DOI: 10.1016/j.jneumeth.2005.12.033.
- Samejima, K. et al. (Nov. 2005). “Representation of Action-Specific Reward Values in the Striatum”. en. In: *Science* 310.5752, pp. 1337–1340. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1115270.
- Scangos, K. W. and V. Stuphorn (Feb. 2010). “Medial Frontal Cortex Motivates But Does Not Control Movement Initiation in the Countermanding Task”. en. In: *The Journal of Neuroscience* 30.5, pp. 1968–1982. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.4509-09.2010.
- Seo, H. and D. Lee (Aug. 2007). “Temporal Filtering of Reward Signals in the Dorsal Anterior Cingulate Cortex during a Mixed-Strategy Game”. en. In: *The Journal of Neuroscience* 27.31, pp. 8366–8377. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.2369-07.2007.

- Shi, W., S. Ballesta, and C. Padoa-Schioppa (Jan. 2022). “Economic Choices under Simultaneous or Sequential Offers Rely on the Same Neural Circuit”. en. In: *The Journal of Neuroscience* 42.1, pp. 33–43. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.1265-21.2021.
- Strait, C. E., T. C. Blanchard, and B. Y. Hayden (June 2014). “Reward Value Comparison via Mutual Inhibition in Ventromedial Prefrontal Cortex”. en. In: *Neuron* 82.6, pp. 1357–1366. ISSN: 08966273. DOI: 10.1016/j.neuron.2014.04.032.
- Sutton, R. S., D. Precup, and S. Singh (Aug. 1999). “Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning”. en. In: *Artificial Intelligence* 112.1-2, pp. 181–211. ISSN: 00043702. DOI: 10.1016/S0004-3702(99)00052-1.
- Valentin, V. V. and J. P. O’Doherty (Dec. 2009). “Overlapping Prediction Errors in Dorsal Striatum During Instrumental Learning With Juice and Money Reward in the Human Brain”. en. In: *Journal of Neurophysiology* 102.6, pp. 3384–3391. ISSN: 0022-3077, 1522-1598. DOI: 10.1152/jn.91195.2008.
- Wunderlich, K., A. Rangel, and J. P. O’Doherty (Aug. 2010). “Economic choices can be made using only stimulus values”. en. In: *Proceedings of the National Academy of Sciences* 107.34, pp. 15005–15010. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1002258107.

Chapter 4

GENERAL DISCUSSION

4.1 Summary of main findings

Both projects in this thesis start with the premise that the brain does not rely on a static value lookup table. Instead, value is a highly dynamic construct, computed on the fly through a multi-faceted, dynamic neural architecture. The central aim of this dissertation was to elucidate the neural architectures that support adaptive valuation, specifically investigating how the brain navigates the computational challenges of context-dependent preferences and distinct operational choice domains. Together Chapter 2 and 3 reveal a valuation system that is both hierarchically organized and able to flexibly represent decision relevant variables in a manner that generalizes to support current choice requirements.

In Chapter 2, I used functional magnetic resonance imaging (fMRI) to investigate how the brain dynamically constructs value for multi-attribute objects under shifting goal-contexts. This study traced the steps of a hierarchical value computation architecture. While visual cortices represented context-independent raw stimulus attributes, regions within the prefrontal cortex transformed those attributes such that they become goal-dependent. Specifically, the ventromedial prefrontal cortex (vmPFC) and orbitofrontal cortex (OFC) were identified as key substrates for computing “attributes in value space,” effectively adapting the value of individual stimulus attributes to their value in the current context. These context-sensitive attribute signals were then integrated into a unified overall subjective value representation, which was predominantly localized along the dorsomedial prefrontal cortex (dmPFC).

Zooming in from this whole brain perspective, Chapter 3 leveraged human single-unit electrophysiology to elucidate how these value computations are implemented at the cellular level across different choice domains. By directly contrasting structurally parallel action-based and stimulus-based reinforcement learning tasks, the results revealed a highly dynamic and temporally shifting representational architecture. Prior to a decision, neurons within regions of the prefrontal cortex (vmPFC, ACC, and pre-supplementary motor area (preSMA)) broadly tracked available option values, facilitating comparison by allowing downstream processes to draw from

a consistent set of neurons. However, when looking at the representation of value locked to the chosen option the network exhibited more regionally separated encoding of value. A double dissociation emerged for chosen value: the ACC operated as a specialized node for maintaining action-based chosen value, whereas the vmPFC selectively tracked stimulus-based chosen value. Bridging this domain-specific divide, the preSMA reliably signaled the identity and value of the chosen option regardless of whether the underlying choice was primarily driven by spatial motor plans or visual stimuli.

4.2 ‘General’ value computation

A broadly debated topic in the study of value-based decision-making centers on whether the brain utilizes an abstracted “common currency” to compare disparate options, or whether valuation architectures are segregated based on the nature/characteristics of the option being evaluated (Cabanac, 1992; Shizgal and Conover, 1996; Montague and Berns, 2002). Various studies have posited that regions like the orbitofrontal cortex (OFC) and ventromedial prefrontal cortex (vmPFC) compute a domain-general economic value independent of sensorimotor contingencies (e.g., Padoa-Schioppa, 2007; Levy and Glimcher, 2012). This hypothesis is supported by functional imaging studies demonstrating overlapping BOLD representations for various types of choice options within the vmPFC and associated networks (Chib et al., 2009; Gläscher, Hampton, and O’Doherty, 2009). This dissertation offers some insight into such common coding since this offers one potential way to achieve flexible computations: the sharing of neural architectures to represent and compare disparate options on a unified scale. This type of adaptability requires the valuation system to first abstract options away from their specific attributes or sensorimotor contingencies. However, abstracting away too much option specific information could lead to problems in both the performing of relevant actions to obtain selected options, and in credit assignment post-decision. As such the brain needs to determine what information to abstract away and what specific option information to retain.

While the findings presented in this dissertation provide some evidence of abstraction in value computations, future research should investigate the existence of a fully abstracted value computation. Non-human primate studies have found such an abstracted value code for neurons in the orbitofrontal and ventromedial prefrontal cortices, which encode economic value entirely independently of the specific physical movements required to execute a choice (Padoa-Schioppa and Assad, 2006;

Strait, Blanchard, and Hayden, 2014). Probing these processes in the human brain, functional imaging work has consistently identified the vmPFC as a core hub for domain-general computation, demonstrating overlapping BOLD activity when participants evaluate diverse reward modalities and choice types (Gläscher, Hampton, and O’Doherty, 2009; Wunderlich, Rangel, and O’Doherty, 2010). Instead, we find individual neurons distributed across the vmPFC, ACC, and preSMA tracked available option values regardless of whether the choice was dictated by visual stimuli or motor actions during the initial evaluation phase (see Chapter 3). Especially in the ACC, there was a substantial subgroup of neurons that coded for value across tasks. While these neurons do not appear to rely on a true shared neural code for value, this finding does provide evidence of a shared neural substrate for representing the value of disparate choice options prior to choice commitment. We do find a small subset of neurons across vmPFC and preSMA that appear to code for the value of both stimuli and actions prior to choice commitment using a shared coding direction. However, given the small number of neurons in this analysis and the apparent unique value coding scheme employed by the majority of neurons in the ACC, our results point to the use of a shared neural substrate rather than a fully abstracted common currency. Future research will have to endeavor to understand the exact geometry of these value representation in the frontal cortex.

Some level of abstraction can similarly be observed in the computation of value for complex, multi-attribute objects (Chapter 2). The valuation system converges on an integrated value signal that is abstracted along two dimensions: the specific option attributes and the goal-context under consideration. This allows for the comparison of options with highly dissimilar attributes across vastly different goal-contexts on one unified scale. Previous research has extensively demonstrated that the vmPFC serves as a critical integration hub capable of collapsing multi-dimensional attributes into a single scalar value. For example, this region effectively integrates diverse attributes like the taste and health of a food item, or the magnitude and probability of a monetary reward (Suzuki, Cross, and O’Doherty, 2017; Lim, O’Doherty, and Rangel, 2013; Farashahi et al., 2019). Research by McNamee, Rangel, and O’Doherty (2013) showed that this integrated value signal is fundamentally abstract. Prefrontal cortex is shown to encode the value of entirely different categories of goods (e.g., food versus apparel) using overlapping, category-independent neural codes (McNamee, Rangel, and O’Doherty, 2013). We can think of this result as relying on a similar computation as discussed in Chapter 2; effectively stripping away the sensory particulars of the stimuli to compute a unified decision variable.

Building upon this foundation and previous work on static multi-attribute valuation (e.g., Kahnt et al., 2011), Chapter 2 implements the concept of abstraction as dynamically shifting goal-contexts and finds a remarkably similar ventral-dorsal gradient along the medial wall of the PFC (where more dorsal regions hold a more abstracted value code). These findings demonstrate that the brain abstracts not just across different physical attributes, but across shifting internal goals.

Despite some evidence supporting the use of a (more or less) abstracted valuation space, several important questions and limitations remain regarding the exact nature and extent of this abstraction in value computation. First, the experimental paradigms employed in this thesis have been specifically designed to isolated specific dimensions of abstraction (i.e., predefined goal-contexts (Chapter 2) or explicit motor versus stimulus mappings (Chapter 3)) but these design decisions inevitably constrain the parameter space. It remains an open question whether the prefrontal cortex utilizes this exact same value-computation architecture when evaluating vastly more complex, unconstrained real-world options, or highly abstract domains such as social and moral valuation (Ruff and Fehr, 2014). Furthermore, the classical conceptualization of a strictly univariate “common currency” or a “single scalar value” may be an oversimplification of the computations at hand. There is an active discussion on the existence of a more high-dimensional population-level abstract value code (see Johnston et al., 2024). Such a geometric perspective could additionally speak to the idea that value is not encoded as a purely one-dimensional scalar. Instead, value might be better understood as low-dimensional shared subspaces within a broader, multiplexed neural population that represents various option attributes (Yoo and Hayden, 2018; Ebitz and Hayden, 2021). Future research leveraging large-scale population recordings across a variety of tasks that elucidate generalized value computation, will be critical to determine the exact nature of value encoding at the level of individual neurons. Additionally, bridging this hypothesis of a ventral-dorsal spatial gradient of abstraction observed via fMRI (Chapter 2 and McNamee, Rangel, and O’Doherty (2013)) with the temporal dynamics at the single-unit level (Chapter 3) will pose a challenge for future research. Finally, the epilepsy patients who participated in this research (Chapter 3) were verified to be behaviorally similar to healthy controls. Additionally, the tasks were designed in consideration of potential cognitive and memory constraints. However, it should be noted that there are limitations to our ability to generalize these results to the healthy population. Patients performed these tasks under highly unusual conditions as part of a multi-day hospital stay and suffered epileptic seizures during this stay. Neural

activity altered by these seizures would not be hypothesized to have any relation to task-relevant variables, but it can be difficult to determine the influence of patients' condition on their overall brain function. Combining these results with general population fMRI studies, non-epileptic single-neuron recordings and non-human primate studies remains important to obtain a holistic understanding of the neural processes at play.

4.3 'Specific' value computation

While the capacity to abstract options into a domain-general "common currency" facilitates the comparison of disparate options, decision-making ultimately requires interacting with a specific physical environment. To execute a choice and learn from its outcomes, the brain must be able to re-anchor abstract values to their specific sensorimotor or perceptual properties. If the valuation system solely utilized a generalized architecture, it would face a severe credit assignment problem: upon receiving a reward, the brain would not know which specific physical action to reinforce or which specific visual stimulus to update in its internal model (e.g., Walton et al., 2011; Jocham et al., 2016). Therefore, a fully adaptive valuation architecture must possess the dynamic capacity to reorganize and retain highly specific option information when required by the phase of the decision process.

The single-neuron recordings presented in Chapter 3 demonstrate exactly this dynamic reorganization. While pre-decision signals were highly generalized, as the decision was being finalized, encoding shifted to a strict domain-specific divide: the ACC specialized in maintaining action-based chosen value, whereas the vmPFC selectively tracked stimulus-based chosen value. This single-neuron-level segregation provides a mechanistic explanation for classical macro-level lesion studies, which have shown that damage to the ACC specifically impairs action-value learning, whereas OFC/vmPFC damage impairs stimulus-value learning (Rudebeck et al., 2008; Camille, Tsuchida, and Fellows, 2011). By demonstrating this post-choice specialization, the results illustrate how the brain solves the credit assignment problem by routing finalized value signals to their relevant operational domains. Furthermore, by isolating purely motor-based choices, these findings clarify that the ACC's role in action valuation is distinct from spatial attention, disentangling a frequent confound in primate electrophysiology (see for discussion Padoa-Schioppa, 2011).

The necessity for retaining specific, context-dependent representations is also evident in the attribute-level fMRI data. In Chapter 2, before reaching full abstraction in

the dmPFC, regions like the vmPFC and OFC computed intermediate “attributes in value space.” This demonstrates that the brain does not immediately discard specific stimulus features during valuation. Instead, it actively maintains and weights these individual attributes according to the current goal-context, ensuring the resulting value is highly specific to the environment’s current demands. Crucially, translating these specifically evaluated, multi-attribute stimuli into the physical execution of a choice requires a dedicated bridging mechanism. Returning to the single-unit data in Chapter 3, the preSMA emerged as a prime candidate for this “good-to-action” transformation (Cai and Padoa-Schioppa, 2014). During stimulus-based decisions, individual neurons in the preSMA encoded the identity of the chosen stimulus significantly earlier than the identity of the chosen action. This cascading representation provides real-time single-unit evidence of the brain actively mapping a specific goal identity onto the precise sensorimotor contingencies required for its execution, expanding upon recent human single-unit findings in the preSMA (Aquino et al., 2023). The fact that this region represents the value of both actions and stimuli (both before and locked onto choice), but does so through largely independent populations of neurons across domains, further supports the idea that the preSMA plays a central role in bridging general evaluation with specific motor execution.

Despite these insights into how the prefrontal cortex routes specific value and identity signals, several questions remain regarding the boundaries of this segregation. Both projects in this thesis utilized relatively drastic goal and domain shifts. For example, the binary distinction between “stimulus” and “action” in Chapter 3, while experimentally necessary to isolate these choice domains, represents two extreme endpoints of a continuum. In naturalistic environments, decisions often require a simultaneous, interactive evaluation of both the specific object desired and the complex energetic costs of the actions required to navigate to it. It remains an open question how the ACC and vmPFC interact dynamically when credit must be assigned jointly to both a stimulus and a complex motor sequence. Future research employing continuous movement trajectories and more naturalistic paradigms (e.g., Yoo, Tu, and Hayden, 2021; Cisek and Kalaska, 2010) will be crucial to understanding unconstrained stimulus and action value interactions.

Similarly, while the results discussed in Chapter 2 closely align with previous research on multi-attribute value computation, future research must explore how the computation of “attributes in value space” shapes choice behavior across a broader spectrum of stimulus categories. Lastly, the fMRI methodology in Chapter

2 precluded analyzing the representations of attributes in value space at the scale of individual neurons. However, guided by regions like the lateral OFC where we observed both integrated value and attribute-in-value-space BOLD encoding, we might hypothesize that abstraction occurs within the same neuron as a function of time leading up to the decision. Such temporal multiplexing would closely mirror findings in non-human primate electrophysiology, where individual OFC neurons have been shown to transition from encoding specific stimulus identities to abstract choice values over the course of a single trial (Rich and Wallis, 2016). Mapping these temporal dynamics in the human brain could ultimately unify the spatial gradients of abstraction with the rapid temporal routing required for flexible valuation.

4.4 Concluding remarks

Ultimately, decision-making in dynamic environments requires balancing two opposing computational principles. The brain must abstract information enough to compare vastly different options, yet retain enough specificity to execute decisions and learn from their outcomes. The central conclusion of this dissertation is that the human prefrontal cortex solves this fundamental computational challenge not through a single, static valuation center, but through a highly dynamic neural architecture that shifts across both space and time.

By integrating whole brain dynamics with single-unit electrophysiology, the research presented here reveals that valuation is a cascading sequence of transformations. The process begins with the hierarchical construction of context-dependent “attributes in value space,” converges into a transient, domain-general shared neural substrate for option comparison, and finally fractures into strictly segregated, domain-specific circuits to facilitate the “good-to-action” transformation and ensure accurate credit assignment.

Understanding these precise biological mechanisms can be helpful beyond providing new insights to debates within neuroeconomics and cognitive neuroscience. Revealing how the brain computes value under these changing conditions can facilitate a richer understanding of the mechanisms underlying even more ‘traditional’ static value-based decision making experiments. In turn, obtaining a deeper understanding of the mechanisms employed by the brain could provide biological constraints for computational models of learning and decision-making. Especially for those who aim to build systems capable of navigating complex, shifting environments, the hierarchical and dynamic valuation architecture of the human brain can serve as a

theoretical foundation to inspire a path towards goal-directed behavior in agents.

References

- Aquino, T. G. et al. (Mar. 2023). “Neurons in human pre-supplementary motor area encode key computations for value-based choice”. en. In: *Nature Human Behaviour* 7.6, pp. 970–985. ISSN: 2397-3374. DOI: 10.1038/s41562-023-01548-2.
- Cabanac, M. (Mar. 1992). “Pleasure: the common currency”. en. In: *Journal of Theoretical Biology* 155.2, pp. 173–200. ISSN: 00225193. DOI: 10.1016/S0022-5193(05)80594-6.
- Cai, X. and C. Padoa-Schioppa (Mar. 2014). “Contributions of Orbitofrontal and Lateral Prefrontal Cortices to Economic Choice and the Good-to-Action Transformation”. en. In: *Neuron* 81.5, pp. 1140–1151. ISSN: 08966273. DOI: 10.1016/j.neuron.2014.01.008.
- Camille, N., A. Tsuchida, and L. K. Fellows (Oct. 2011). “Double Dissociation of Stimulus-Value and Action-Value Learning in Humans with Orbitofrontal or Anterior Cingulate Cortex Damage”. en. In: *Journal of Neuroscience* 31.42, pp. 15048–15052. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3164-11.2011.
- Chib, V. S. et al. (Sept. 2009). “Evidence for a Common Representation of Decision Values for Dissimilar Goods in Human Ventromedial Prefrontal Cortex”. en. In: *Journal of Neuroscience* 29.39, pp. 12315–12320. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.2575-09.2009.
- Cisek, P. and J. F. Kalaska (June 2010). “Neural Mechanisms for Interacting with a World Full of Action Choices”. en. In: *Annual Review of Neuroscience* 33.1, pp. 269–298. ISSN: 0147-006X, 1545-4126. DOI: 10.1146/annurev.neuro.051508.135409.
- Ebitz, R. B. and B. Y. Hayden (Oct. 2021). “The population doctrine in cognitive neuroscience”. en. In: *Neuron* 109.19, pp. 3055–3068. ISSN: 08966273. DOI: 10.1016/j.neuron.2021.07.011.
- Farashahi, S. et al. (Sept. 2019). “Flexible combination of reward information across primates”. en. In: *Nature Human Behaviour* 3.11, pp. 1215–1224. ISSN: 2397-3374. DOI: 10.1038/s41562-019-0714-3.
- Gläscher, J., A. N. Hampton, and J. P. O’Doherty (Feb. 2009). “Determining a Role for Ventromedial Prefrontal Cortex in Encoding Action-Based Value Signals During Reward-Related Decision Making”. en. In: *Cerebral Cortex* 19.2, pp. 483–495. ISSN: 1460-2199, 1047-3211. DOI: 10.1093/cercor/bhn098.
- Jocham, G. et al. (Apr. 2016). “Reward-Guided Learning with and without Causal Attribution”. en. In: *Neuron* 90.1, pp. 177–190. ISSN: 08966273. DOI: 10.1016/j.neuron.2016.02.018.

- Johnston, W. J. et al. (Nov. 2024). “Semi-orthogonal subspaces for value mediate a binding and generalization trade-off”. en. In: *Nature Neuroscience* 27.11, pp. 2218–2230. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/s41593-024-01758-5.
- Kahnt, T. et al. (May 2011). “Decoding different roles for vmPFC and dlPFC in multi-attribute decision making”. en. In: *NeuroImage* 56.2, pp. 709–715. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2010.05.058.
- Levy, D. J. and P. W. Glimcher (2012). “The root of all value: a neural common currency for choice”. en. In: *Current Opinion in Neurobiology* 22.6, pp. 1027–1038. DOI: 10.1016/j.conb.2012.06.001.
- Lim, S.-L., J. P. O’Doherty, and A. Rangel (May 2013). “Stimulus Value Signals in Ventromedial PFC Reflect the Integration of Attribute Value Signals Computed in Fusiform Gyrus and Posterior Superior Temporal Gyrus”. en. In: *The Journal of Neuroscience* 33.20, pp. 8729–8741. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.4809-12.2013.
- McNamee, D., A. Rangel, and J. P. O’Doherty (2013). “Category-dependent and category-independent goal-value codes in human ventromedial prefrontal cortex”. en. In: *Nature Neuroscience* 16.4, pp. 479–485. DOI: 10.1038/nn.3337.
- Montague, P. and G. S. Berns (Oct. 2002). “Neural Economics and the Biological Substrates of Valuation”. en. In: *Neuron* 36.2, pp. 265–284. ISSN: 08966273. DOI: 10.1016/S0896-6273(02)00974-1.
- Padoa-Schioppa, C. (July 2011). “Neurobiology of Economic Choice: A Good-Based Model”. en. In: *Annual Review of Neuroscience* 34.1, pp. 333–359. ISSN: 0147-006X, 1545-4126. DOI: 10.1146/annurev-neuro-061010-113648.
- Padoa-Schioppa, C. and J. A. Assad (May 2006). “Neurons in the orbitofrontal cortex encode economic value”. en. In: *Nature* 441.7090, pp. 223–226. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature04676.
- Padoa-Schioppa, C. (Dec. 2007). “Orbitofrontal Cortex and the Computation of Economic Value”. en. In: *Annals of the New York Academy of Sciences* 1121.1, pp. 232–253. ISSN: 0077-8923, 1749-6632. DOI: 10.1196/annals.1401.011.
- Rich, E. L. and J. D. Wallis (July 2016). “Decoding subjective decisions from orbitofrontal cortex”. en. In: *Nature Neuroscience* 19.7, pp. 973–980. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn.4320.
- Rudebeck, P. H. et al. (Dec. 2008). “Frontal Cortex Subregions Play Distinct Roles in Choices between Actions and Stimuli”. en. In: *The Journal of Neuroscience* 28.51, pp. 13775–13785. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3541-08.2008.
- Ruff, C. C. and E. Fehr (Aug. 2014). “The neurobiology of rewards and values in social decision making”. en. In: *Nature Reviews Neuroscience* 15.8, pp. 549–562. ISSN: 1471-003X, 1471-0048. DOI: 10.1038/nrn3776.

- Shizgal, P. and K. Conover (Apr. 1996). “On the Neural Computation of Utility”. en. In: *Current Directions in Psychological Science* 5.2, pp. 37–43. ISSN: 0963-7214, 1467-8721. DOI: 10.1111/1467-8721.ep10772715.
- Strait, C. E., T. C. Blanchard, and B. Y. Hayden (June 2014). “Reward Value Comparison via Mutual Inhibition in Ventromedial Prefrontal Cortex”. en. In: *Neuron* 82.6, pp. 1357–1366. ISSN: 08966273. DOI: 10.1016/j.neuron.2014.04.032.
- Suzuki, S., L. Cross, and J. P. O’Doherty (2017). “Elucidating the underlying components of food valuation in the human orbitofrontal cortex”. en. In: *Nature Neuroscience* 20.12, pp. 1780–1786. DOI: 10.1038/s41593-017-0008-x.
- Walton, M. E. et al. (Dec. 2011). “Giving credit where credit is due: orbitofrontal cortex and valuation in an uncertain world”. en. In: *Annals of the New York Academy of Sciences* 1239.1, pp. 14–24. ISSN: 0077-8923, 1749-6632. DOI: 10.1111/j.1749-6632.2011.06257.x.
- Wunderlich, K., A. Rangel, and J. P. O’Doherty (Aug. 2010). “Economic choices can be made using only stimulus values”. en. In: *Proceedings of the National Academy of Sciences* 107.34, pp. 15005–15010. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1002258107.
- Yoo, S. B. M. and B. Y. Hayden (Aug. 2018). “Economic Choice as an Untangling of Options into Actions”. en. In: *Neuron* 99.3, pp. 434–447. ISSN: 08966273. DOI: 10.1016/j.neuron.2018.06.038.
- Yoo, S. B. M., J. C. Tu, and B. Y. Hayden (Mar. 2021). “Multicentric tracking of multiple agents by anterior cingulate cortex during pursuit and evasion”. en. In: *Nature Communications* 12.1, p. 1985. ISSN: 2041-1723. DOI: 10.1038/s41467-021-22195-z.

BIBLIOGRAPHY

- Liu, Z. et al. (June 2016). “DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations”. en. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Las Vegas, NV, USA: IEEE, pp. 1096–1104. ISBN: 978-1-4673-8851-1. DOI: 10.1109/CVPR.2016.124.
- Tzourio-Mazoyer, N. et al. (Feb. 2020). *Automated Anatomical Labeling (AAL1) atlas*. DOI: 10.25493/WJHA-2NY.
- Wang, L. et al. (Oct. 2015). “Probabilistic Maps of Visual Topography in Human Cortex”. en. In: *Cerebral Cortex* 25.10, pp. 3911–3931. ISSN: 1047-3211, 1460-2199. DOI: 10.1093/cercor/bhu277.

Appendix A

SUPPLEMENTARY MATERIAL - CHAPTER 2

Supplementary Information




Figure A.1: Online experiment design. Participants performed a rating task under four instructed goal-contexts (tropical island vacation, ski trip, job interview and summer wedding). Participants indicated how much they would like to wear the item shown given the goal-context on a scale from 1 through 9 (1; don't want to purchase at all, 9; want to purchase a lot). This rating was reported using the 1 through 9 keys at the top of a participants keyboard. Image order was fully randomized and trials under a certain goal-context were presented in a blocked manner (20 trials before switching). Trials were self-paced and participants moved on to the next stimulus-context combination when a rating was entered. Images shown are placeholders, actual stimuli were from Liu et al., 2016.




Figure A.2: **A)** Distribution of ratings across participants and contexts for the MRI sample (n=35). **B)** Ratings split by gender for the MRI sample. **C)** Ratings distribution for the online sample. **D)** Ratings split by gender for the online sample. **E)** Attribute ratings split by gender. Note this includes only the MRI sample since no attribute ratings were obtained for the online sample. **F)** Average pairwise correlations for all attributes rated by the MRI sample. **G)** Spearman rank correlation. Outliers ($z > 3$) removed for plotting only.

Context	Attribute	Estimate	Standard error	T-value	P-value
Island	warmth	-0.398	0.011	-36.4	1.115605e-234
Island	comfort	0.167	0.019	8.7	4.082655e-18
Island	formalness	-0.236	0.012	-19.7	1.065373e-80
Island	festiveness	0.146	0.016	9.2	5.781847e-20
Ski	warmth	0.486	0.010	48.1	0.000000e+00
Ski	comfort	0.3119	0.019	16.5	4.469856e-58
Ski	formalness	-0.098	0.013	-7.5	1.092899e-13
Ski	festiveness	-0.057	0.016	-3.4	5.804354e-04
Job	warmth	0.090	0.014	6.4	1.559250e-10
Job	comfort	-0.380	0.019	-20.2	4.266328e-84
Job	formalness	0.561	0.008	72.9	0.000000e+00
Job	festiveness	-0.209	0.016	-12.8	2.047208e-36

Table A.1: Attribute main effects across subjects within each goal-context.

Raw attribute	Kruskal - Wallis H	P-value	Group
Formalness	70.7	4.45E-16	fMRI
Warmth	84.9	3.72E-19	fMRI
Festiveness	4.6	0.10	fMRI
Comfort	15.6	0.0004	fMRI
Formalness	176.70	4.35E-39	Online
Warmth	212.0	9.18E-47	Online
Festiveness	6.8	0.03	Online
Comfort	8.1	0.02	Online

Table A.2: Statistics of Kruskal - Wallis tests for coefficients fit using the Logistic model.




Figure A.3: **A)** Leave-one-out model fits as in Figure 3.1D but for only the Online experiment. Generalization across individuals, with models trained on N-1 and tested on left out 1 subject. **B)** Across subject generalization performance for fMRI subjects. **C)** Leave-one-out generalization results as in Figure 3.1E. Model trained on all subjects from one gender and tested on each of the opposite gendered subjects one-by-one. Results shown for Online experiment only. **D)** Across-gender generalization results for fMRI experiment group.




Figure A.4: ROI locations as used in generalization analysis. The two visual ROIs are derived from Wang et al., 2015 and are shown in green and brown. Frontal ROIs are from Tzourio-Mazoyer et al., 2020 shown in blue and orange. For more details see Methods.




Figure A.5: Proportion of significant voxels along posterior-to-anterior ROIs in the dorsal visual stream.




Figure A.6: **A)** Generalization decoding accuracies for festiveness and comfort in ventral and dorsal visual ROIs. We observe good generalization across attributes and ROIs. **B)** Generalization accuracies for frontal ROIs for festiveness and comfort. We see only festiveness can be decoded in the training context. Additionally, we observe no significant generalization across ROIs for both festiveness and comfort. **C)** Averaged decoding accuracies for all train-test combinations show a similar pattern as **A)**. We observe no differences between in- and out-of-context decoding accuracies. **D)** Similarly for the comfort attribute, we observe no differences in decoding accuracies between models tested on the training context and those tested on another context.




Figure A.7: **A)** Generalization decoding for integrated value for ventral and dorsal visual ROIs. We observe no significant decoding in the training context for the ventral visual ROI, indicating there is no information about integrated value held in this region. Additionally, while we do observe significant decoding in the training context for the dorsal visual ROI, this decoding does not generalize the test contexts. **B)** Within subject decoding accuracies of integrated value for each of the ROIs along the ventral visual pathway. For these decoding analyses all trials across goal-contexts were used (and train-test splits were set up, see Methods). Box-and-whiskers show across subject medians and quartiles similar to previous figures. Significance of these group-level accuracies is computed through comparison to a permuted null and indicated as follows: *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$. **C)** Similarly, for the ROIs along the dorsal visual pathway we show the within-subject decoding accuracies and group-level statistics for the integrated value ratings across the whole task. **D)** Encoding model of integrated value using the behavioral model estimates from the linear model (see Methods) instead of the participant rating. Max voxel in LPFC: [25, 81, 43], $z=6.5$

Appendix B

SUPPLEMENTARY MATERIAL - CHAPTER 3

Supplementary Information

Patient ID	Sex	Age	Ethnicity	Race
P97CS	M	32	Non-Hispanic	Black
P94CS	F	26	Hispanic	White
P92CS	F	30	Non-Hispanic	White
P90CS	M	32	Non-Hispanic	Black
P87CS	F	26	Hispanic	White
P86CS	F	39	Hispanic	White
P84CS	M	63	Non-Hispanic	White
P81CS	F	28	Hispanic	Other
P98CS	F	30	Hispanic	White
P99CS	F	39	Hispanic	White
P100CS	M	25	Hispanic	White
P103CS	M	53	Non-Hispanic	White
P102CS	F	55	Non-Hispanic	White
P106CS	M	37	Non-Hispanic	White
P108CS	F	44	Non-Hispanic	White
P110CS	M	38	Hispanic	White

Table B.1: Patient demographics.

	Action Task	Stimulus Task
Choice Kernel	7426.4	4541.0
RL 2 parameter	7464.5	4429.5
RL 4 parameter	6757.6	4444.3

Table B.2: Total BIC scores across subjects.




Figure B.1: Locations of ACC, vmPFC, preSMA and amygdala electrodes across patients.




Figure B.2: **A)** Individual session's model performance for the Action Task. Distance in BIC between the best (0 distance) and other models. Larger BIC distance indicates worse model performance. **B)** Similarly, Stimulus Task individual session's model performance. **C), D), G), H)** 4 param RL model parameter recovery in the stimulus task. **E), F), I), J)** 4 param RL model parameter recovery in action task.




Figure B.3: Consistency in coding direction of the largest beta. Neurons are selected to only include those that significantly ($p > 0.05$) code for both stimulus value in the stimulus task and action value in the action task. Model fitted beta coefficients with the largest effect (i.e., the strongest stimulus- or action-value code)




Figure B.4: **A)** Proportion of significant neurons encoding the chosen action during the action task across four regions of interest: ventromedial prefrontal cortex (vmPFC), pre-supplementary motor area (preSMA), anterior cingulate cortex (ACC), and Amygdala. Color conventions match main manuscript **B)** Time-resolved proportion of significant units encoding the chosen action relative to response execution (dashed vertical line at 0 ms) in the action task. Solid lines denote the proportion for each brain region computed using a sliding window approach (see Methods). **C)** Proportion of significant neurons encoding the chosen stimulus during the stimulus task. Conventions follow panel A. **D)** Time-resolved proportion of significant units encoding the chosen stimulus relative to the response during the stimulus task. Conventions follow panel B. **E)** Proportion of significant neurons encoding the chosen action during the stimulus task. Conventions follow panel A. **F)** Time-resolved proportion of significant units encoding the chosen action relative to the response during the stimulus task. Conventions follow panel B