

Title

Hypothesis Spaces: How Structural Beliefs Shape Uncertainty in the Anxious Mind

Corresponding Authors

Julia Benita Senfft¹
julia.senfft@uni-bielefeld.de

Ondrej Zika^{1,2,3}
ondrej.zika@ucd.ie

Author affiliations

1 Faculty of Psychology and Sports Science, Department of Psychology, AE14: Biological and Cognitive Neurosciences, Bielefeld University, Bielefeld, Germany

2 School of Psychology, University College Dublin, Dublin, Ireland

3 Centre for Psychiatry Research, Department of Clinical Neuroscience, Karolinska Institutet, & Stockholm Health Care Services, Region Stockholm, Stockholm, Sweden

Keywords

uncertainty; structural inference; anxiety; hypothesis spaces; structure learning

Abstract

While uncertainty is ubiquitous, not all uncertainty is the same. In this review, we dissociate between uncertainty in aspects of known problem spaces (e.g. how long are hospital waiting times; “parameter uncertainty”) and uncertainty arising from unknown structure or even the nature of a problem (e.g. how do I get medical help in a new country; “structural uncertainty”). We suggest that much everyday uncertainty is structural: uncertainty about which model of the world applies in the first place. We bring research on uncertainty and structure learning together through the concept of hypothesis spaces - sets of candidate generative structures that an individual considers plausible, together with distribution of beliefs over their likelihood. We propose that properties of these spaces, including their size, complexity and entropy, provide a systematic framework for understanding structural uncertainty and may help explain characteristic features of anxious thought such as worry, overthinking, catastrophizing and overgeneralization. This perspective also offers a way to refine broad constructs such as intolerance of uncertainty by linking them to specific representational and inferential processes and their properties. We suggest that revisiting the role of uncertainty in anxiety from the viewpoint of structural inference holds the key for mechanistic understanding of clinical observations and closer alignment with fundamental research in cognitive neuroscience.

Introduction

Successful learning, decision making and planning require us to internally represent structures of the external world, such as sequences of actions that will lead us to a specific goal. For example, when needing to visit a doctor after moving to a new country, we first have to register with the insurance provider and receive an insurance card, before booking an appointment at the doctor's office, and finally, attending the appointment. Building adequate internal representations which reliably predict and map onto external events is critical for abstract thought and flexible behaviour.

However, we often do not know which internal representation best captures the true structure of the external world. Instead of a single option we must consider a number of hypothetical scenarios to make sense of limited knowledge. This generates structural uncertainty – uncertainty over structural rules (see Fig. 2 A-C for an example). This type of uncertainty differs from parameter uncertainty, which refers to an uncertain aspect of a known world model.

Here we review the recent literature on uncertainty and structural inference and argue that a large proportion of everyday uncertainty is structural rather than parametric: the challenge is often more about inferring which generative structure applies, than managing unpredictability within known structures (Fig. 1). Arguably, the first is also more anxiety-provoking. We bring uncertainty and structural inference together through the concept of hypothesis spaces – probabilistic beliefs over candidate structures, and consider how their properties may constitute a useful framework for conceptualizing structural inference, as well as its implications for anxiety. We suggest that viewing uncertainty from the perspective of probabilistic hypothesis spaces can provide a mechanistic account of distinctions already recognized phenomenologically in clinical settings, thereby bridging cognitive neuroscience research with anxious thought in clinical settings as well as daily life.




Figure 1 Schematic depiction of parameter and structural uncertainty. Blue points represent abstract events, black arrows depict directed causal relationships. Superimposed opaque blue points represent uncertainty about the value or outcome of that node, i.e. structural uncertainty (bottom row). Multiple chains in the same area represent the agent representing multiple causal explanations at once (structural uncertainty). Both types of uncertainty can be present (bottom-right) or absent (top left).

Not all uncertainty is the same

Acting adaptively requires dissociating signals from noise(s). A surprising event, for example, may reflect a genuine change in the environment or merely stochastic variability in observations. While we should update our beliefs when an outcome arises due to a signal, we should disregard events due to chance. This problem maps onto two complementary sources of uncertainty: aleatoric uncertainty arising from stochastic variability in observations, and epistemic uncertainty arising from lack of information (not knowing). These dimensions are commonly operationalized by varying two properties of the process generating outcomes: its stochasticity, the spread of outcomes around a fixed contingency, and its volatility, the rate at which that contingency changes [1–4]. A related line of research has focused on analyzing behaviour in temporally structured environments [5–7]. In such switching environments, the learner must dissociate whether a surprising event occurred due to chance (oddball) or a context shift (increase in epistemic uncertainty). While this approach introduces hidden structure into the task, the main focus lies in estimating parameters of adaptation to a single hypothesized structure (e.g. the rate of state switch, known as a hazard rate). Together, these approaches are helpful in understanding how individuals infer hidden parameters of the environment where a generative structure is known.

Beliefs about structure determine uncertainty estimates

Critically, while signal and noise are given by the environment, the agent's beliefs may not always map on the ground truth: a meaningful event can be dismissed as noise or, vice versa, noise can be treated as signal, both leading to miscalibrated beliefs. In many daily life and experimental contexts, evidence is compatible with multiple causal explanations (i.e. the data-generating process itself is uncertain). As illustrated in Fig 2C, each possible pathway embodies a different assumption about how the system operates. Whether observed variability is treated as noise (aleatoric) or signal (epistemic) depends on the model used to explain it [8]. The same surprising observation may signal stochastic noise under one model but a change in the underlying state under another. Parameter estimates are consequently conditional on structural assumptions. Whether an observation warrants inferring environmental change (and adapting action accordingly) is thus conditional on the generative model an agent considers plausible in the first place [9]. Because observations can only be interpreted relative to an inferred generative structure, structural beliefs sit at the core of most cognitive problems. Yet, while parameter uncertainty has received ample attention in the literature, investigations of uncertainty over possible structures are less common.

Structural Inference

Cumulative evidence suggests that categorization is 'baked' into the brain [10]: we infer task rules by sequentially testing hypotheses [11] to construct future behavior compositionally from

existing building-blocks [12], implying that multiple candidate generative structures are maintained and evaluated over time.

Critically, structure learning is fundamentally under-constrained by environmental evidence alone: indefinitely many structures remain compatible with any finite set of observations. Recent work thus reframes structure learning as a constraint-satisfaction problem: plausible structures emerge from the joint satisfaction of environmental and prior constraints [13]. Which structures are considered plausible at all is conditional on structural expectations driven by priors [14,15]. Demonstrably, people bring structured expectations even to simple learning tasks [16].

Cognitive neuroscience has formalized a number of mechanisms to capture structure learning, including Bayesian latent cause models [17,18,14] or model-based learning within the reinforcement learning framework [19]. Recent accounts employed Bayesian models that capture how individuals learn temporal structures in terms of transition matrices over latent states, corresponding to inferring the hyperparameters of a generative process [20]. This statistical learning can be flexibly adjusted based on uncertainty [21], reflecting the active adaptation of learning strategies conditional to the degree and complexity of structural information available in uncertain sequences.

Furthermore, neural studies provide converging evidence for the representational capacities required by structure learning. Research on cognitive maps, latent-state representations, and belief states suggests that the brain maintains representations of latent task structures extending beyond immediate observations [22–24]. The orbitofrontal cortex (OFC), ventromedial prefrontal cortex (vmPFC), and the hippocampal formation (HPC) encode latent task structure as abstract state spaces that support flexible behaviour [25–27]. Contextual information shapes task-specific representational geometry: frontoparietal cortex represents distinct contexts on orthogonal manifolds, while states linked by a shared structure become compressed along structure-relevant dimensions, keeping the operative structure separable from its alternatives [28,29]. Compositional representations of abstract task structures function as building blocks that can be constructed, recombined, and reused through compositional learning and generative replay, enabling flexible generalization [12,30–35].

Reformulating structural uncertainty as inference over branching hypothesis spaces

Having seen that humans represent and use structural knowledge in the form of cognitive maps or contextual beliefs, how do they manage competing or simultaneously plausible hypotheses? We outline the idea of hypothesis spaces: sets of candidate generative structures a learner currently considers plausible. Coming from the angle of psychopathology, we highlight the subjectivity of such spaces, as they do not constitute normative probabilistic beliefs (as belief states; [36,37]). Hypotheses may differ qualitatively in their internal organization, ranging from simple linear causal chains to hierarchically nested and branching causal processes (Fig. 2 A-C). Structure learning involves maintaining and updating beliefs over these competing generative structures. While established accounts of structure learning, such as cognitive map learning,

can update a representation of environmental structure [22], hypothesis-space inference can further represent uncertainty over alternative structural explanations, including alternative explanations for structural change. In particular, the learner maintains a probability distribution over possible structures [38], similarly to Markovian belief states [37,39] or probabilistic successor representations [18].




Figure 2 Hypothesis spaces and their properties, Panels A-C depict a hypothetical scenario of taking a real-world action, planning a doctor visit. (A) lists the possible steps that might be taken; (B) presents three possible chains of actions, i.e. (hypotheses) varying in complexity (from top to bottom). (C) Then presents a hypothetical set of hypotheses one may hold. Panels D-F exemplify different properties of hypothesis spaces. (D) presents three scenarios, an individual hypothesis space with just one option under consideration (top/green), a hypothesis space with three scenarios under consideration (middle/blue), and a full hypothesis space of seven scenarios (purple/bottom). Structural uncertainty plot (vertical bars) symbolically highlight how the amount of structural uncertainty increases with the size of a hypothesis space. (E) highlights how the same-sized hypothesis spaces can differ in their weighing in terms of their complexity, top row depicting an individual assigning more probability to low complexity, and lower row showing someone assigning more probabilities to more complex hypotheses. (F) depicts different distributions of beliefs (i.e. probabilities) across hypotheses, so called entropy.

Structural inference and anxiety

Anxiety has been consistently linked to increased sensitivity and aversion to uncertainty [40–43]. Common cognitive symptoms and behaviours in anxiety such as excessive worry, overthinking, overgeneralization or catastrophizing have close links to uncertainty [44]. For

example, the object of pathological worry is often *hypothetical* rather than real. Similarly, catastrophizing is about overweighing the probability of *unlikely* scenarios. In experimental tasks involving learning under uncertainty, research has found that anxious individuals select uncertainty-reducing choices over reward maximization [45]. Trait anxiety has also been linked to misestimation of aleatoric and epistemic uncertainty, including altered adaptation to volatility and changepoints [46–49], a tendency to attribute stochasticity to volatility [50], and elevated volatility priors [51–53]. However, replication has proven inconsistent [4,6,54]. While these results broadly point to alternations in learning under uncertainty in anxiety, the evidence remains mixed and inconsistent. Interestingly, research has largely operationalized uncertainty as inference over parameters, assuming known generative structure of the task, or inducing structure without accounting for it in the analysis [55].

Recent work has identified links between anxiety and aberrant inference of latent structures [7,56,57]. Anxiety in non-clinical populations has been linked to increased propensity to infer and switch between hidden states [7], altered model-based control [57], and rigid choice behavior [58]. Furthermore, anxiety has been linked to differences in acquired higher-order representations determining which behavioural strategies are considered [56]. Collectively, these empirical results hint that anxiety impacts structural inference (or vice versa). This has led to a number of theoretical proposals reformulating anxiety in terms of altered internal representation of structures [59,60]. These accounts closely align with clinical accounts of anxiety. Psychological treatments such as Acceptance and Commitment Therapy (ACT) or Cognitive & Behavioural Therapy (CBT) emphasize understanding of structural beliefs and (in some cases) their modification [61–63].

We suggest that revisiting the role of uncertainty in anxiety from the viewpoint of structure learning holds the key for a closer alignment of clinical observations with fundamental research in cognitive neuroscience. In the following section, we suggest how specific biases at the level of structural inference readily account for common cognitive symptoms of anxiety including worry, overthinking or overgeneralization. Furthermore, we suggest that umbrella terms such as “intolerance of uncertainty” can be reformulated using concepts from structure learning such as hypothesis spaces.

Hypothesis spaces in anxiety

Anxious thought manifests as worry, rumination, overthinking or catastrophizing. These symptoms may all reflect distortions in structure and weighing of an individual's hypothesis space. For example, worry and overthinking likely relate to persistent generation and retention of alternative explanations for what might happen and why. In catastrophizing, even low-probability possibilities may remain under consideration rather than being discarded. Paradoxically, attempts to anticipate every possible contingency may expand rather than resolve uncertainty: as the number of candidate explanations increases, belief becomes distributed across competing models of the world instead of concentrating on a small set of plausible structures.

Anxious thought may thus become self-reinforcing, with structural uncertainty promoting further hypothesis generation, which in turn maintains structural uncertainty. This perspective suggests that anxious cognition might relate to properties of the underlying subjective hypothesis space: namely its size, complexity and entropy.

Because indefinitely many structures are compatible with finite observations, structural inference is governed by a trade-off between explanatory flexibility and computational tractability. Maintaining more candidate hypotheses (i.e. greater size of a hypothesis space) increases explanatory flexibility by allowing unexpected observations to be accommodated by at least one hypothesis, thereby reducing expected surprise. Yet, considering too many hypotheses simultaneously is cognitively demanding and it can lead to rumination and worry, factors linked to elevated stress levels [64]. While no study has explicitly looked at the size of hypothesis spaces and anxious symptoms, trait anxiety has been linked to increased propensity to represent multiple hidden states as potential causes of threat, producing steeper expectation switches following reversals while discounting isolated oddballs [7]. This is further supported by evidence of choice overload and indecisiveness in anxiety [65]. In Fig 2D we abstract the doctor-visit example to illustrate how the size of hypothesis space can lead to different levels of structural uncertainty.

Second, hypotheses differ in their internal causal organization and its complexity (Fig 2E). Accessing care might be represented as a simple chain of action-outcomes, or as a branching structure in which referral depends on insurance status, insurance status on residency, and each step opens further contingencies. Inference ordinarily favours parsimonious models, penalizing the excess flexibility of richer models [66]. However, individuals may vary in their tendency to infer and weigh simpler versus more complex models, even if this is maladaptive. Shifting this prior towards richer models introduces additional latent causes and branching contingencies.

More complex as well as larger hypothesis spaces both naturally lead to increased probability of occurrence of a specific event. More complex candidate hypotheses contain more nodes at which a given feature can occur, so a larger proportion of the maintained space contains any particular feature somewhere in its causal sequence. Thus, a single aversive event is objectively more likely. This has direct consequences for aspects of anxious thought such as overestimation of aversive probabilities or catastrophizing. Exaggerated beliefs about aversive likelihoods may arise from different size or complexity of hypothesis spaces in anxiety. An alternative possibility is that anxious individuals may also specifically construct hypothesis spaces that are more likely to contain an aversive event (Fig 3).




Figure 3 Likelihood of aversive events in differently sized and complex hypothesis spaces. Each panel shows a hypothesis space over the same domain. Each hypothesis ($H_1 \dots H_n$) is a directed path through events or actions (circles); magenta nodes are aversive. The left panel holds four simple, i.e. largely linear hypotheses (14 nodes, mean 3.5 per hypothesis); the right panel holds seven hypotheses, three of them longer and more branched (35 nodes, mean 5.0). The upper summary row in each panel gives the proportion of *nodes* that are aversive: two of fourteen on the left, five of thirty-five in the right, i.e. one in seven in both hypothesis spaces. The lower row shows the proportion of hypotheses containing the aversive event, and this is not identical: it rises from two of four to five of seven. In a richer hypothesis space (right), a fixed aversive rate is distributed across a larger fraction of the entertained possibilities, so aversive events become reachable along more of the paths the agent is holding open. Two consequences follow. Elevated threat expectations require no bias in event probabilities: an agent with perfectly calibrated parameter estimates can arrive here purely through the shape of its hypothesis space. Second, since widening and elaborating the space is the ordinary response to unresolved uncertainty, the increase is self-generated by the very inference intended to resolve it.

A third property of hypothesis spaces is how belief likelihoods are concentrated across hypotheses, i.e. belief entropy (Fig 2F). High entropy spaces are characterized by more distributed beliefs and higher overall uncertainty, while low entropy of beliefs exemplifies commitment to a specific hypothesis. Entropy over hypotheses is directly linked to structural uncertainty, which increases entropy leading to more uncertainty. In turn, entropy acts as a prior, and, as such, it shapes how much evidence is required before beliefs are revised. Retaining several threat structures with belief concentrated on one supports rapid, confident revision when that hypothesis is disconfirmed; the same hypotheses with belief distributed across them support none strongly and sustain parallel evaluation. In contrast to generally smaller hypothesis spaces where some options simply do not exist, entropy accounts for cases where, hypotheses retained at low probability remain recoverable, leaving extinguished threat models available for reinstatement [7,67], whereas probability sustained across many candidates

simultaneously corresponds, phenomenologically, to the proliferation of possibilities characteristic of worry.

Across the three properties, flexibility is purchased at inferential cost. Which candidates are sampled is itself determined by prior evidence, so structural inference requires assessing evidence across the maintained space, and this assessment scales with how many candidates are held, how evenly belief is distributed over them, and how richly each is structured. Under such costs people restrict what they evaluate, sampling a limited consideration set rather than assessing every option [68]. Information therefore acquires value beyond its immediate instrumental utility even where outcomes can no longer be influenced, because it constrains what has to be assessed in the first place. Consistently, trait somatic anxiety is associated with greater information seeking that reveals whether an event will occur despite being unable to influence the outcome [69,70]. Where such information is unavailable, candidate structures must be evaluated internally, as in uncertainty-driven exploration during planning [71], offering a computational description of overthinking.

In summary, we suggest that much of the uncertainty encountered in everyday life is structural: uncertainty about which model of the world applies, rather than only about unknown parameters within a known model. We propose hypothesis spaces as a systematic framework for studying this problem by formalizing the candidate structures an individual considers plausible and the distribution of belief across them. This perspective may help unpack the broad notion of “uncertainty” traditionally captured by intolerance of uncertainty, while providing a closer mapping between experimental findings and everyday anxious thought. In particular, properties such as the size, complexity and entropy of hypothesis spaces offer testable dimensions through which worry, overthinking, catastrophizing and generalization may be examined. Future work should therefore manipulate these properties systematically while independently varying structural and parametric uncertainty. At present, however, the mapping between specific computational mechanisms and particular symptoms remains largely unknown.

Moreover, hypothesis-space inference is unlikely to provide a complete account of anxious cognition: processes such as attention and memory likely influence which hypotheses are generated, maintained or retrieved. We have also largely set valence aside, although our generalization example makes clear that whether represented outcomes are neutral, rewarding or aversive is likely to be critical. A fuller account of anxious thought will therefore require understanding how structural beliefs interact with valence and other cognitive mechanisms.

Acknowledgments

JBS was supported by Studienstiftung des Deutschen Volkes awarded to Julia Benita Senfft. We acknowledge the support of Tina B. Lonsdorf (Bielefeld University), who provided funding for student research assistance and made relevant resources available to support this project. OZ was supported by an ERC Preparatory Fellowship awarded to O.Z. by Bielefeld University. We would like to thank Dr. Shany Grossman for providing helpful comments on the manuscript.

References

- [1] Piray P, Daw ND. Computational processes of simultaneous learning of stochasticity and volatility in humans. *Nat Commun* 2024;15:9073. <https://doi.org/10.1038/s41467-024-53459-z>.
- [2] Fang X, Piray P. Inferring the causes of noise from binary outcomes: A normative theory of learning under uncertainty 2026:2026.03.01.708925. <https://doi.org/10.64898/2026.03.01.708925>.
- [3] Pulcu E, Browning M. Humans adapt rationally to approximate estimates of uncertainty. *eLife* 2025;14:RP103734. <https://doi.org/10.7554/eLife.103734>.
- [4] McCoy B, Lawson RP. The relationship between anxious traits and learning about changes in stochasticity and volatility. *PLOS Comput Biol* 2025;21:e1013646. <https://doi.org/10.1371/journal.pcbi.1013646>.
- [5] Nassar MR, Bruckner R, Frank MJ. Statistical context dictates the relationship between feedback-related EEG signals and learning. *eLife* 2019;8:e46975. <https://doi.org/10.7554/eLife.46975>.
- [6] Satti MH, Wille K, Nassar MR, et al. Absence of Systematic Effects of Internalizing Psychopathology on Learning Under Uncertainty 2025:2025.05.12.653409. <https://doi.org/10.1101/2025.05.12.653409>.
- [7] Zika O, Wiech K, Reinecke A, et al. Trait anxiety is associated with hidden state inference during aversive reversal learning. *Nat Commun* 2023;14:4203. <https://doi.org/10.1038/s41467-023-39825-3>.
- [8] Hüllermeier E, Waegeman W. Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods. *Mach Learn* 2021;110:457–506. <https://doi.org/10.1007/s10994-021-05946-3>.
- [9] Weiss A, Chambon V, Lee JK, et al. Interacting with volatile environments stabilizes hidden-state inference and its brain signatures. *Nat Commun* 2021;12:2228. <https://doi.org/10.1038/s41467-021-22396-6>.
- [10] Barrett LF, Miller EK. Categorization is ‘baked’ into the brain. *Nat Rev Neurosci* 2026;27:435–56. <https://doi.org/10.1038/s41583-026-01036-2>.

- [11] Bähner F, Popov T, Boehme N, et al. Abstract rule learning promotes cognitive flexibility in complex environments across species. *Nat Commun* 2025;16:5396. <https://doi.org/10.1038/s41467-025-60943-7>.
- [12] Bakermans JJW, Warren J, Whittington JCR, et al. Constructing future behavior in the hippocampal formation through composition and replay. *Nat Neurosci* 2025;28:1061–72. <https://doi.org/10.1038/s41593-025-01908-3>.
- [13] Lundbak Olesen C, Mikuš N, Hansen M, et al. It Is What It Isn't: Introducing a Constraint-Based Approach to Structure Learning. *Entropy* 2026;28:534. <https://doi.org/10.3390/e28050534>.
- [14] Goddu MK, Gopnik A. The development of human causal learning and reasoning. *Nat Rev Psychol* 2024;3:319–39. <https://doi.org/10.1038/s44159-024-00300-5>.
- [15] Katayama R, Yoshida W, Amemori K, et al. Neural mechanisms of confidence propagation in hierarchical partially observable decision-making. *iScience* 2025;28. <https://doi.org/10.1016/j.isci.2025.112782>.
- [16] Hedrich NL, Schulz E, Hall-McMaster S, et al. An inductive bias for slowly changing features in human reinforcement learning. *PLOS Comput Biol* 2024;20:e1012568. <https://doi.org/10.1371/journal.pcbi.1012568>.
- [17] Gershman SJ, Blei DM, Niv Y. Context, learning, and extinction. *Psychol Rev* 2010;117:197–209. <https://doi.org/10.1037/a0017808>.
- [18] Geerts JP, Gershman SJ, Burgess N, et al. A probabilistic successor representation for context-dependent learning. *Psychol Rev* 2024;131:578–97. <https://doi.org/10.1037/rev0000414>.
- [19] Daw ND, Gershman SJ, Seymour B, et al. Model-Based Influences on Humans' Choices and Striatal Prediction Errors. *Neuron* 2011;69:1204–15. <https://doi.org/10.1016/j.neuron.2011.02.027>.
- [20] Razmi N, Dai XC, Bakst L, et al. A normative account of human temporal structure learning 2026:2026.01.10.698785. <https://doi.org/10.64898/2026.01.10.698785>.
- [21] Daikoku T, Yumoto M. Order of statistical learning depends on perceptive uncertainty. *Curr Res Neurobiol* 2023;4:100080. <https://doi.org/10.1016/j.crneur.2023.100080>.
- [22] Garvert MM, Saanum T, Schulz E, et al. Hippocampal spatio-predictive cognitive maps adaptively guide reward generalization. *Nat Neurosci* 2023;26:615–26. <https://doi.org/10.1038/s41593-023-01283-x>.
- [23] Moneta N, Grossman S, Schuck NW. Representational spaces in orbitofrontal and ventromedial prefrontal cortex: task states, values, and beyond. *Trends Neurosci* 2024;47:1055–69. <https://doi.org/10.1016/j.tins.2024.10.005>.

- [24] Whittington JCR, Muller TH, Mark S, et al. The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. *Cell* 2020;183:1249-1263.e23. <https://doi.org/10.1016/j.cell.2020.10.024>.
- [25] Niv Y. Learning task-state representations. *Nat Neurosci* 2019;22:1544–53. <https://doi.org/10.1038/s41593-019-0470-8>.
- [26] Schuck NW, Cai MB, Wilson RC, et al. Human Orbitofrontal Cortex Represents a Cognitive Map of State Space. *Neuron* 2016;91:1402–12. <https://doi.org/10.1016/j.neuron.2016.08.019>.
- [27] Moneta N, Garvert MM, Heekeren HR, et al. Task state representations in vmPFC mediate relevant and irrelevant value signals and their behavioral influence. *Nat Commun* 2023;14:3156. <https://doi.org/10.1038/s41467-023-38709-w>.
- [28] Muhle-Karbe PS, Sheahan H, Pezzulo G, et al. Goal-seeking compresses neural codes for space in the human hippocampus and orbitofrontal cortex. *Neuron* 2023;111:3885-3899.e6. <https://doi.org/10.1016/j.neuron.2023.08.021>.
- [29] Flesch T, Juechems K, Dumbalska T, et al. Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron* 2022;110:1258-1270.e11. <https://doi.org/10.1016/j.neuron.2022.01.005>.
- [30] He L, Wang X, Zhang J, et al. Human hippocampal ripples coordinate planning sequences and compositional representations in neocortex. *Nat Neurosci* 2026;29:1711–21. <https://doi.org/10.1038/s41593-026-02291-3>.
- [31] Luettgau L, Chen N, Erdmann T, et al. A neural mechanism for compositional generalization of structure in humans. *eLife* 2025;14. <https://doi.org/10.7554/eLife.107162.1>.
- [32] Nagy DG, Shen T, Zhou H, et al. Analogy making as amortised model construction 2025. <https://doi.org/10.48550/arXiv.2507.16511>.
- [33] Piray P, Daw ND. Reconciling flexibility and efficiency: medial entorhinal cortex represents a compositional cognitive map. *Nat Commun* 2025;16:7444. <https://doi.org/10.1038/s41467-025-62733-7>.
- [34] Schwartenbeck P, Baram A, Liu Y, et al. Generative replay underlies compositional inference in the hippocampal-prefrontal circuit. *Cell* 2023;186:4885-4897.e14. <https://doi.org/10.1016/j.cell.2023.09.004>.
- [35] Tsvividis P, Loula J, Burga J, et al. Human-level learning of complex novel tasks as theory-based modelling, exploration and planning. *Philos Trans R Soc Math Phys Eng Sci* 2026;384:20240529. <https://doi.org/10.1098/rsta.2024.0529>.
- [36] Bruckner R, Heekeren HR, Nassar MR. Understanding learning through uncertainty and bias. *Commun Psychol* 2025;3:24. <https://doi.org/10.1038/s44271-025-00203-y>.

- [37] Dayan P, Daw ND. Decision theory, reinforcement learning, and the brain. *Cogn Affect Behav Neurosci* 2008;8:429–53. <https://doi.org/10.3758/CABN.8.4.429>.
- [38] Wu CM, Meder B, Schulz E. Unifying Principles of Generalization: Past, Present, and Future. *Annu Rev Psychol* 2025;76:275–302. <https://doi.org/10.1146/annurev-psych-021524-110810>.
- [39] Muller TH, Mars RB, Behrens TE, et al. Control of entropy in neural models of environmental state. *eLife* 2019;8:e39404. <https://doi.org/10.7554/eLife.39404>.
- [40] Brown VM, Price R, Dombrovski AY. Anxiety as a disorder of uncertainty: implications for understanding maladaptive anxiety, anxious avoidance, and exposure therapy. *Cogn Affect Behav Neurosci* 2023;23:844–68. <https://doi.org/10.3758/s13415-023-01080-w>.
- [41] Grupe DW, Nitschke JB. Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nat Rev Neurosci* 2013;14:488–501. <https://doi.org/10.1038/nrn3524>.
- [42] Morriss J. Psychological mechanisms underpinning change in intolerance of uncertainty across anxiety-related disorders: New insights for translational research. *Neurosci Biobehav Rev* 2025;173:106138. <https://doi.org/10.1016/j.neubiorev.2025.106138>.
- [43] Sandhu TR, Xiao B, Lawson RP. Transdiagnostic computations of uncertainty: towards a new lens on intolerance of uncertainty. *Neurosci Biobehav Rev* 2023;148:105123. <https://doi.org/10.1016/j.neubiorev.2023.105123>.
- [44] Verra L, Spitzer B, Schuck NW, et al. Increased generalisation in trait anxiety is driven by aversive value transfer. *Commun Psychol* 2026;4:46. <https://doi.org/10.1038/s44271-026-00415-w>.
- [45] Aberg KC, Toren I, Paz R. A neural and behavioral trade-off between value and uncertainty underlies exploratory decisions in normative anxiety. *Mol Psychiatry* 2022;27:1573–87. <https://doi.org/10.1038/s41380-021-01363-z>.
- [46] Browning M, Behrens TE, Jocham G, et al. Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nat Neurosci* 2015;18:590–6. <https://doi.org/10.1038/nn.3961>.
- [47] Gagne C, Zika O, Dayan P, et al. Impaired adaptation of learning to contingency volatility in internalizing psychopathology. *eLife* 2020;9:e61387. <https://doi.org/10.7554/eLife.61387>.
- [48] Gibbs-Dean T, Katthagen T, Hu R, et al. The joint estimation of uncertainty and its relationship with psychotic-like traits and psychometric schizotypy. *Npj Ment Health Res* 2025;4:40. <https://doi.org/10.1038/s44184-025-00146-6>.
- [49] Pulcu E, Browning M. The Misestimation of Uncertainty in Affective Disorders. *Trends Cogn Sci* 2019;23:865–75. <https://doi.org/10.1016/j.tics.2019.07.007>.

- [50] Yan X, Ebitz RB, Grissom N, et al. Distinct Computational Mechanisms of Uncertainty Processing Explain Opposing Exploratory Behaviors in Anxiety and Apathy. *Biol Psychiatry Cogn Neurosci Neuroimaging* 2025;10:954–63. <https://doi.org/10.1016/j.bpsc.2025.01.005>.
- [51] Duda JM, Cannon TD, Paskewitz S, et al. Anhedonia and Anxious Arousal Are Associated With Distinct Expectations About the Statistics of a Volatile Environment. *Biol Psychiatry Glob Open Sci* 2026;6:100736. <https://doi.org/10.1016/j.bpsgos.2026.100736>.
- [52] Guitart-Masip M, Walsh A, Dayan P, et al. Anxiety associated with perceived uncontrollable stress enhances expectations of environmental volatility and impairs reward learning. *Sci Rep* 2023;13:18451. <https://doi.org/10.1038/s41598-023-45179-z>.
- [53] Hein TP, Gong Z, Ivanova M, et al. Anterior cingulate and medial prefrontal cortex oscillations underlie learning alterations in trait anxiety in humans. *Commun Biol* 2023;6:271. <https://doi.org/10.1038/s42003-023-04628-1>.
- [54] Ting C-C, Palminteri S, Lebreton M, et al. The elusive effects of incidental anxiety on reinforcement-learning. *J Exp Psychol Learn Mem Cogn* 2022;48:619–42. <https://doi.org/10.1037/xlm0001033>.
- [55] Zika O. The relationship between latent state inference and (intolerance of) uncertainty. *Neurosci Biobehav Rev* 2023;152:105321. <https://doi.org/10.1016/j.neubiorev.2023.105321>.
- [56] Schiff H, Sharp PB. Anxiety is Associated With Biases in Task Generalization. *Top Cogn Sci* 2026;n/a:e70051. <https://doi.org/10.1111/tops.70051>.
- [57] Sookud S, Martin I, Gillan CM, et al. Impaired Goal-Directed Planning in Transdiagnostic Compulsivity Is Explained by Uncertainty About Learned Task Structure. *Biol Psychiatry Cogn Neurosci Neuroimaging* 2025;11:357–65. <https://doi.org/10.1016/j.bpsc.2025.10.005>.
- [58] Lamba A, Frank MJ, FeldmanHall O. Keeping an Eye Out for Change: Anxiety Disrupts Adaptive Resolution of Policy Uncertainty. *Biol Psychiatry Cogn Neurosci Neuroimaging* 2024;9:1188–98. <https://doi.org/10.1016/j.bpsc.2024.07.015>.
- [59] Sharp PB, Thomas JG, Yadid N, et al. From task models to world models: A person-specific computational psychology of anxiety 2026.
- [60] Shi Y, Garvert M. Forming, selecting and updating cognitive maps to support flexible behavior: relevance for anxiety 2026. <https://doi.org/10.1098/rstb.2025.0246>.
- [61] Barnes-Holmes Y, Barnes-Holmes D, McHugh L, et al. Relational Frame Theory: Some implications for understanding and treating human psychopathology. *Int J Psychol Psychol Ther* 2004;4:355–76.

- [62] Dozois DJA, Beck AT. Chapter 6 - Cognitive Schemas, Beliefs and Assumptions. In: Dobson KS, Dozois DJA, editors. *Risk Factors Depress.*, San Diego: Elsevier; 2008, p. 119–43. <https://doi.org/10.1016/B978-0-08-045078-0.00006-X>.
- [63] Hayes SC, Luoma JB, Bond FW, et al. Acceptance and Commitment Therapy: Model, processes and outcomes. *Behav Res Ther* 2006;44:1–25. <https://doi.org/10.1016/j.brat.2005.06.006>.
- [64] Zoccola PM, Dickerson SS. Assessing the relationship between rumination and cortisol: A review. *J Psychosom Res* 2012;73:1–9. <https://doi.org/10.1016/j.jpsychores.2012.03.007>.
- [65] Hu X, Turel O, Chen W, et al. The effect of trait-state anxiety on choice overload: the mediating role of choice difficulty. *DECISION* 2023;50:143–52. <https://doi.org/10.1007/s40622-023-00345-0>.
- [66] Piasini E, Liu S, Chaudhari P, et al. How Occam’s razor guides human decision-making 2025:2023.01.10.523479. <https://doi.org/10.1101/2023.01.10.523479>.
- [67] Berwian IM, Hitchcock P, Pisupati S, et al. Using computational models of learning to advance cognitive behavioral therapy. *Commun Psychol* 2025;3:72. <https://doi.org/10.1038/s44271-025-00251-4>.
- [68] Liu S, Gershman SJ. Action subsampling supports policy compression in large action spaces. *PLOS Comput Biol* 2025;21:e1013444. <https://doi.org/10.1371/journal.pcbi.1013444>.
- [69] Bari BA, Gershman SJ. The Value of Non-Instrumental Information in Anxiety: Insights from a Resource-Rational Model of Planning. *Comput Psychiatry* 2025;9. <https://doi.org/10.5334/cpsy.124>.
- [70] Charpentier CJ, Cogliati Dezza I, Vellani V, et al. Anxiety increases information-seeking in response to large changes. *Sci Rep* 2022;12:7385. <https://doi.org/10.1038/s41598-022-10813-9>.
- [71] Fan H, Callaway F, Gershman SJ. Uncertainty-driven exploration during planning. *Decision* 2025;12:334–55. <https://doi.org/10.1037/dec0000267>.