Unifying principles of generalization: past, present, and future

Charley M. Wu,[1,2,3,*] Björn Meder,[4] and Eric Schulz[2]

[1]Human and Machine Cognition Lab, University of Tübingen, 72074 Tübingen, Germany; email: charley.wu@uni-tuebingen.de
[2]Department of Computational Neuroscience, Max Planck Institute of Biological Cybernetics, 72074 Tübingen, Germany
[3]Center for Adaptive Rationality, Max Planck Institute for Human Development, 14195 Berlin, Germany
[4]Institute for Mind, Brain and Behavior, Department of Psychology, Health and Medical University Potsdam, 14471 Potsdam, Germany

This is a preprint. When citing this paper, please use the following:
Wu CM, Meder B, Schulz, E. 2024. Unifying principles of generalization: past, present, and future. Annu. Rev. Psych. 76: Submitted.
DOI: 10.1146/annurev-psych-021524-110810

Annu. Rev. Psych 2024. 76:1–33

https://doi.org/10.1146/annurev-psych-021524-110810

Keywords

generalization, concept learning, function learning, value approximation, reinforcement learning, structure induction

Abstract

Generalization, defined as applying limited experiences to novel situations, represents a cornerstone of human intelligence. Our review traces the evolution and continuity of psychological theories of generalization, from origins in concept learning (categorizing stimuli) and function learning (learning continuous input-output relationships), to domains such as reinforcement learning and latent structure learning. Historically, there have been fierce debates between rule-based mechanisms—using explicit hypotheses about environmental structure—and similarity-based mechanisms—leveraging comparisons to prior instances. Each approach has unique advantages: rules support rapid knowledge transfer, while similarity is computationally simple and flexible. Today, these debates have culminated in the development of hybrid models grounded in Bayesian principles, effectively marrying the precision of rules with the flexibility of similarity. The ongoing success of hybrid models not only bridges past dichotomies but also underscores the importance of integrating both rules and similarity for a comprehensive understanding of human generalization.

Contents

1. INTRODUCTION

Generalization: The process of applying previously acquired knowledge to new, unfamiliar situations

Concept learning: Learning to apply discrete category labels to objects or events.

Function learning: Learning to understand and predict the continuous relationship between input and output variables.

Rules: Explicit hypotheses about the structure of the environment that can guide generalization

Similarity: A comparison of new situations to previous experiences, as a basis for generalization

In the psychological landscape of theories and models, the study of how people generalize past experiences to novel situations has occupied a central and domain-bridging role. Given the unending flux and flow of new experiences and novel situations, generalization stands as a testament to the flexibility of human intelligence (Lake et al. 2017; Chollet 2019), and is widely studied in psychology (Shepard 1987; Chater & Vitányi 2003; Tenenbaum & Griffiths 2001; Wu et al. 2018), neuroscience (Taylor et al. 2021; Norbury et al. 2018; Poggio & Bizzi 2004), and machine learning (Zhang et al. 2016; Jäkel et al. 2008a; Geirhos et al. 2018). Here, we bridge traditional psychological theories with modern computational approaches, providing new perspectives for both old problems and enduring challenges. While the computational methods are certainly new, the theoretical underpinnings and core questions are very familiar to psychology and can be traced back to foundational research in concept and function learning.

Over the years, debates about the mechanisms underlying human generalization have spanned multiple domains. Research in concept learning has studied how people generalize learned category labels when asked to classify new instances, for example, identifying the breed of a dog or deciding whether a hotdog is a sandwich. Meanwhile, research in function learning has studied how people generalize by learning the relationship between inputs and outputs, allowing for interpolation or extrapolation beyond observed data, such as predicting how much study time is needed to pass a test or anticipating how much you will enjoy a new menu item at your favorite restaurant. In both domains, theories about the underlying mechanisms of generalization have largely coalesced around two ingredients: extracting regularities of the environment in the form of generic *rules* to apply in novel settings and using *similarity* to compare new situations to previously encountered instances, with the expectation that similar outcomes will result from similar situations.

While fierce historical debates have raged over which ingredient is more central, today these arguments have largely been settled in favor of hybrid models, which have both rule- and similarity-based interpretations and are frequently based on Bayesian principles (Tenenbaum & Griffiths 2001; Lucas et al. 2015). While a duality of interpretations suggests an exchangeability between rule- and similarity-based representations (Goodman et al. 2008), the computations used by hybrid models typically operate over either one or the other—over hypothesized rules or over representations of similarity (Hahn & Ramscar 2001)—each conferring distinct advantages. Rules unlock compositionality and rapid transfer, while similarity is easy to compute and can flexibly capture various relationships in the environment.

In this review, we revisit the distinction between similarity- and rule-based mechanisms of generalization. Our approach seeks to bridge the past and the present, by emphasizing the continuity of these two mechanisms in theories of generalization. We first explore the development of generalization research in concept learning and function learning, where in each domain there have been converging trajectories toward hybrid models that integrate rule-based and similarity-based approaches. Second, we establish connections between function learning theories and contemporary methods for value generalization in reinforcement learning. This requires integrating a new dimension of uncertainty-directed exploration to guide adaptive learning. Third, we highlight inherent relations between Bayesian concept learning and theories of structure induction, which support generalization by inferring hidden environmental structure. We conclude by proposing new directions for further integrating similarity and rules, combining their relative advantages to unlock faster and more efficient generalization in increasingly complex problems.

2. COMMON PRINCIPLES FOR GENERALIZATION

We first review foundational psychological theories of generalization in *concept learning* and *function learning*, which broadly map onto the distinction between classification and regression problems, as they are commonly referred to in statistics and machine learning. A child distinguishing dogs from cats based on characteristics like barking or meowing is a type of classification problem used in concept learning, while a teacher predicting students' test scores based on study habits and past performance is a type of regression problem used in function learning. Research in these two domains has largely progressed in distinct, parallel tracks. Yet they share a similar historical trajectory of debates about the main mechanisms supporting generalization. The proposed mechanisms can be categorized as *rule-based approaches*, which focus on extracting regularities or generic 'rules' from the environment, and *similarity-based approaches*, which compare new situations to previously encountered instances.

In this section, we examine the evolution of theories about generalization across concept learning and function learning. In both domains, these theories have largely culminated in hybrid models, often using Bayesian principles to unify rule-based and similarity-based approaches. We then show how these hybrid approaches provide the foundations for scaling up to increasingly more complex and real-world problems, drawing connections between theories of function learning and modern approaches to value generalization in reinforcement learning, and from Bayesian concept learning to theories of structure learning.

2.1. Concept Learning

A chief aim of cognitive psychology has been to understand how individuals categorize and differentiate between different elements of the "blooming and buzzing confusion" (James 1890) of the environment (Bruner et al. 1956). Research in the domain of concept learning has long used classification problems with discrete stimuli as a means to study generalization

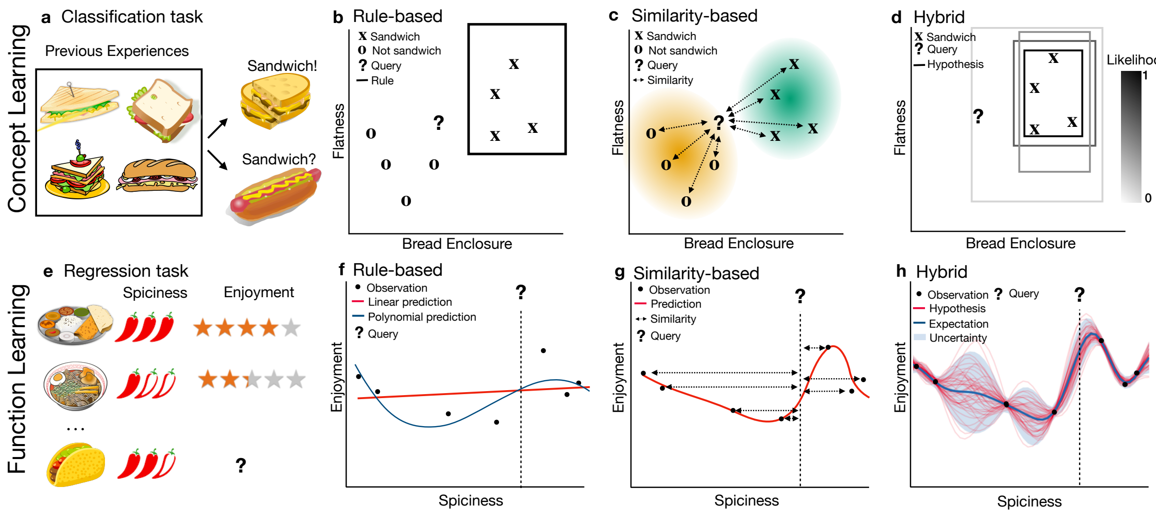

Figure 1

Generalization in Concept Learning and Function Learning. **a**) Concept learning is often studied based on classifying discrete stimuli (e.g., sandwich vs. not sandwich). **b**) Rule-based methods describe explicit category boundaries (rectangle), while **c**) similarity-based methods utilize similarity (arrows) to previous exemplars (data points) or learned prototypes (centroid of colored oval). **d**) Bayesian concept learning (Tenenbaum & Griffiths 2001) provides a hybrid approach by defining a distribution over rules (rectangles), which produce patterns of generalization consistent with similarity-based accounts (Shepard 1987; Tversky 1977). The likelihood favors narrower hypotheses, which is indicated by the shading of lines. **e**) Function learning is studied based on learning a mapping between inputs (e.g., spiciness) and outputs (e.g., enjoyment). **f**) Rule-based methods describe specific parametric families of functions (e.g., linear or polynomial), while **g**) similarity-based methods commonly use Artificial Neural Networks (ANNs) to approximate nonlinear functions, where the influence of each data point is proportional to their similarity (i.e., inverse distance; arrows). **h**) Gaussian Process regression provides a hybrid approach using kernel similarity to describe a distribution over hypothesized functions (red lines), which are summarized in terms of an expectation (blue line) and uncertainty (blue ribbon). Food images are from OpenClipArt under CC0 1.0.

(Rosch 1973; Medin & Schaffer 1978; Smith & Medin 1981; Nosofsky 1986). For instance, learning the category "sandwich" from examples of paninis and subs, and then generalizing confidently when shown a grilled cheese for the first time, but perhaps hesitating when shown a hotdog (Figure 1a). Important debates in this field have concerned which representations are learned and the mechanisms used for generalizing about novel stimuli (Erickson & Kruschke 1998; Ashby & Maddox 2005; Johansen & Palmeri 2002; Bowman et al. 2020). Here, we broadly categorize different influential approaches into rule-based and similarity-based approaches.

Rule-based Concept Learning. One influential class of theories proposed that concepts are defined based on rules that describe the explicit boundaries of category membership (Bruner et al. 1956; Ashby & Gott 1988; Rouder & Ratcliff 2006, rectangle in Fig. 1b). For instance, one might describe the necessary and sufficient features (Smith & Medin 1981) of a sandwich as "food flattened between two pieces of bread", and thus classify any novel food that satisfies this rule as a sandwich. The specificity of rules facilitates rapid generalization, while their compositionality (i.e., the ability to combine multiple rules) makes them infinitely productive (Goodman et al. 2008).

However, for the same reasons, rules can also be inflexible (what about open-faced sandwiches?) and difficult to learn, since infinite productivity also implies an infinite hypothesis space of candidate rules to consider. Even with mechanisms for learning exceptions to rules for added flexibility (Nosofsky et al. 1994), rule-based methods only seem to offer partial explanations of human category learning (Tenenbaum 2000), and perform best when paired together with other learning mechanisms (Ashby et al. 1998; Erickson & Kruschke 1998; Love et al. 2004). Nevertheless, the basic mechanisms of rule-based generalization (i.e., proposing explicit hypotheses) play an important role in modern theories of structure learning (Kemp & Tenenbaum 2008) and program induction (Lake et al. 2015, 2017; Rule et al. 2020; Fränken et al. 2022; Zhao et al. 2023), which use a probabilistic framework to add flexibility to the rigid structure of rules.

Universal Law of Generalization: The probability of a response for one stimulus being generalized to another is a function of the distance between the two stimuli in a psychological space

Bayesian concept learning: A probabilistic approach to concept learning, using a distribution over rule-like hypotheses about based concept boundaries, producing similarity-like generalization patterns.

Bayesian size principle: Smaller, more specific hypotheses are preferred over broader ones, given consistent evidence.

Similarity-based Concept Learning. Another class of theories uses *similarity*-based methods for predicting the category of novel stimuli (Figure 1c). Early theories introduced the notion of a psychological space (Torgerson 1952; Ekman 1954), where stimuli are embedded as geometric coordinates and a measure of distance (e.g., Euclidean distance) serves to represent the (dis-)similarity between stimuli. The most influential example is Shepard's (1987) "Universal Law of Generalization", which used confusability (i.e., the probability of responding to stimulus $\mathbf{x}$ when shown stimulus $\mathbf{x}'$) to construct a psychological space using Multidimensional Scaling (Shepard 1962; Kruskal 1964). Intuitively, stimuli producing similar responses are embedded in similar locations, such that the same unit of distance in any direction corresponds to the same level of generalization. Stimuli located further apart in psychological space are thus less likely to yield the same response, becoming exponentially less likely as their distance increases (Figure 2a).

At the core of Shepard's theory is the assumption that representations about categories correspond to a "consequential region" in psychological space (Figure 2a). Generalization thus arises due to uncertainty about the extent of these regions. As the distance between stimuli $\mathbf{x}$ and $\mathbf{x}'$ increases, they are less likely to belong to the same region and therefore less likely to produce similar outcomes, thus producing a smooth gradient of generalization. Other similarity-based approaches are consistent with this notion of a psychological space, where comparison to either previously encountered exemplars (Medin & Schaffer 1978; Nosofsky 1986), or to a learned prototype (Rosch 1973; Smith & Minda 1998) aggregated over multiple experiences, provides the basis for generalization (arrows in Figure 1c).

Yet the notion of similarity has been famously criticized for being too flexible, with endless and arbitrary ways to define similarity for any pair of stimuli (Goodman 1972; Medin et al. 1993; Hahn & Ramscar 2001). Modern theories address this challenge by providing new approaches for describing the psychological mechanisms people use to construct context-relevant similarity representations (see Radulescu et al. 2021, for a review), forming a rational rather than arbitrary basis for computing similarity. Furthermore, advances in similarity-based approaches to generalization are now able to capture rich relational structure (Wu et al. 2021; Whittington et al. 2020) and represent the temporal dynamics of the environment (Stachenfeld et al. 2017; Garvert et al. 2023).

Hybrid Concept Learning Using Bayesian Principles. Today, the most prolific theories of concept learning are considered hybrids and have a duality of both rule- and similarity-based interpretations (Pettine et al. 2023). One influential example is the *Bayesian concept* *learning* framework (Figure 1d; Tenenbaum & Griffiths 2001), which uses a distribution over hypothesized category boundaries (boxes in Figure 1d) to categorize novel stimuli (Sidebar 2.1).

Bayesian Concept Learning

Bayes' rule is used to describe the posterior probability that each hypothesis h captures category C given a set of positive observations $\mathbf{x}_i \in \mathcal{X}$:

$$p(h|\mathcal{X}) \propto p(h)p(\mathcal{X}|h), \tag{1}$$

This posterior integrates prior beliefs $p(h)$ and the likelihood of the data $p(\mathcal{X}|h)$, where the prior is usually assumed to be uniform, while the likelihood makes use of the *Bayesian size principle* (Tenenbaum & Griffiths 2001) to favor narrower hypotheses that are still consistent with the data:

$$p(\mathcal{X}|h) = \begin{cases} \frac{1}{|h|^n} & \text{if } \mathbf{x}_1, \ldots, \mathbf{x}_n \in h \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

Having defined the posterior probability of a single hypothesis h, the goal is to predict whether a novel stimulus $\mathbf{x}_*$ falls within the same category C as previously observed examples $\mathcal{X}$. Bayesian concept learning defines this probabilistically, by aggregating over all hypotheses h (i.e., category boundaries) consistent with $\mathbf{x}_*$ belonging to C:

$$p(\mathbf{x}_* \in C|\mathcal{X}) = \sum_{h:\mathbf{x}_* \in C} p(h|\mathcal{X}). \tag{3}$$

This represents a sum of posterior probabilities $p(h|\mathcal{X})$ for different hypotheses that encapsulate $\mathbf{x}_*$, where the contribution of each hypothesis is weighted by the size principle (Eq. 2).

A key concept is the *Bayesian size principle* (Tenenbaum 1999, 2000; Tenenbaum & Griffiths 2001), where under the assumption of "strong sampling", greater likelihoods are assigned to narrower hypotheses consistent with the data (darker shading for smaller rectangles in Figure 1d). Strong sampling assumes that rather than being completely random, the data $\mathcal{X}$ are explicitly sampled from positive examples of the category C, as is commonly the case in pedagogical settings (Csibra & Gergely 2009), where a parent or a teacher will provide informative examples of categories, such as "plane", "dog", or "sandwich'. Consequently, the distribution of the observed data $\mathcal{X}$ is expected to reflect the range of the category boundary, thus preferring narrower hypotheses consistent with the data, where the strength of this preference increases with more observations.

Bayesian concept learning thus uses computations over rule-based representations of explicit category boundaries, yet can replicate behavioral patterns of several influential similarity-based theories, such as Shepard's (1987) smooth generalization gradients and is also equivalent to a special case of Tversky's set-theoretic model of similarity (Tversky 1977). And while other hybrid models advocated for a "separate-but-equal" approach (Erickson & Kruschke 1998) by incorporating rules and similarity as separate mechanisms, Bayesian concept learning represents a "unified" approach, where rules and similarity are seen as two sides of the same coin (Pothos 2005; Goodman et al. 2008; Tenenbaum 2000;

Austerweil et al. 2015).

This core Bayesian framework—based on describing a distribution of hypotheses and adapting them to new data—has since proliferated computational theories across a wide range of phenomena, such as causal learning (Meder et al. 2014; Griffiths & Tenenbaum 2005, 2009a), word learning (Xu & Tenenbaum 2007), structure induction (Kemp & Tenenbaum 2008), and the learning of compositional programs (Lake et al. 2015, 2017; Ellis et al. 2023; Fränken et al. 2022; Zhao et al. 2023). A distinct advantage of operating over rule-based representations is the ability to reason compositionally, by syntactically manipulating and combining multiple rules (Piantadosi et al. 2016). Yet given an expressive hypothesis space, exact Bayesian inference is usually intractable, with most approaches relying on sample-based (Tenenbaum & Griffiths 2001; Kemp & Tenenbaum 2008; Ellis et al. 2023) or variational approximations (Dasgupta et al. 2020). Thus, it remains an open question how humans achieve the power and productivity of rule learning, but with limited cognitive resources (van Rooij et al. 2019; Sanborn et al. 2010; Rubino et al. 2023).

2.2. Function Learning

Beyond discrete category membership, generalization has also been studied in the domain of function learning (Figure 1e) based on inferring a continuous relationship between inputs and outputs (Carroll 1963; Brehmer 1974; Lucas et al. 2015; Koh & Meyer 1991; Kalish et al. 2007; Busemeyer et al. 1997). For example, learning how spiciness (input) relates to one's enjoyment of a meal (output), or how the amount of time spent studying (input) predicts test scores (output).

Pioneering research by Carroll (1963) used function learning to show that human generalization goes beyond merely predicting previously observed outcomes, in contrast to the domain of concept learning, where participant responses are typically limited to previously learned category labels, even if the stimuli presented are novel. Rather, Carroll's (1963) work on function learning showed that people can extrapolate beyond their past experiences, generalizing not only to new inputs but also predicting new outputs (e.g., an off-the-charts food experience). While largely operating along a separate research tradition, the domain of function learning is also characterized by a parallel debate between rule- and similarity-based theories, which has culminated in hybrid formalizations (Busemeyer et al. 1997; Kalish et al. 2007; Lucas et al. 2015).

Gaussian process regression: A probabilistic approach to function learning, using a distribution over hypothesized functions, with both rule- and similarity-based interpretations.

Kernel similarity: A similarity metric defined for any pair of stimuli, used in Gaussian Processes.

Rule-based Function Learning. Many early theories of function learning are considered rule-based, by assuming people use a specific parametric model (e.g., a linear or polynomial function), and then learn by optimizing the parameters to best explain the data (Carroll 1963; Brehmer 1976, Figure 1f). In function learning, rules correspond to a hypothesized relationship between variables, much like the law of gravity describes a polynomial relationship between mass and distance, or fitting a linear regression assumes a linear relationship. While rule-based methods can capture the systematicity of human extrapolation patterns (i.e., strong linear assumptions; DeLosh et al. 1997), they lack the flexibility of humans, who can learn to interpolate almost any function with enough training (McDaniel & Busemeyer 2005).

Similarity-based Function Learning. To better account for the flexibility of human generalization, similarity-based models (Figure 1g) of function learning were developed, often using artificial neural networks (ANNs) to encode the generic principle that similar inputs produce similar outputs (McClelland et al. 1986; Busemeyer et al. 1997). The influence of previous observations decreases as a function of distance (arrows in Figure 1g) to a given input, with nearby observations exerting a larger influence. ANNs are universal function approximators (Hornik 1991; Cybenko 1989), and can approach arbitrarily low error in fitting a given function, given sufficient neurons in the hidden layers. But while this flexibility aligns with similar human capabilities in interpolation tasks, ANNs fail to match the specific inductive biases humans exhibit when extrapolating beyond observed data (Schulz et al. 2017; **?**). For instance, humans tend to extrapolate functions with strong linear expectations (Kwantes & Neal 2006; Kalish et al. 2004), a tendency not inherently captured by standard ANNs. This distinction underscores the need to further refine neural network models to more accurately mirror human cognitive processes in both interpolation and extrapolation.

Hybrid Function Learning Using Bayesian Principles. To combine the rule-like systematicity of human extrapolation with the similarity-like flexibility of interpolation, hybrid function learning models were developed. One notable example is Gaussian Process regression (Figure 1h; Rasmussen & Williams 2005), which can account for many empirical patterns of human function learning (Lucas et al. 2015; Griffiths et al. 2008) while using similar Bayesian computations as hybrid models of concept learning (Sidebar 2.2).

Bayesian Function Learning

Gaussian Process regression (Rasmussen & Williams 2005) provides a Bayesian approach to function learning by mapping inputs $\mathcal{X}$ to real-valued outputs y through a distribution over hypothesized functions h. A prior over functions takes the form of a multivariate Gaussian distribution:

$$p(h) \sim \mathcal{GP}\left(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\right), \qquad 4.$$

defined by a prior mean $m(\mathbf{x})$, which is typically set to 0 without loss of generality (Rasmussen & Williams 2005), and a covariance function $k(\mathbf{x}, \mathbf{x}')$, which is defined by a choice of kernel. A common choice is the radial basis function (RBF) kernel:

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(\frac{-||\mathbf{x} - \mathbf{x}'||^2}{2\lambda^2}\right), \qquad 5.$$

capturing the inductive bias that similar inputs are expected to produce similar outputs, with similarity defined as an exponentially decaying function of distance in feature space (Figure 2b). The posterior distribution is then defined by conditioning on observed data $\mathcal{D} = \{\mathcal{X}, \mathbf{y}\}$ of encountered inputs $\mathbf{x}_i \in \mathcal{X}$ and outputs $y_i \in \mathbf{y}$. The posterior is also Gaussian, with predictions for any input $\mathbf{x}_*$ characterized by posterior mean $m(\mathbf{x}_*|\mathcal{D})$ and variance $v(\mathbf{x}_*|\mathcal{D})$:

$$p\left(h(\mathbf{x}_*)\,|\mathcal{D}\right) \sim \mathcal{GP}\left(m(\mathbf{x}_*|\mathcal{D}), v(\mathbf{x}_*|\mathcal{D})\right) \qquad 6.$$

Gaussian Process regression provides a Bayesian approach to function learning (see Schulz et al. 2018a, for a tutorial), based on a distribution over hypothesized functions that explain the data (red lines in Figure 1h). In contrast to Bayesian concept learning, the Gaussian assumptions of Gaussian Process regression provide an analytically tractable posterior distribution, characterized by an expected outcome for any new input (blue line; Figure 1h), but also an uncertainty estimate (blue ribbon; Figure 1h). These analytically tractable computations also have exact equivalencies to artificial neural networks in the limit of an infinite number of hidden units (Neal 1996).

A key ingredient in Gaussian Process regression is the choice of kernel function, which provides an explicit similarity metric between any pair of inputs $\mathbf{x}$ and $\mathbf{x}'$ with desirable mathematical properties (Schölkopf & Smola 2002). This can captures inductive biases present in similarity-based theories (Figure 1g), such as the common RBF kernel (Eq. 5), which assumes similar inputs are likely to produce similar outcomes (Figure 2b). However, there is a rich set of kernel functions to choose from, capturing different forms of inductive biases (Duvenaud et al. 2013). For instance, linear kernels make strong assumptions about linear relationships, periodic kernels encode cyclical trends, and graph kernels capture relational structure between discrete nodes on a graph (e.g., a diffusion kernel; Wu et al. 2021; Kondor & Lafferty 2002, Figure 2c).

Like Bayesian concept learning, Gaussian Process regression has been described as a hybrid model because it has both similarity- and rule-based interpretations (Lucas et al. 2015; Austerweil et al. 2015). The similarity-based interpretation is straightforward, since the kernel explicitly encodes similarity between data points, facilitating computations that operate directly over similarity representations. However, the framework also lends support to two rule-based interpretations.

The first is based on a mathematical property known as Mercer's (1909) theorem, describing how any kernel can be decomposed into a combination of basis functions (Lucas et al. 2015; Austerweil et al. 2015), each corresponding to an abstract rule. Just as any color can be decomposed into red, green, and blue components, the basis functions that collectively constitute a kernel form the rule-like building blocks that allow Gaussian processes to express a potentially unlimited range of functions. Thus, inversely analogous to Bayesian concept learning, Gaussian Processes operate on similarity-based computations but provide equivalent rule-based interpretations.

There is also a second rule-based interpretation, with more direct applications based on the compositionality of Gaussian Process kernels (Schulz et al. 2017; Duvenaud et al. 2013). Multiple kernels can be combined via addition or multiplication operations to produce new kernels. Since each kernel can be seen as providing rule-like biases about the hypothesized form of a function (e.g., a linear kernel for linear relationships, or a periodic kernel for periodic functions), compositional kernels thus allow for new compositional biases (e.g., a linear periodic relationship describing our alarming climate trends), similar to how rules can be combined to create new composite rules. Composing multiple kernels thus allows for aggregating multiple hypotheses about the hidden structure of the environment. Notably, earlier work by Brehmer (1974) already proposed people learn functions by testing how well increasingly more complex rules can explain the observed data. The Gaussian Process framework further formalizes this idea and injects the ability to reason about compositional rules as well.

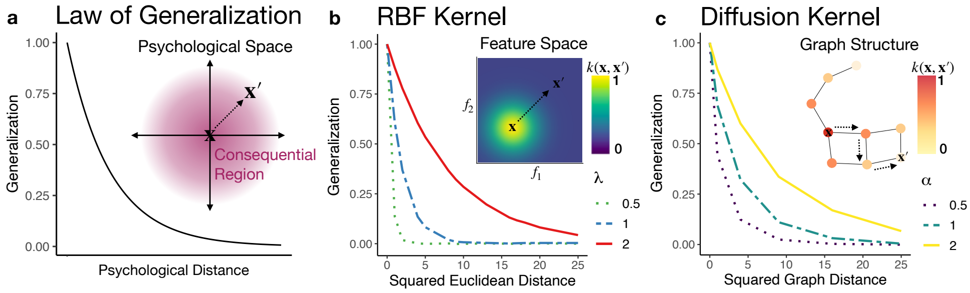

Figure 2

Principles of Generalization. **a**) Shepard's (1987) Law of Generalization describes generalization as a function of distance (dashed arrow) between stimuli in psychological space (inset). The smooth gradient of generalization arises due to uncertainty about the extent of a consequential region, with more distant stimuli less likely to belong to the same region. **b**) An RBF kernel provides a similarity metric based on the squared Euclidean distance between stimuli in feature space (inset: dashed arrow), producing similar generalization gradients as Shepard's model (quantified using Pearson correlation between expected outputs). The lengthscale parameter λ governs the rate at which generalization decays as a function of distance. **c**) In structured environments, a diffusion kernel (Kondor & Lafferty 2002; Wu et al. 2021) offers an analogous similarity metric based on the connectivity structure of a graph, where the diffusion parameter α governs the rate that previous observations "diffuse" over the graph.

2.3. Converging Historical Traditions

Given similar historical developments in concept and function learning, there is much to be gained from integrating theories of generalization across domains. In concept learning (Figure 1a-d), Shepard's (1987) "Universal Law of Generalization" provides an influential similarity-based approach, where generalization is characterized as distance in "psychological space", with stimuli embedded at closer distances are more likely to produce the same responses (Figure 2a). Yet through a probabilistic application of rule-based mechanisms, characterized by explicit hypotheses about the boundaries of a category, a hybrid Bayesian concept learning framework (Figure 1d; Tenenbaum & Griffiths 2001) can reproduce the same smooth gradient of generalization, showing how rule- and similarity-based mechanisms can be interpreted as two sides of the same coin (Pothos 2005; Austerweil et al. 2015).

In the domain of function learning, there has been an analogous trajectory of rule- and similarity-based theories culminating in hybrid approaches using Bayesian principles (Figure 1e-h). Current hybrid theories of function learning based on Gaussian Process regression utilize similarity-based mechanisms implemented through "kernel functions" that capture inductive biases (e.g., stimuli with similar features will yield similar outputs), yet also provide rule-based interpretations and allow for compositional operations over different kernels (Lucas et al. 2015; Schulz et al. 2017). An RBF kernel is a common choice, defining a similarity metric based on distance in feature space and producing similar generalization gradients as Shepard's theory (Figure 2b), However, a wide range of possible kernels provide other inductive biases not present in typical similarity-based theories, such as assumptions about linear, periodic, or graph-structured relationships (Figure 2c; Wu et al. 2021), which can be compositionally combined to yield new rule-like assumptions (Duvenaud et al. 2013; Schulz et al. 2017).

Ultimately, this convergence of concept and function learning has leveraged the strengths of both rule- and similarity-based approaches to illuminate the rich tapestry of human generalization. Building on these historical developments, we now turn to examine how principles of concept learning and function learning have informed new domains of generalization, tackling increasingly complex and challenging domains.

Reinforcement learning: A framework for understanding learning through trial-and-error feedback from the environment

Value function approximation: A key method for generalization in reinforcement learning, by expected value of different states or actions.

Bandit problem: Experimental paradigm used to investigate the trade-off between exploration and exploitation. Decision-makers repeatedly choose among options to accumulate payoffs, where each option yields probabilistic rewards.

Directed exploration: A strategic approach in learning or decision-making where exploration is guided by specific goals or hypotheses, such as reducing subjective uncertainty.

3. FROM LEARNING FUNCTIONS TO ACTING ON THE WORLD

Compared to concept learning, function learning has received relatively less attention and produced fewer experiments. Yet there has been a revival of interest, given the importance of *value function approximation* (Wang et al. 2020; Schölkopf 2015; Tesauro 1995) for generalization in Reinforcement Learning (RL) problems (Sutton & Barto 2018). RL provides a computational framework for understanding learning in both biological and artificial systems and can trace its origins to early research on associative and instrumental learning (Thorndike 1911; Pavlov 1927; Skinner 1938). Characterizing learning as a trial-and-error process, RL agents learn to associate different actions with expectations of reward through feedback from the environment, leading to gradual improvements in selecting reward-maximizing behaviors. However, no biological or artificial agent can try every possible action in most real-world settings, highlighting the necessity to generalize feedback from past experiences to novel settings. And while simple two-alternative choice tasks (e.g., a 2-armed bandit) are still commonly used in RL to study human learning through repeated experience with a small number of alternatives, there is a growing impetus to better understand how humans generalize in more real-world contexts, where the space of possible outcomes is too vast to be experienced exhaustively (Wu et al. 2018).

Research in RL has long grappled with the need to generalize to novel actions and states (Tesauro 1995). In complex games such as Go (Silver et al. 2016), the number of possible game states vastly outnumbers the number of atoms in the known universe. Thus, in scaling up to solve increasingly complex tasks, more contemporary RL algorithms commonly learn a *value function* mapping a vast space of potential actions or states to expectations of reward (i.e., value) (Wang et al. 2020; Sutton et al. 2000; Tesauro 1995; Silver et al. 2016). This estimated value function can then be used to generalize a limited number of experienced outcomes to a vast and potentially infinite space of possibilities, guiding efficient exploration and action selection. Here, we review recent advances in understanding human generalization in RL settings that do not permit exhaustive exploration and connect these findings to theories from function learning.

3.1. Generalization in Reinforcement Learning Using Value Function Approximation

Several recent studies have investigated human generalization in RL problems (Wimmer et al. 2012; Wu et al. 2018; Schulz et al. 2018b; Norbury et al. 2018; Stojić et al. 2020; Giron et al. 2023; Witt et al. 2023, Figure 3a). A common feature of these problems is the use of structured rewards, where participants can leverage features of the environment, such as spatial location (Wu et al. 2018), abstract features (Wu et al. 2020; Stojić et al. 2020; Norbury et al. 2018), or nodes on a graph (Wu et al. 2021) to predict the value of novel actions or stimuli.

In these structured bandit problems, participants are given the goal of maximizing rewards by iteratively selecting options (e.g., tiles, Gabor patches, or nodes on a graph;

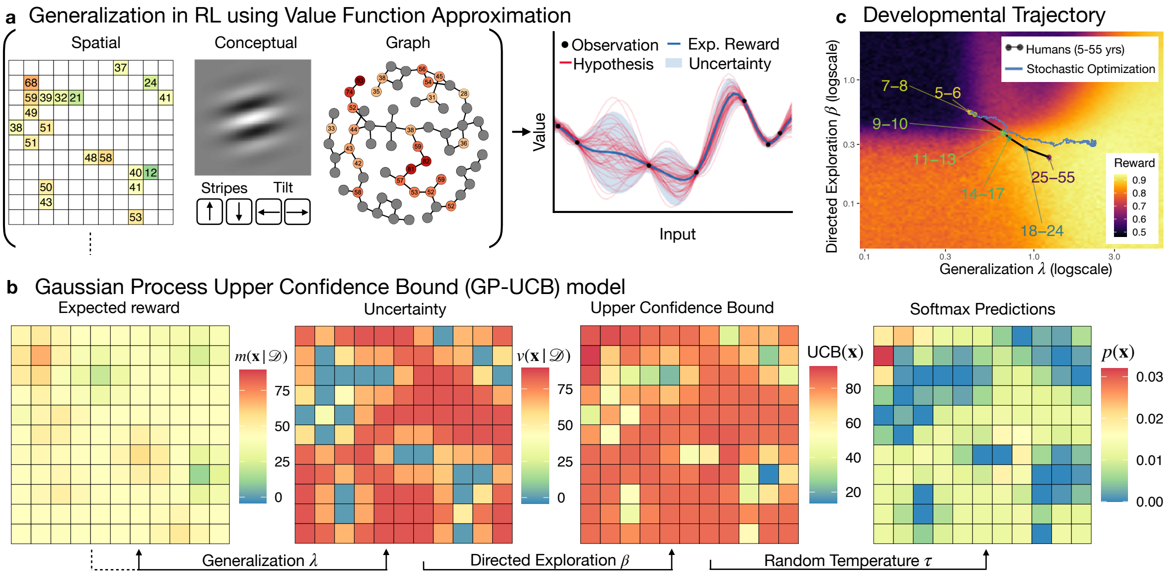

Figure 3

Value Generalization in RL. **a**) Reward generalization in Reinforcement Learning as Value Function Approximation. Left: Bandit tasks with structured rewards, where similar locations, feature combinations, or connected nodes generate similar rewards. Right: Generalization can be modeled using Gaussian Process regression to infer a value function, mapping a potentially infinite range of actions or states to probabilistic predictions about expected reward and subjective uncertainty. **b**) Overview of the Gaussian Process Upper Confidence Bound (GP-UCB) model in a spatial bandit task. Conditioned on the observations in panel a (left), the Gaussian Process model makes predictions about expected reward $m(\mathbf{x})$ and uncertainty $v(\mathbf{x})$, where the parameter λ governs the extent that past observations generalize to new options. Upper confidence bound (UCB) sampling combines the expected rewards $m(\mathbf{x})$ and uncertainty $v(\mathbf{x})$ using a weighted sum, where the parameter β defines the value of exploring uncertain options relative to exploiting high reward expectations. Lastly, UCB values are transformed into probabilistic predictions of where the participant will search next using a softmax function, where the temperature τ governs the amount of random exploration. **c**) The developmental trajectory of human learners (5-55yrs) resembles stochastic optimization over GP-UCB parameters. The labeled dots are the median parameter estimates from human subjects, while the blue line is the trajectory of the best-performing stochastic optimization algorithm computed over a fitness landscape of 1 million parameter combinations (Giron et al. 2023).

Figure 3a), each yielding stochastic rewards. To emphasize the need for generalization, participants are typically given a search horizon that is greatly limited compared to the number of unique options. However, rather than each option having independent reward distributions—as is commonly the case in bandit tasks—here, the expected rewards are correlated, such that options in similar spatial locations, with similar features, or well-connected to one another on a graph, are expected to yield similar rewards. This correlated reward structure provides traction for generalization, allowing participants to guide the selection of actions toward promising regions of the search space. In this body of research, human generalization is typically best characterized by the same Gaussian Process model as in traditional function learning tasks, but based on the implicit learning of a value function, which is then used to predict actions (Figure 3b).

One important distinction between generalization in an RL setting compared to traditional function learning is that the goal in RL is not necessarily to learn the true underlying value function. Rather, one needs to balance the *explore-exploit dilemma* (Mehlhorn et al.

2015), by both exploring uncertain options to acquire information, while also exploiting options with high expectations of reward to maximize immediate gains. Thus, a fundamental challenge in RL is to determine what information should be acquired given current beliefs (i.e., active learning; Settles 2009; Nelson 2005). In this RL setting, the Gaussian Process model addresses this challenge by making Bayesian predictions about the expected reward that include quantifications of uncertainty (blue line and ribbons, respectively, in Figure 3a right). These two components can be used to implement and test different sampling strategies in their ability to balance the exploration-exploitation dilemma and predict human behavior (Sidebar 3.1). Two prominent mechanisms of exploration are uncertainty-directed exploration and random exploration, which play a dissociable role in human exploration (Wu et al. 2018; Wilson et al. 2014; Wu et al. 2022a; Cogliati Dezza et al. 2019), with different neural signatures (Zajkowski et al. 2017) and distinct developmental trajectories (Giron et al. 2023; Meder et al. 2021; Schulz et al. 2019; Somerville et al. 2017a).

Value Generalization in Active Learning

Efficient learning in RL requires a balance between exploring new options for information and exploiting known high-value options for immediate reward. In complex problems with numerous choices, an added dimension is determining *where* to explore. Gaussian Process regression offers a Bayesian approach (Figure 3a-b) to predict both expected rewards (posterior mean $m(\mathbf{x}|\mathcal{D})$) and the associated uncertainty (posterior variance $v(\mathbf{x}|\mathcal{D})$). These predictions can be integrated with an *Upper Confidence Bound* (UCB) sampling strategy, which uses a weighted sum of reward expectations and uncertainty estimates to assign value to each option:

$$\mathrm{UCB}(\mathbf{x}) = m(\mathbf{x}|\mathcal{D}) + \beta\sqrt{v(\mathbf{x}|\mathcal{D})}. \tag{7.}$$

Here, the "exploration bonus" β dictates the value given to exploring uncertain options relative to exploiting immediate rewards (i.e., uncertainty-directed exploration). These UCB values can then form probabilistic predictions about which option an agent will choose next using a softmax function:

$$p(\mathbf{x}) \propto \exp\left(\mathrm{UCB}(\mathbf{x})/\tau\right), \tag{8.}$$

where options are selected proportional to their UCB value (Figure 3b). The temperature parameter τ governs the randomness of these predictions, providing an additional, undirected form of exploration.

Figure 3b illustrates how the reward expectations and uncertainty estimates of a Bayesian function learning model are combined to predict choices in a spatially correlated bandit problem (Wu et al. 2018). The best account of human choice behavior combines Gaussian Process regression as a model of value generalization with Upper Confidence Bound (UCB) sampling (Auer 2002), which quantifies the value of a choice option by adding an uncertainty-based "exploration bonus" to reward expectations (Srinivas et al. 2009). Together, the GP-UCB model demonstrates how generalization and exploration mechanisms interact to guide decision-making in RL. Additionally, several studies have fit the GP-UCB model on choices and then used it to predict out-of-sample judgments participants made about the expected reward of unexplored options, along with subjective confidence ratings (Wu et al. 2020, 2021). Thus, participants not only select actions "as-if" they are using a form of Bayesian value function approximation, but the same computations can also be used to predict their judgments, showing a correspondence between implicit value generalization in RL with explicit function learning in psychology.

3.2. Developmental Changes in Generalization and Exploration

By integrating generalization via value function learning with both uncertainty-directed and random exploration for active learning, the GP-UCB model has provided a powerful lens for understanding developmental changes in learning. Human development is often likened to a "cooling off" process (Gopnik et al. 2017, 2015), in analogy to mechanisms of stochastic optimization used in modern machine learning models. Like a heated piece of metal that becomes harder to manipulate as it cools off, stochastic optimization algorithms start off highly flexible and open to a wide range of solutions, even those that might not seem very good at first. But as they "cool down", they become less flexible and more selective in favoring only local improvements. This analogy is appealing, since children are highly stochastic and flexible learners, with the randomness of their choices (Bonawitz et al. 2014) and hypotheses (Buchsbaum et al. 2012; Lucas et al. 2014; Gopnik et al. 2017; Denison et al. 2013) gradually diminishing over the lifespan.

Yet, there is ambiguity in how to interpret this verbal analogy, with the most common being "cooling off" as a uni-dimensional transition from exploration to exploitation (Gopnik 2020), focusing solely on a reduction in random exploration. And while past work has indeed found differences in random exploration between different age-groups (Somerville et al. 2017a; Blanco & Sloutsky 2021; Schulz et al. 2019; Meder et al. 2021), this only seems to be part of the picture, with developmental differences in more systematic, uncertainty-directed exploration (Schulz et al. 2019; Blanco & Sloutsky 2021; Somerville et al. 2017b), along with various aspects of belief integration and generalization about novel options (Van den Bos et al. 2012; Blanco et al. 2016).

To provide a concrete test of the "cooling off" analogy, Giron et al. (2023) directly compared the trajectory of human learners (aged 5 to 55) against that of various optimization algorithms (Figure 3c). The results showed that "cooling off" does not only apply to the single dimension of randomness. Rather, development resembles an optimization process in the space of learning strategies: what begins as large tweaks in the parameters that define learning during childhood, plateaus and converges in adulthood. While the developmental trajectory of human learning strategies is strikingly similar to the best-performing algorithms (Figure 3c), none discovered reliably better regions of the strategy space than adult participants, suggesting a remarkable efficiency of human development.

In sum, investigating exploration and exploitation within RL settings has proven to be a productive approach for understanding human generalization across the lifespan in increasingly complex scenarios. The use of different kinds of structured reward environments constitutes an important step towards more complex and naturalistic learning and decision problems that capture the interrelatedness of real-world experiences (Wise et al. 2024). Such tasks help to better understand what computational and psychological principles guide action selection in expansive state spaces, where the complexity of the state and action space does not permit exhaustive search and learning, naturally requiring predictive generalization and leveraging uncertainty to efficiently act on the world.

4. FROM LEARNING CONCEPTS TO LEARNING STRUCTURE

Human generalization is much deeper than just comparing features at face value. Rather, generalization also depends on the relational structure and temporal dynamics of the environment, which are often hidden and need to be inferred. Research on structure learning can be broadly divided into two traditions. The first originates from Tolman's (1948) pioneering notion of a "cognitive map". Research in this domain has extensively studied spatial navigation in the hippocampal-entorhinal system (Whittington et al. 2022; Epstein et al. 2017; Moser et al. 2014), which has since been extended to a wide range of non-spatial modalities and domains (Behrens et al. 2018). The second tradition, known as Bayesian structure induction (Kemp & Tenenbaum 2008, 2009), builds on a similar formalism as Bayesian concept learning (Tenenbaum & Griffiths 2001), where explicit, rule-like hypotheses about structure can be inferred from observed data, reflecting our ability to discern patterns and regularities in the environment. While these traditions are based on different theoretical foundations, here we show that they share a common framework of similarity-based mechanisms for learning rule-like hypotheses about structure.

Cognitive Map: A mental representation of the structure of the environment, used for navigation, learning, and generalization

Successor representation: An RL model using anticipated future states of the environment for predictive generalization

Structure induction: The process of inferring underlying structure from observed data, often using Bayesian principles

4.1. Cognitive Maps

Originally proposed as an alternative to stimulus-response learning, Tolman (1948) found that rats could rapidly adapt to new situations (e.g., choosing the second shortest path in a maze when the shortest path is blocked) and to new goals (e.g., efficiently navigating to food rewards placed in novel locations of a familiar maze). These results suggested the rats had generalized their experiences based on establishing a "field map of the environment" (Tolman 1948, pg. 2), Today, this notion of a cognitive map is grounded in neural evidence (in humans and other animals) relating the activity of specialized cells in the hippocampal-entorhinal system to computations facilitating navigation and self-location, such as encoding spatial orientations, boundaries, and distance to objects (see Moser et al. 2014; Peer et al. 2021; Epstein et al. 2017, for reviews). And as Tolman originally speculated, cognitive maps are not only restricted to representing spatial structure. Rather, the same neural machinery used for spatial navigation also encodes relational and structural knowledge across a wide range of domains, including social relationships (Tavares et al. 2015), smells (Bao et al. 2019), abstract visual features (Constantinescu et al. 2016), and the connectivity of hidden graph structures (Garvert et al. 2017).

One influential account of structure learning in the hippocampal-entorhinal system is the Successor Representation (SR; Dayan 1993; Stachenfeld et al. 2017; Momennejad 2020). Originally developed as a method to improve the generalization of Temporal Difference (TD) learning (Sutton & Barto 2018), the SR describes a decomposition of the value function into a similarity matrix and the singular rewards of each state (Figure 4a). The similarity matrix quantifies the similarity between each pair of states based on expected future state transitions, influenced by both the structure of the environment and the agent's behavioral policy (i.e., how the agent moves around in the environment), and thus corresponds to an explicit, graph-like representation of the environment (Peer et al. 2021). The value generalizations predicted by the SR—taking the form of a linear combination of state similarities and reward observations—captures the underlying transition dynamics and connectivity structure of the environment, with stronger generalizations between well-connected states. Related methods using kernel similarity (Gershman et al. 2017; Wu et al. 2021, Figure 2c) rather than SR similarity, operate on similar principles, with exact equivalencies in special cases (Machado et al. 2018). For example, Garvert et al. 2023 showed that the Gaussian Process kernel can be approximated by the successor representation of visited states in an open environment and then used to successfully predict human choices in a bandit task, illustrating a continuity between cognitive maps and value generalization using function learning.

The SR provides a candidate mechanism for how the structure of the environment can be learned through experience (i.e., "on-policy" learning). This on-policy method learns the similarity matrix encoding environmental structure using the familiar computations of prediction-error learning (Rescorla & Wagner 1972; Sutton & Barto 2018), but by predicting future state transitions rather than predicting rewards (Dayan 1993; Russek et al. 2017a). However, some of the most convincing demonstrations of the SR as a model of the hippocampal-entorhinal system are based on "off-policy" methods, which sidestep the problem of learning latent structure by simply assuming a random policy over infinite time, allowing for an analytic solution (Stachenfeld et al. 2017; Momennejad et al. 2017). One motivation for using off-policy methods is that on-policy structure learning is slow, requiring exhaustive exploration of an environment before an accurate model of the environment can develop. In contrast, humans can rapidly learn new structures with relatively few experiences in on-policy settings (Mark et al. 2020; Rubino et al. 2023). For instance, consider how you might deftly navigate a foreign airport based on intuitions about previously experienced airport layouts or how you might transfer domain knowledge about bread baking to a new problem such as making steamed buns. In addition to these limitations, the SR only makes point estimates about expected reward (but see Geerts et al. 2019; Madarasz & Behrens 2019; Mark et al. 2020, for Bayesian extensions), which limits the active learning mechanisms it has access to (Sidebar 3.1).

Overall, the SR provides an elegant and simple theory of structure learning within the RL framework, where similarity-based representations acquired through associative learning enable structure-informed generalization through value function approximation. However, the slowness of this learning process may fall short of explaining the full efficiency with which humans learn relational structure. Other recent theories of cognitive map learning such as the Tolman Eichenbaum Machine (TEM; Whittington et al. 2020) combine path integration (Mittelstaedt & Mittelstaedt 1980) with conjunctive memory (Manns & Eichenbaum 2006) to more efficiently learn latent structure. And while the TEM is capable of transferring learned structures to new environments, it cannot infer entirely novel structures. In contrast, humans can reason compositionally about new relational structures that they have never experienced before. Consider how you can imagine novel food combinations that have never been observed (e.g., tea-flavored jelly, Barron et al. 2013, or broccoli-flavored ice cream, Gershman et al. 2017) or novel configurations of previously encountered structures, such as predicting where your gate might be when racing through a foreign airport to catch a connecting flight. This highlights the necessity for more explicit, rule-based theories of cognitive map learning, which is an area ripe for further exploration.

4.2. Structure Induction

Structure induction (Kemp & Tenenbaum 2008; Meder et al. 2014; Griffiths & Tenenbaum 2009b; Kemp & Tenenbaum 2009; Lynn & Bassett 2020; Ke et al. 2022) provides an alternative approach to inferring the underlying structure or organizational pattern in a set of observations or data. For example, inferring the taxonomy of different animals based on

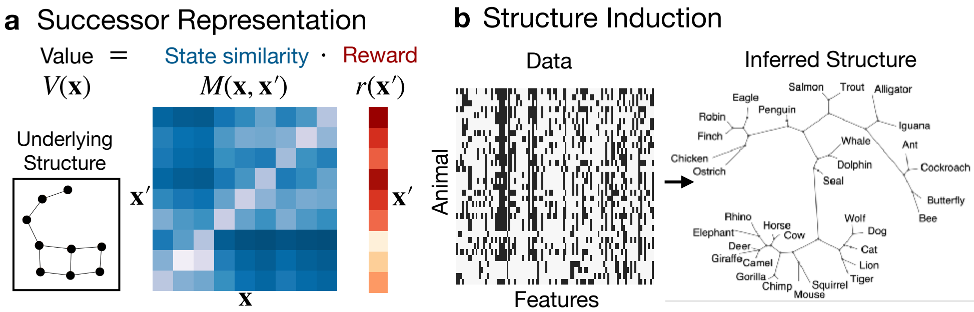

Figure 4

Structure Learning. **a**) The Successor Representation (SR; Dayan 1993) defines a decomposition of a TD-learning (Sutton & Barto 2018) value function $V(\mathbf{x})$ into a similarity matrix $M(\mathbf{x}, \mathbf{x}')$ based on expected future state transitions, and the singular rewards of each state $r(\mathbf{x}')$. The state similarities are a function of the underlying structure (left) and the agent's policy, and allow for generalization via a linear form of value function approximation. **b**) Bayesian structure induction (Kemp & Tenenbaum 2008) uses Bayesian principles to infer the underlying structure (e.g., a taxonomy) that gave rise to observed relational data (e.g., animals and their shared features).

their shared features (Figure 4b). Animals with similar features can be expected to occupy closely connected positions in a taxonomy, yet there is a large hypothesis space of possible configurations. Here, Bayesian structure induction (Kemp & Tenenbaum 2008) uses a similar mathematical formalism as Bayesian concept learning (Tenenbaum & Griffiths 2001), based on describing a distribution of rule-like hypotheses, which are evaluated based on their similarity to the observed data. Instead of defining hypotheses about category boundaries, structure induction defines an inference process operating on hypotheses about structural configurations (i.e., different graph structures). A prior over hypothesized graphs encodes a preference for simpler structures, with each hypothesis weighted according to its likelihood of generating the observed data (Kemp & Tenenbaum 2008).

Although Bayesian inference about latent structure is often intractable when scaled to complex problems, hybrid models of structure induction can circumvent this problem by incorporating similarity-based mechanism. One notable method (Kemp & Tenenbaum 2008, 2009) used to evaluate the likelihood of each candidate hypothesis is identical to a Gaussian Process[1]. Here, each candidate hypothesis is used to parameterize a graph kernel (Zhu et al. 2003), and simulated observations are sampled from the Gaussian Process prior (Eq. 4). Higher similarity between generated and observed data corresponds to a higher likelihood for the hypothesized structure.

Thus, structure learning and function learning can be seen as complementary problems to each other, with a hybrid approach to Bayesian structure induction relying on the same Gaussian Process computations as in function learning (Lucas et al. 2015) and value generalization settings (Wu et al. 2021). Hypotheses about rule-like structures are used to define a similarity metric based on a Gaussian Process kernel, and used to simulate data. Comparisons between the simulated and observed data facilitates inference about which

[1]The authors refer to this as a Gaussian Markov Random Field (Zhu et al. 2003), which is a multivariate Gaussian distribution identical to a Gaussian Process prior (Eq. 4).

structures are most likely. Once the structure has been inferred, these same computations can then be reused to generalize about novel outcomes (Kemp & Tenenbaum 2008) and guide exploration (Wu et al. 2021; Ludwig et al. 2022) in structured environments. For instance, this complementary relationship has also been leveraged to generalize about novel properties of the data (i.e., property induction; Kemp & Tenenbaum 2009). Given a set of binary features of various animals (Figure 4b), structure induction can be used to infer the underlying taxonomy structure. Once a posterior distribution over structured has been defined, the same Gaussian Process function learning approach (with an additional binarization of outcome variables) is used to infer the probability of novel features. If you were to learn a new fact about squirrels (e.g., their front teeth never stop growing), you might be more likely to generalize this fact to similar animals, such as mice, but less likely to generalize it to more dissimilar animals, such as penguins.

In summary, structure induction offers a prime example of the complementarity between rule- and structured-based mechanisms. Rule-based computations over a distribution of hypothesized structures offer the possibility of rapid generalization. Yet the intractability of Bayesian inference can be side-stepped through sample-based approximations, using Bayesian function learning operating over similarity-based computations. Together, these complementary approaches to generalization support both the inference of latent structure and the use of this structure to infer new features and outcomes.

5. GENERAL DISCUSSION

We have traced the development of psychological theories of generalization, from foundational research on concept learning and function learning to more modern domains of RL and latent structure induction. Throughout this long history, fierce debates between rule- and similarity-based theories have been reconciled through the development of hybrid models, often based on Bayesian principles. The ongoing success of hybrid models suggests that accommodating both rule- and similarity-based representations is central to explaining human generalization.

Yet, each approach makes computational commitments to a specific representational format, offering distinct advantages. Similarity provides a flexible and efficient approach to generalization, relating new situations to prior experiences, and leveraging relational knowledge when the underlying structure is known. In turn, rules unlock compositionality, facilitating generalization and inference about novel structures, which is exemplified in Bayesian structure induction. However, there may also be exchangeability between rule-based and similarity-based mechanisms of generalization, suggesting a dynamic interplay that enables adaptive learning through hybrid approaches that blend both strategies. We first explore these themes, before plotting out a trajectory for the future of research on generalization.

Rules Unlock Compositionality, but Are Challenging to Learn. Rule-based mechanisms are foundational to our understanding of generalization, drawing upon a rich history of theoretical and empirical research (Simon & Lea 1974; Bruner et al. 1956; Ashby & Gott 1988; see Ashby & Maddox 2005, for a review). These mechanisms are particularly effective in structured domains, where the precision of rules facilitates rapid, one-shot generalization (Dasgupta et al. 2022). Whether taught pedagogically (e.g., "i before e except after c") or learned through experience (e.g., "talking loudly in the library is forbidden"), rules represent explicit hypotheses about regularities of the environment extracted from data (Reber & Lewis 1977). In concept learning, rules can represent hypotheses about the boundaries between categories, while in function learning, rules can represent hypotheses about the (parametric) relationship between inputs and outputs. Signatures of rule-based mechanisms can also be seen in theory-based RL (Tsividis et al. 2021; Deisenroth & Rasmussen 2011; Allen et al. 2020), where agents generate hypotheses about the underlying rules governing its environment to inform learning and exploration (e.g., "keys open doors, but only if the colors match" in game environments; Pouncy et al. 2021).

This ability to reason about and use rules thus unlocks an unrivaled capacity of human intelligence, since rules allow for compositional and syntactic manipulation (Piantadosi et al. 2016; Dehaene et al. 2022). Indeed, the power of logic and mathematics can be thought of as nothing more than the manipulation of syntactic rules (Newell & Simon 1976). Thus, rule-based mechanisms unlock the ability to compositionally combine multiple rules or substructures to generate an infinitely productive space of potential hypotheses. Recent advances in program induction (Rule et al. 2020; Ellis et al. 2023; Lake et al. 2017)—using similar computations as Bayesian concept learning (Tenenbaum & Griffiths 2001) and structure induction (Kemp & Tenenbaum 2008)—indicate a promising framework for modeling how humans infer generative rule-like structure from data, providing a modern interpretation of Fodor's (1975) Language of Thought (LoT). However, the compositionality of rules also creates a combinatorial explosion of possible hypotheses, making search and inference increasingly difficult (Fränken et al. 2022). Thus, despite the utility of rule-based mechanisms, open challenges lie in their complexity and the demands they place on cognitive resources for generating and testing new hypotheses (Rubino et al. 2023).

Similarity Is Flexible, but Can Be Arbitrary. Similarity-based mechanisms for generalization are ubiquitous in psychology (Tversky 1977; Shepard 1987; Tenenbaum & Griffiths 2001; Chater & Vitányi 2003; Jäkel et al. 2008a; Gershman & Daw 2017; Botvinick et al. 2019). The notion that stimuli with similar features or occurring in similar contexts are more likely to belong to the same category or yield comparable outputs is a powerful principle of generalization, and can be flexibly applied to a wide range of domains. While historically defined based on feature comparisons or by appeal to some abstract psychological space (Shepard 1987), recent advances have expanded these mechanisms to capture rich relational structures (Wu et al. 2021) based on network connections (Lau et al. 2020; Tavares et al. 2015) or environmental dynamics (Stachenfeld et al. 2017; Machado et al. 2018), thus extending similarity-based theories of generalization to increasingly structured environments.

However, these mechanisms are not without drawbacks. It is far from straightforward to simply go out into the world to measure how similar things are to one another. Consider how naturalistic stimuli have a host of different features and relationships, offering a potentially unlimited number of ways by which similarity can be computed (Goodman 1972). Should an apple be compared to an orange on the basis of color, shape, taste, or country of origin? Thus, one must specify with respect to which features (or via which relationships) the stimuli are being compared (Medin et al. 1993). This is often dependent on the underlying context: when at a fruit orchard, color might provide a useful comparison on the basis of ripeness, whereas, at a customs office, country of origin is more relevant for determining the amount of tax to levy. Thus, the endless ways in which different stimuli can be compared has led to the criticism that similarity is too flexible (Murphy & Medin 1985), potentially undermining its utility as a concept in psychology. The context-dependent nature of human similarity judgments can also lead to paradoxical conclusions, as illustrated by violations of logical axioms like the law of triangle inequality (Tversky 1977). And while recent theories of rational attention have proposed associative learning mechanisms for gradually ignoring reward-irrelevant features (Radulescu et al. 2021), this approach is only feasible for simple stimuli with a handful of predefined features. In more naturalistic settings, stimuli may have a potentially innumerable set of features, making it infeasible to gradually prune irrelevant features from an infinite set. These complexities illustrate the nuanced and sometimes contradictory nature of similarity-based generalization in human cognition.

Similar challenges also apply when defining similarity representations over latent structure, which share context- and goal-dependent assumptions about which features are relevant. For instance, the development of Darwin's Tree of Life was rooted in targeted observations about features that were shared or differed between species (Doolittle & Bapteste 2007), while dimensional accounts of psychopathology similarly aim to capture shared symptom patterns across mental illnesses based on a targeted subset of features (Kotov et al. 2021, 2017). Thus, while latent structure plays a pivotal role in generalization by complementing similarity-based inferences, it shares some of the very same challenges that arise in defining relevant features for computing similarity. This intertwined nature of similarity and structure highlights both their importance and the enduring challenges concerning their role in human cognition and generalization.

Integrating Rules and Similarity. We have highlighted the relative advantages and disadvantages of rule- and similarity-based mechanisms of generalization. However, the success of hybrid approaches suggests it is not necessarily one or the other. Rather, there is likely a degree of exchangeability between rules and similarities, involving transformations from one currency to the other (Cushman 2020). This is not a new concept. In RL, model-based representations of the environment can be used to rationally plan out actions (Miller et al. 2017), but in the process, new value and policy representations are constructed, supporting future model-free action selection (Kool et al. 2018). In social learning, observed actions can be "unpacked" via inverse reinforcement learning (IRL; Jara-Ettinger 2019) to infer latent model-free and model-based representations assumed to have generated the behavior (Wu et al. 2022b). Thus, the caching of past computations (i.e., amortization; Dasgupta et al. 2018) and inference via IRL provide two mechanisms by which the representations involved in model-free and model-based RL are exchanged and combined with one another (Cushman 2020; Wu et al. 2022b). Our current theories in this domain suggest that we use a mixture of strategies, composing elements from each mechanism into an adaptive mixture of representations (Russek et al. 2017b; Keramati et al. 2016; Huys et al. 2015).

Model-free learning: Category of RL methods using reward outcomes to learn a behavioral policy and value function, without simulating future scenarios.

Model-based learning: More complex form of RL, which builds a model of the environment to simulate and plan future actions.

Are rules and similarity-based representations exchangeable in a similar sense? Rule-based representations about category boundaries, functional forms, or the structure of the environment can inform or be directly used to define similarity representations. We have shown how rule-like hypotheses about the structure of some latent graph can be used to define a similarity matrix using a graph kernel (Figure 2c), to infer rule-like representations about the latent structure (Kemp & Tenenbaum 2008), predict novel features outside of the training data (Kemp & Tenenbaum 2009), or to perform value generalization in an RL setting (Wu et al. 2021). In the other direction, we have also shown how similarity-based representations support the inference of rule-like hypotheses about latent structure. The SR (Dayan 1993) leverages simple associative learning mechanisms to learn a similarity matrix, corresponding to a rule-like hypothesis about the underlying latent structure of the environment. Even more directly, hybrid models of Bayesian structure induction (Kemp & Tenenbaum 2008) have relied on similarity-based computations using Gaussian Process kernels to simulate data under each hypothesized graph structure. Thus, learned rules can be "cached" as similarity representations, facilitating rapid and efficient generalization. Meanwhile, inferring rules and structure can be supported by sample-based approximations, where each candidate hypothesis can be used to construct a similarity representation to perform tractable inference.

SUMMARY POINTS

1. Rules and similarity are foundational concepts across the entire expanse of psychological research on how humans generalize from limited experiences to novel situations.
2. Hybrid models include elements of both rule- and similarity-based approaches, providing a unified computational framework for investigating human generalization across diverse contexts.
3. Gaussian Process function learning coupled with uncertainty-directed exploration provides a model of generalization and active learning in a wide range of reinforcement learning problems with large decision-spaces.
4. Structure learning supports similarity-based generalization by representing latent relational structure and the temporal dynamics of the environment, while conversely, similarity-based mechanisms may play a key role in learning latent structure.
5. Rule- and similarity-based representations have complementary advantages, with an exchangeability between these representations offering insights into how humans simultaneously display flexible and compositional generalization.

5.1. The Future of Generalization

Having surveyed the past and present, we now turn our attention to the future. Psychology—like many sciences—is often at its best when combining different approaches, where recent advances demonstrate the potential of combining computational and psychological theory. At the same time, several challenges remain to account for the unparalleled flexibility and efficiency of human generalization.

First, we propose a new integration of rule- and similarity-based mechanisms for structure learning in RL settings, combining their relative strengths and leveraging the exchangeability of representations to achieve more a comprehensive framework of generalization. Second, we point out fundamental connections between Gaussian Process regression and theories of episodic memory, which suggest the potential for developing boundedly rational models of generalization to account for cognitive limitations. Third, there is still a need to explore generalization in environments that more closely resemble real-world conditions, requiring the integration of individual and social information. By addressing these issues, future research will continue a long and always central line of research seeking to understand how humans adapt and continually improvise and adapt to novel situations.

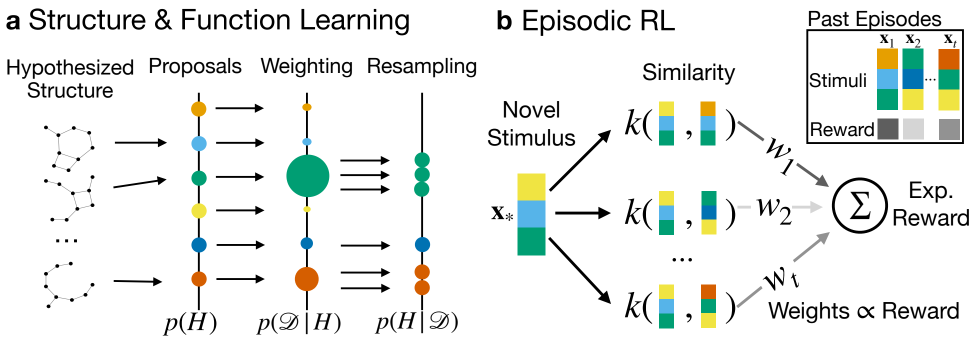

Figure 5

Future Directions. **b**) Integrating structure learning with function learning under a common framework, using a particle filter. **a**) The Episodic RL framework provides a different conceptualization of the computations in a Gaussian Process, with exact equivalencies to expected reward predictions (but not uncertainty). Here, similarity is computed between new stimuli and past episodes, which are then weighted by rewards and summed up.

Particle filter: Approximate Bayesian technique, whereby hypotheses (particles) are sequentially refined by weighting them according to the likelihood of data and resampling.

Episodic reinforcement learning: A framework for value function approximation, which compares novel states, actions or stimuli to previously encountered episodes.

Combining Structure Learning and Function Learning. We propose that rather than accepting a duality of interpretation as the final synthesis of rule- and similarity-based mechanisms, future models of generalization could provide a more complete unification, utilizing each mechanism to its strengths. We have advocated for Gaussian Process function learning as a candidate model of human value generalization in many domains, where the kernel provides a similarity metric based on a given representation of the environment. Yet we currently lack a model that simultaneously infers structure while performing predictive generalization. Since Gaussian Processes play a key role in the computations of Bayesian structure induction (Kemp & Tenenbaum 2008), a future model (Figure 5a) could simultaneously perform inference over candidate structures and generate predictions about novel outcomes. Rule-based mechanisms can be used to propose hypotheses about structure (e.g., proposing different graph configurations by leveraging previously learned schemas; Wingate et al. 2013; Kemp & Tenenbaum 2008; Ellis et al. 2023; Le et al. 2021; Fränken et al. 2022; Rubino et al. 2023). Each hypothesized structure can then be used to parameterize a graph kernel (Fig. 2c), where a Gaussian Process using similarity-based mechanisms can be used to both predict new outcomes and to evaluate the likelihood of a given hypothesis (as in Kemp & Tenenbaum 2008). Lastly, Bayesian principles will allow us to describe a distribution over hypotheses (about structure), which can adapt to the observed data. With rules providing the structure and similarity providing the canvas, generalization combining both mechanisms can achieve both flexibility and efficiency.

One candidate algorithm to implement this proposal is a particle filter (Doucet et al. 2009; Speekenbrink 2016). A particle filter uses a finite set of hypotheses (i.e., particles), which are refined and updated based on new data to provide an approximation to Bayesian inference. Here, particles can represent different hypotheses about latent structure—for instance, a specific graph configuration. When encountering new data, particles are reweighted by their likelihood (e.g., by generating simulated observations from a Gaussian Process; Kemp & Tenenbaum 2008), and then resampled in proportion to this likelihood. Thus, inaccurate hypotheses die out, while more accurate hypotheses proliferate, and are refined. To take into account uncertainty about the underlying hypotheses (i.e., variance across particles), one could tap into the compositional nature of kernels (Rasmussen & Williams 2005; Schulz et al. 2017), and combine all current hypotheses (i.e., the resampled population of particles) into a composite kernel by averaging across particles. Such an approach could propagate uncertainty about the underlying structure through to uncertainty about potential outcomes, facilitating active learning at both levels.

Generalization with Limited Resources. While originating as a machine learning technique, Gaussian Process regression has direct links to psychological theories integrating RL mechanisms with episodic memory (Lengyel & Dayan 2007; Gershman & Daw 2017). In this light, Gaussian Processes can be understood as a Bayesian extension of *Episodic RL* (Gershman & Daw 2017; Botvinick et al. 2019). In Episodic RL (Figure 5a), an agent stores episodic memories about previously encountered stimuli and their associated rewards. To predict the value of some novel stimuli, one first computes similarity to each previously encountered "episode". Then, the reward value for each episode is multiplied by its similarity to the novel stimuli and then summed up. In other words, generalization is performed through inferring similarity-weighted expectations, where more similar episodes exert more influence on how their rewards generalize to the novel stimuli.

When using a kernel function to compute similarity (Gershman & Daw 2017), these predictions are equivalent to the posterior mean of a Gaussian Process (Jäkel et al. 2008b; Wu et al. 2021). This use of a similarity-weighted sum for generalization is reminiscent of classic exemplar-based theories of concept learning (Nosofsky 1986; Kruschke 1992; Medin & Schaffer 1978), while also having direct equivalencies to computational methods used in function learning. Furthermore, when using an RBF kernel (Figure 2b) as the similarity metric, Episodic RL is equivalent to an RBF network, which has featured prominently in machine learning approaches to value function approximation (Sutton & Barto 2018; Jäkel et al. 2008b) and as a theory of human generalization in the visual and motor systems (Poggio & Bizzi 2004). Thus, while the mathematics of Gaussian Process regression may seem unfamiliar to psychology, the underlying computations reoccur in numerous psychological theories of learning and generalization. However, a crucial difference is that the Gaussian Process—being a Bayesian model—also makes predictions with uncertainty, which play an essential role in describing human exploration (Wu et al. 2018; Giron et al. 2023; Wilson et al. 2014; Gershman 2018) and subjective confidence judgments (Wu et al. 2020, 2021).

The relationships between Gaussian Process regression and Episodic RL provide pathways for further integrating psychological and computational theories. For instance, to investigate the role of memory limitations in value generalization, one can induce memory load by removing information about previous choices and their outcomes (e.g., withholding observations from the grid shown in Figure 3a; Breit et al. 2022). In this case, learners would be reliant on episodic memory of past choices to generalize previous experiences. This scenario offers opportunities to study "Episodic Generalization", by applying the kernel function to a sample of episodic memories to make predictive generalizations. Such a paradigm would also facilitate the utilization of process tracing techniques, such as eye movements serving as indicators of memory retrieval (Spivey & Geng 2001; Johansson & Johansson 2014; Scholz et al. 2015). For instance, when answering a question about a stimulus previously presented in a particular spatial location, eye movements tend to be directed towards this (now empty) location (Spivey & Geng 2001). This "looking at nothing" suggests that episodic memory representations include spatial or oculomotor information, and that eye movements are indicative of the retrieval of episodic memories. In an RL task with spatially correlated rewards and memory load, eye movements may be directed towards previously sampled reward locations, forming the basis for similarity-based value generalization operating on the costly retrieval of memory traces.

What Is Still Missing?. Here, we have explored an expanding core of psychological research on generalization. From early work studying stimulus categorization and function learning, we have traced a continuity of mechanisms to new domains, such as active learning in RL and latent structure learning. However, the full scope of human generalization is certainly still beyond our current theories. Consider a chef figuring out how to substitute a missing ingredient in a recipe or a biologist identifying new species in an unexplored habitat. Generalization in both settings is informed by an interplay of rules and similarity—about the interactions between different foods and cooking techniques or about the interplay of biological traits, ecological niches, and reproductive success. Yet, the open-ended complexity of features to evaluate (Wise et al. 2024) and actions to consider (Moskvichev et al. 2023) present open challenges for our current theories. Additionally, chefs, biologists, and humans in all walks of life primarily learn from one another. While psychological research has often focused on studying isolated individuals learning from the environment (imagine a Skinner box as a canonical example), there is evidence of distinct mechanisms when learning from other people (Ho et al. 2017), compared to learning from the environment. Thus, future theories of human generalization must also account for more open-ended and socially embedded problems.

On one hand, psychological research has been continually expanding to investigate learning and generalization in more complex, and open-ended problems. While 2-alternative forced choice problems are still prevalent, experimental studies have begun to use increasingly complex tasks with higher-dimensional stimuli (Meagher & Nosofsky 2023), depleting rewards (Wu et al. 2023b), balancing rewards with avoiding risky outcomes (Schulz et al. 2018b), and non-stationary environments that change over time (Speekenbrink & Konstantinidis 2015). Here, open questions concern how relevant features of aspects of the latent environment structure are identified, and to which extent inductive biases simplify these inferences through strong, prior assumptions. In addition, there is currently great interest in studying generalization in the Abstraction and Reasoning Corpus (ARC; Chollet 2019). The ARC challenge is comprised of visual grids representing an abstract concept (input), with the decision-maker tasked with constructing an output grid corresponding to the input. This can be seen as type of function learning problem, requiring strong inductive biases about the generative process, since one needs to generate solution grids instead of only selecting from possible answers. These challenges may play a key role in explaining why AI approaches, including Large Language Models (LMMs), have yet to come close to human performance (Moskvichev et al. 2023). Thus, there is a promising future for efforts directed toward studies of generalization that integrate the complexity and open-endedness inherent in real-world decision-making environments.

On the other hand, a promising yet under-explored area is the integration of individual and social generalization mechanisms (Witt et al. 2023; Wu et al. 2023b). Far from being a peripheral feature, the capacity for social learning (and by extension social generalization; Witt et al. 2023) is often proposed as being the defining characteristics of human intelligence (Henrich 2016; Heyes 2018), differentiating us from other animals and AI (Wu et al. 2022b). Yet, research on generalization has commonly focused on individual learning in a vacuum. In many real-world contexts, however, we are surrounded by social information, which can greatly inform our generalization and decision-making processes. For instance, observing which menu items other customers order in a restaurant, and using that to inform your own choices. In such scenarios, individual and social learning mechanisms exhibit a dynamic interplay (Wu et al. 2023b), working in tandem to achieve efficient generalization. Here, there are also new applications for familiar concepts from individual generalization, since social information cannot always be taken verbatim, but needs to account for differences in individual preferences, abilities, and goals (Witt et al. 2023). Additionally, our ability to communicate via language in social settings offers new advantages for rule-based mechanisms, since they can be easily transmitted to one other (Wu et al. 2023a). Such scenarios offer a promising avenue for investigating how humans generalize and make decisions in real-world contexts, where social information plays a vital role in shaping adaptive behavior.

FUTURE ISSUES

1. The mechanisms underlying the integration of rule- and similarity-based generalization are still unknown, and the dynamic interplay between these processes should be explored in different learning contexts.
2. Combining structure induction with models of active learning is a promising direction for developing more comprehensive models of generalization that leverage the advantages of both rule- and similarity-based mechanisms.
3. Exploring the relationship between Gaussian Process regression and Episodic Reinforcement Learning provides a foundation for investigating how cognitive constraints, like working memory load, influence generalization under bounded rationality.
4. Investigating how humans and computational models navigate and generalize in high-dimensional spaces, will require new methods for identifying and prioritizing relevant features and relevant hypotheses.
5. The role of social learning has been under-represented in theories of generalization, with a need for new research studying how social and cultural contexts influence the mechanisms of generalization.

5.2. Conclusions

Human generalization has long been considered a hallmark of our unique cognitive abilities (Lake et al. 2017), with Roger Shepard famously proclaiming that the first general law of psychology should be a law of generalization (Shepard 1987). Here, we have traced the development of theories of generalization, illustrating a continuity of formerly competing mechanisms—rules and similarity—culminating in hybrid approaches. Ultimately, the future of generalization will hold new and exciting ideas, but still carry echoes of perennially reoccurring principles from history.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

All authors contributed to the conception of the paper and approved the final version. CMW wrote the first draft and created the visualizations, with contributions from BM. CMW is supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A, and funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy–EXC2064/1–390727645. We thank the members of the Human and Machine Cognition Lab for their helpful feedback.

LITERATURE CITED

Allen KR, Smith KA, Tenenbaum JB. 2020. Rapid trial-and-error learning with simulation supports flexible tool use and physical reasoning. *Proc. Natl. Acad. Sci. U. S. A.* 117(47):29302–29310

Ashby FG, Alfonso-Reese LA, Turken AU, Waldron EM. 1998. A neuropsychological theory of multiple systems in category learning. *Psychol. Rev.* 105(3):442–481

Ashby FG, Gott RE. 1988. Decision rules in the perception and categorization of multidimensional stimuli. *J. Exp. Psychol. Learn. Mem. Cogn.* 14(1):33–53

Ashby FG, Maddox WT. 2005. Human category learning. *Annu. Rev. Psychol.* 56:149–178

Auer P. 2002. Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.* 47:235–256

Austerweil JL, Gershman SJ, Tenenbaum JB, Griffiths TL. 2015. Structure and flexibility in bayesian models of cognition. *Oxford Handbook of Computational and Mathematical Psychology*

Bao X, Gjorgieva E, Shanahan LK, Howard JD, Kahnt T, Gottfried JA. 2019. Grid-like neural representations support olfactory navigation of a Two-Dimensional odor space. *Neuron* 102(5):1066–1075.e5

Barron HC, Dolan RJ, Behrens TEJ. 2013. Online evaluation of novel choices by simultaneous representation of multiple memories. *Nat. Neurosci.* 16(10):1492–1498

Behrens TEJ, Muller TH, Whittington JCR, Mark S, Baram AB, et al. 2018. What is a cognitive map? organizing knowledge for flexible behavior. *Neuron* 100(2):490–509

Blanco NJ, Love BC, Ramscar M, Otto AR, Smayda K, Maddox WT. 2016. Exploratory decision-making as a function of lifelong experience, not cognitive decline. *Journal of Experimental Psychology: General* 145(3):284

Blanco NJ, Sloutsky VM. 2021. Systematic exploration and uncertainty dominate young children's choices. *Developmental Science* 24(2):e13026

Bonawitz E, Denison S, Gopnik A, Griffiths TL. 2014. Win-stay, lose-sample: A simple sequential algorithm for approximating bayesian inference. *Cognitive Psychology* 74:35–65

Botvinick M, Ritter S, Wang JX, Kurth-Nelson Z, Blundell C, Hassabis D. 2019. Reinforcement learning, fast and slow. *Trends Cogn. Sci.* 23(5):408–422

Bowman CR, Iwashita T, Zeithamova D. 2020. Tracking prototype and exemplar representations in the brain across learning. *eLife* 9:e59360

Brehmer B. 1974. Hypotheses about relations between scaled variables in the learning of probabilistic inference tasks. *Organizational Behavior and Human Performance* 11(1):1–27

Brehmer B. 1976. Learning complex rules in probabilistic inference tasks. *Scand. J. Psychol.* 17(1):309–312

Breit S, Sakaki M, Murayama K, Wu CM. 2022. In search of lost memories: modeling forgetful generalization, In *Proceedings of the 15th Biannual Conference of the German Society for Cognitive Science*

Bruner JS, Goodnow JJ, Austin GA. 1956. A Study of Thinking. New York: John Wiley & Sons, Inc

Buchsbaum D, Bridgers S, Skolnick Weisberg D, Gopnik A. 2012. The power of possibility: Causal learning, counterfactual reasoning, and pretend play. *Philosophical Transactions of the Royal Society B: Biological Sciences* 367(1599):2202–2212

Busemeyer JR, Byun E, Delosh EL, McDaniel MA. 1997. Learning functional relations based on experience with input-output pairs by humans and artificial neural networks. In *Knowledge, Concepts and Categories*, vol. 464. Cambridge, MA, US: The MIT Press, xiii, 408–437

Carroll JD. 1963. Functional learning: The learning of continuous functional mappings relating stimulus and response continua. *ETS Research Bulletin Series* 1963(2):i–144

Chater N, Vitányi PMB. 2003. The generalized universal law of generalization. *J. Math. Psychol.* 47(3):346–369

Chollet F. 2019. On the measure of intelligence. *arXiv preprint arXiv:1911.01547*

Cogliati Dezza I, Cleeremans A, Alexander W. 2019. Should we control? the interplay between cognitive control and information integration in the resolution of the exploration-exploitation dilemma. *Journal of Experimental Psychology: General* 148(6):977

Constantinescu AO, O'Reilly JX, Behrens TEJ. 2016. Organizing conceptual knowledge in humans with a gridlike code. *Science* 352(6292):1464–1468

Csibra G, Gergely G. 2009. Natural pedagogy. *Trends in Cognitive Sciences* 13(4):148–153

Cushman F. 2020. Rationalization is rational. *Behav. Brain Sci.* 43:e28

Cybenko G. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems* 2(4):303–314

Dasgupta I, Grant E, Griffiths T. 2022. Distinguishing rule and exemplar-based generalization in learning systems, In *Proceedings of the 39th International Conference on Machine Learning*, eds. K Chaudhuri, S Jegelka, L Song, C Szepesvari, G Niu, S Sabato, vol. 162 of *Proceedings of Machine Learning Research*, pp. 4816–4830, PMLR

Dasgupta I, Schulz E, Goodman ND, Gershman SJ. 2018. Remembrance of inferences past: Amortization in human hypothesis generation. *Cognition* 178:67–81

Dasgupta I, Schulz E, Tenenbaum JB, Gershman SJ. 2020. A theory of learning to infer. *Psychological review* 127(3):412

Dayan P. 1993. Improving generalization for temporal difference learning: The successor representation

Dehaene S, Al Roumi F, Lakretz Y, Planton S, Sablé-Meyer M. 2022. Symbols and mental programs: a hypothesis about human singularity. *Trends in Cognitive Sciences*

Deisenroth M, Rasmussen CE. 2011. PILCO: A model-based and data-efficient approach to policy search, In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pp. 465–472

DeLosh EL, Busemeyer JR, McDaniel MA. 1997. Extrapolation: the sine qua non for abstraction in function learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 23(4):968–986

Denison S, Bonawitz E, Gopnik A, Griffiths TL. 2013. Rational variability in children's causal inferences: The sampling hypothesis. *Cognition* 126(2):285–300

Doolittle WF, Bapteste E. 2007. Pattern pluralism and the tree of life hypothesis. *Proceedings of the National Academy of Sciences* 104(7):2043–2049

Doucet A, Johansen AM, Others. 2009. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of nonlinear filtering* 12(656-704):3

Duvenaud D, Lloyd JR, Grosse R, Tenenbaum JB, Ghahramani Z. 2013. Structure discovery in nonparametric regression through compositional kernel search

Ekman G. 1954. Dimensions of color vision. *J. Psychol.* 38(2):467–474

Ellis K, Wong L, Nye M, Sable-Meyer M, Cary L, et al. 2023. Dreamcoder: growing generalizable, interpretable knowledge with wake–sleep bayesian program learning. *Philosophical Transactions of the Royal Society A* 381(2251):20220050

Epstein RA, Patai EZ, Julian JB, Spiers HJ. 2017. The cognitive map in humans: spatial navigation and beyond. *Nat. Neurosci.* 20(11):1504–1513

Erickson MA, Kruschke JK. 1998. Rules and exemplars in category learning. *J. Exp. Psychol. Gen.*

127(2):107–140
Fodor JA. 1975. The language of thought, vol. 5. Harvard university press
Fränken JP, Theodoropoulos NC, Bramley NR. 2022. Algorithms of adaptation in inductive inference. *Cognitive Psychology* 137:101506
Garvert MM, Dolan RJ, Behrens TEJ. 2017. A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *Elife* 6:e17086
Garvert MM, Saanum T, Schulz E, Schuck NW, Doeller CF. 2023. Hippocampal spatio-predictive cognitive maps adaptively guide reward generalization. *Nature Neuroscience* 26(4):615–626
Geerts JP, Stachenfeld KL, Burgess N. 2019. Probabilistic successor representations with kalman temporal differences
Geirhos R, Medina Temme CR, Rauber J, Schütt HH, Bethge M, Wichmann FA. 2018. Generalisation in humans and deep neural networks
Gershman SJ. 2018. Deconstructing the human algorithms for exploration. *Cognition* 173:34–42
Gershman SJ, Daw ND. 2017. Reinforcement learning and episodic memory in humans and animals: An integrative framework. *Annu. Rev. Psychol.* 68:101–128
Gershman SJ, Malmaud J, Tenenbaum JB. 2017. Structured representations of utility in combinatorial domains. *Decisions* 4(2):67–86
Giron AP, Ciranka S, Schulz E, van den Bos W, Ruggeri A, et al. 2023. Developmental changes in exploration resemble stochastic optimization. *Nature Human Behaviour* :1–13
Goodman N. 1972. Seven strictures on similarity. In *Problems and Projects*. Bobs-Merril
Goodman ND, Tenenbaum JB, Feldman J, Griffiths TL. 2008. A rational analysis of rule-based concept learning. *Cogn. Sci.* 32(1):108–154
Gopnik A. 2020. Childhood as a solution to explore–exploit tensions. *Philosophical Transactions of the Royal Society B* 375(1803):20190502
Gopnik A, Griffiths TL, Lucas CG. 2015. When younger learners can be better (or at least more open-minded) than older ones. *Current Directions in Psychological Science* 24(2):87–92
Gopnik A, O'Grady S, Lucas CG, Griffiths TL, Wente A, et al. 2017. Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences* 114(30):7892–7899
Griffiths TL, Lucas CG, Williams JJ, Kalish ML. 2008. Modeling human function learning with gaussian processes, In *Proceedings of the 21st International Conference on Neural Information Processing Systems*, pp. 553–560
Griffiths TL, Tenenbaum JB. 2005. Structure and strength in causal induction. *Cogn. Psychol.* 51(4):334–384
Griffiths TL, Tenenbaum JB. 2009a. Theory-based causal induction. *Psychol. Rev.* 116(4):661–716
Griffiths TL, Tenenbaum JB. 2009b. Theory-based causal induction. *Psychological review* 116(4):661–716
Hahn U, Ramscar M. 2001. Similarity and categorization. Oxford University Press
Henrich J. 2016. The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter. princeton University press
Heyes C. 2018. Cognitive gadgets: The cultural evolution of thinking. Harvard University Press
Ho MK, MacGlashan J, Littman ML, Cushman F. 2017. Social is special: A normative framework for teaching with and learning from evaluative feedback. *Cognition* 167:91–106
Hornik K. 1991. Approximation capabilities of multilayer feedforward networks. *Neural networks* 4(2):251–257
Huys QJ, Lally N, Faulkner P, Eshel N, Seifritz E, et al. 2015. Interplay of approximate planning strategies. *Proceedings of the National Academy of Sciences* 112(10):3098–3103
Jäkel F, Schölkopf B, Wichmann FA. 2008a. Generalization and similarity in exemplar models of categorization: insights from machine learning. *Psychon. Bull. Rev.* 15(2):256–271
Jäkel F, Schölkopf B, Wichmann FA. 2008b. Similarity, kernels, and the triangle inequality. *Journal of Mathematical Psychology* 52(5):297–303

James W. 1890. The principles of psychology

Jara-Ettinger J. 2019. Theory of mind as inverse reinforcement learning. *Current Opinion in Behavioral Sciences* 29:105–110

Johansen MK, Palmeri TJ. 2002. Are there representational shifts during category learning? *Cognitive Psychology* 45(4):482–553

Johansson R, Johansson M. 2014. Look here, eye movements play a functional role in memory retrieval. *Psychological Science* 25(1):236–242

Kalish ML, Griffiths TL, Lewandowsky S. 2007. Iterated learning: intergenerational knowledge transmission reveals inductive biases. *Psychon. Bull. Rev.* 14(2):288–294

Kalish ML, Lewandowsky S, Kruschke JK. 2004. Population of linear experts: knowledge partitioning and function learning. *Psychol. Rev.* 111(4):1072–1099

Ke NR, Chiappa S, Wang J, Goyal A, Bornschein J, et al. 2022. Learning to induce causal structure. *arXiv preprint arXiv:2204.04875*

Kemp C, Tenenbaum JB. 2008. The discovery of structural form. *Proc. Natl. Acad. Sci. U. S. A.* 105(31):10687–10692

Kemp C, Tenenbaum JB. 2009. Structured statistical models of inductive reasoning. *Psychological Review* 116(1):20–58

Keramati M, Smittenaar P, Dolan RJ, Dayan P. 2016. Adaptive integration of habits into depth-limited planning defines a habitual-goal–directed spectrum. *Proceedings of the National Academy of Sciences* 113(45):12868–12873

Koh K, Meyer DE. 1991. Function learning: Induction of continuous stimulus-response relations

Kondor RI, Lafferty J. 2002. Diffusion kernels on graphs and other discrete input spaces, In *Proceedings of the 19th International Conference on Machine Learning*, pp. 3121–3322

Kool W, Gershman SJ, Cushman FA. 2018. Planning complexity registers as a cost in metacontrol. *Journal of cognitive neuroscience* 30(10):1391–1404

Kotov R, Krueger RF, Watson D, Achenbach TM, Althoff RR, et al. 2017. The hierarchical taxonomy of psychopathology (HiTOP): A dimensional alternative to traditional nosologies. *Journal of Abnormal Psychology* 126(4):454–477

Kotov R, Krueger RF, Watson D, Cicero DC, Conway CC, et al. 2021. The hierarchical taxonomy of psychopathology (HiTOP): A quantitative nosology based on consensus of evidence. *Annual Review of Clinical Psychology* 17:83–108

Kruschke JK. 1992. ALCOVE: An exemplar-based connectionist model of category learning. *Psychol. Rev.*

Kruskal JB. 1964. Nonmetric multidimensional scaling: A numerical method

Kwantes PJ, Neal A. 2006. Why people underestimate y when extrapolating in linear functions. *J. Exp. Psychol. Learn. Mem. Cogn.* 32(5):1019–1030

Lake BM, Salakhutdinov R, Tenenbaum JB. 2015. Human-level concept learning through probabilistic program induction. *Science* 350(6266):1332–1338

Lake BM, Ullman TD, Tenenbaum JB, Gershman SJ. 2017. Building machines that learn and think like people. *Behav. Brain Sci.* 40:e253

Lau T, Gershman SJ, Cikara M. 2020. Social structure learning in human anterior insula. *Elife* 9

Le TA, Collins KM, Hewitt L, Ellis K, Siddharth N, et al. 2021. Hybrid memoised Wake-Sleep: Approximate inference at the Discrete-Continuous interface

Lengyel M, Dayan P. 2007. Hippocampal contributions to control: the third way. *Advances in neural information processing systems* 20

Love BC, Medin DL, Gureckis TM. 2004. SUSTAIN: a network model of category learning. *Psychol. Rev.* 111(2):309–332

Lucas CG, Bridgers S, Griffiths TL, Gopnik A. 2014. When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition* 131(2):284–299

Lucas CG, Griffiths TL, Williams JJ, Kalish ML. 2015. A rational model of function learning.

Psychon. Bull. Rev. 22(5):1193–1215
Ludwig T, Wu CM, Schulz E. 2022. Connecting exploration, generalization, and planning in correlated trees, In *Proceedings of the 44th Annual Conference of the Cognitive Science Society*, eds. J Culbertson, A Perfors, H Rabagliati, V Ramenzoni. Toronto, Canada: Cognitive Science Society
Lynn CW, Bassett DS. 2020. How humans learn and represent networks. *Proceedings of the National Academy of Sciences* 117(47):29407–29415
Machado MC, Rosenbaum C, Guo X, Liu M, Tesauro G, Campbell M. 2018. Eigenoption discovery through the deep successor representation, In *International Conference on Learning Representations*
Madarasz T, Behrens T. 2019. Better transfer learning with inferred successor maps, In *Advances in Neural Information Processing Systems*, eds. H Wallach, H Larochelle, A Beygelzimer, F dAlché-Buc, E Fox, R Garnett, vol. 32. Curran Associates, Inc.
Manns JR, Eichenbaum H. 2006. Evolution of declarative memory. *Hippocampus* 16(9):795–808
Mark S, Moran R, Parr T, Kennerley SW, Behrens TEJ. 2020. Transferring structural knowledge across cognitive maps in humans and models. *Nat. Commun.* 11(1):4783
McClelland JL, Rumelhart DE, Group PR, Others. 1986. Parallel distributed processing, vol. 2. MIT press Cambridge, MA
McDaniel MA, Busemeyer JR. 2005. The conceptual basis of function learning and extrapolation: comparison of rule-based and associative-based models. *Psychon. Bull. Rev.* 12(1):24–42
Meagher BJ, Nosofsky RM. 2023. Testing formal cognitive models of classification and old-new recognition in a real-world high-dimensional category domain. *Cognitive Psychology* 145:101596
Meder B, Mayrhofer R, Waldmann MR. 2014. Structure induction in diagnostic causal reasoning. *Psychol. Rev.* 121(3):277–301
Meder B, Wu CM, Schulz E, Ruggeri A. 2021. Development of directed and random exploration in children. *Developmental Science* 24(4):e13095
Medin DL, Goldstone RL, Gentner D. 1993. Respects for similarity. *Psychol. Rev.* 100(2):254–278
Medin DL, Schaffer MM. 1978. Context theory of classification learning. *Psychol. Rev.* 85(3):207–238
Mehlhorn K, Newell BR, Todd PM, Lee MD, Morgan K, et al. 2015. Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decisions* 2(3):191–215
Mercer J. 1909. Functions of positive and negative type and their connection with the theory of integral equations. *Philosophical Transactions of the Royal Society A* 209:441–458
Miller KJ, Botvinick MM, Brody CD. 2017. Dorsal hippocampus contributes to model-based planning. *Nature neuroscience* 20(9):1269–1276
Mittelstaedt ML, Mittelstaedt H. 1980. Homing by path integration in a mammal
Momennejad I. 2020. Learning structures: Predictive representations, replay, and generalization. *Current Opinion in Behavioral Sciences* 32:155–166
Momennejad I, Russek EM, Cheong JH, Botvinick MM, Daw ND, Gershman SJ. 2017. The successor representation in human reinforcement learning. *Nat Hum Behav* 1(9):680–692
Moser EI, Roudi Y, Witter MP, Kentros C, Bonhoeffer T, Moser MB. 2014. Grid cells and cortical representation. *Nat. Rev. Neurosci.* 15(7):466–481
Moskvichev A, Odouard VV, Mitchell M. 2023. The conceptarc benchmark: Evaluating understanding and generalization in the arc domain. *arXiv preprint arXiv:2305.07141*
Murphy GL, Medin DL. 1985. The role of theories in conceptual coherence. *Psychol. Rev.* 92(3):289–316
Neal RM. 1996. Priors for infinite networks. *Bayesian learning for neural networks* :29–53
Nelson JD. 2005. Finding useful questions: on bayesian diagnosticity, probability, impact, and information gain. *Psychol. Rev.* 112(4):979–999
Newell A, Simon HA. 1976. Computer science as empirical inquiry: Symbols and search. *Commun. ACM* 19(3):113–126

Norbury A, Robbins TW, Seymour B. 2018. Value generalization in human avoidance learning. *Elife* 7

Nosofsky RM. 1986. Attention, similarity, and the identification–categorization relationship. *J. Exp. Psychol. Gen.* 115(1):39–57

Nosofsky RM, Palmeri TJ, McKinley SC. 1994. Rule-plus-exception model of classification learning. *Psychol. Rev.* 101(1):53–79

Pavlov IP. 1927. Conditional reflexes: an investigation of the physiological activity of the cerebral cortex. Oxford University Press

Peer M, Brunec IK, Newcombe NS, Epstein RA. 2021. Structuring knowledge with cognitive maps and cognitive graphs. *Trends Cogn. Sci.* 25(1):37–54

Pettine WW, Raman DV, Redish AD, Murray JD. 2023. Human generalization of internal representations through prototype learning with goal-directed attention. *Nature Human Behaviour* 7(3):442–463

Piantadosi ST, Tenenbaum JB, Goodman ND. 2016. The logical primitives of thought: Empirical foundations for compositional cognitive models. *Psychological review* 123(4):392

Poggio T, Bizzi E. 2004. Generalization in vision and motor control. *Nature* 431(7010):768–774

Pothos EM. 2005. The rules versus similarity distinction. *Behav. Brain Sci.* 28(1):1–14; discussion 14–49

Pouncy T, Tsividis P, Gershman SJ. 2021. What is the model in model-based planning? *Cogn. Sci.* 45(1):e12928

Radulescu A, Shin YS, Niv Y. 2021. Human representation learning. *Annu. Rev. Neurosci.*

Rasmussen CE, Williams CKI. 2005. Gaussian processes for machine learning

Reber AS, Lewis S. 1977. Implicit learning: An analysis of the form and structure of a body of tacit knowledge. *Cognition* 5(4):333–361

Rescorla RA, Wagner AR. 1972. A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning: current research and theory*, vol. 2

Rosch EH. 1973. Natural categories. *Cogn. Psychol.* 4(3):328–350

Rouder JN, Ratcliff R. 2006. Comparing exemplar-and rule-based theories of categorization. *Current Directions in Psychological Science* 15(1):9–13

Rubino V, Hamidi M, Dayan P, Wu CM. 2023. Compositionality under time pressure, In *Proceedings of the 45th Annual Conference of the Cognitive Science Society*, eds. M Goldwater, F Anggoro, B Hayes, D Ong. Sydney, Australia: Cognitive Science Society

Rule JS, Tenenbaum JB, Piantadosi ST. 2020. The child as hacker. *Trends in cognitive sciences* 24(11):900–915

Russek EM, Momennejad I, Botvinick MM, Gershman SJ, Daw ND. 2017a. Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS Comput. Biol.* 13(9):e1005768

Russek EM, Momennejad I, Botvinick MM, Gershman SJ, Daw ND. 2017b. Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS computational biology* 13(9):e1005768

Sanborn AN, Griffiths TL, Navarro DJ. 2010. Rational approximations to rational models: alternative algorithms for category learning. *Psychol. Rev.* 117(4):1144–1167

Schölkopf B. 2015. Learning to see and act. *Nature* 518(7540):486–487

Schölkopf B, Smola AJ. 2002. Learning with kernels: support vector machines, regularization, optimization, and beyond. MIT press

Scholz A, von Helversen B, Rieskamp J. 2015. Eye movements reveal memory processes during similarity- and rule-based decision making. *Cognition* 136:228–246

Schulz E, Speekenbrink M, Krause A. 2018a. A tutorial on gaussian process regression: Modelling, exploring, and exploiting functions. *J. Math. Psychol.* 85:1–16

Schulz E, Tenenbaum JB, Duvenaud D, Speekenbrink M, Gershman SJ. 2017. Compositional inductive biases in function learning. *Cogn. Psychol.* 99:44–79

Schulz E, Wu CM, Huys QJM, Krause A, Speekenbrink M. 2018b. Generalization and search in risky environments. *Cogn. Sci.* 42(8):2592–2620

Schulz E, Wu CM, Ruggeri A, Meder B. 2019. Searching for rewards like a child means less generalization and more directed exploration. *Psychol. Sci.* 30(11):1561–1572

Settles B. 2009. Active learning literature survey. Tech. Rep. 1648, University of Wisconsin–Madison

Shepard RN. 1962. The analysis of proximities: Multidimensional scaling with an unknown distance function. II

Shepard RN. 1987. Toward a universal law of generalization for psychological science. *Science* 237(4820):1317–1323

Silver D, Huang A, Maddison CJ, Guez A, Sifre L, et al. 2016. Mastering the game of go with deep neural networks and tree search. *Nature* 529(7587):484–489

Simon HA, Lea G. 1974. Problem solving and rule induction: A unified view. In *Knowledge and cognition*, ed. LW Gregg, vol. 321. Lawrence Erlbaum

Skinner BF. 1938. The behavior of organisms: An experimental analysis. Appleton-Century, New York

Smith EE, Medin DL. 1981. Categories and concepts

Smith JD, Minda JP. 1998. Prototypes in the mist: The early epochs of category learning. *J. Exp. Psychol. Learn. Mem. Cogn.* 24(6):1411–1436

Somerville LH, Sasse SF, Garrad MC, Drysdale AT, Abi Akar N, et al. 2017a. Charting the expansion of strategic exploratory behavior during adolescence. *J. Exp. Psychol. Gen.* 146(2):155–164

Somerville LH, Sasse SF, Garrad MC, Drysdale AT, Abi Akar N, et al. 2017b. Charting the expansion of strategic exploratory behavior during adolescence. *Journal of experimental psychology: general* 146(2):155

Speekenbrink M. 2016. A tutorial on particle filters. *J. Math. Psychol.* 73:140–152

Speekenbrink M, Konstantinidis E. 2015. Uncertainty and exploration in a restless bandit problem. *Topics in cognitive science* 7(2):351–367

Spivey MJ, Geng JJ. 2001. Oculomotor mechanisms activated by imagery and memory: Eye movements to absent objects. *Psychological Research* 65:235–241

Srinivas N, Krause A, Kakade SM, Seeger M. 2009. Gaussian process optimization in the bandit setting: No regret and experimental design

Stachenfeld KL, Botvinick MM, Gershman SJ. 2017. The hippocampus as a predictive map. *Nat. Neurosci.* 20(11):1643–1653

Stojić H, Schulz E, P Analytis P, Speekenbrink M. 2020. It's new, but is it good? how generalization and uncertainty guide the exploration of novel options. *J. Exp. Psychol. Gen.* 149(10):1878–1907

Sutton RS, Barto AG. 2018. Reinforcement learning, second edition: An introduction. MIT Press

Sutton RS, McAllester DA, Singh SP, Mansour Y. 2000. Policy gradient methods for reinforcement learning with function approximation, In *Advances in neural information processing systems*, pp. 1057–1063

Tavares RM, Mendelsohn A, Grossman Y, Williams CH, Shapiro M, et al. 2015. A map for social navigation in the human brain. *Neuron* 87(1):231–243

Taylor JE, Cortese A, Barron HC, Pan X, Sakagami M, Zeithamova D. 2021. How do we generalize?

Tenenbaum JB. 1999. Bayesian modeling of human concept learning. *Adv. Neural Inf. Process. Syst.* :59–68

Tenenbaum JB. 2000. Rules and similarity in concept learning. *Adv. Neural Inf. Process. Syst.* 12:59–65

Tenenbaum JB, Griffiths TL. 2001. Generalization, similarity, and bayesian inference. *Behav. Brain Sci.* 24(4):629–40; discussion 652–791

Tesauro G. 1995. Temporal difference learning and TD-Gammon. *Commun. ACM* 38(3):58–68

Thorndike EL. 1911. Animal intelligence: Experimental studies. Hafner, Darien, CT

Tolman EC. 1948. Cognitive maps in rats and men

Torgerson WS. 1952. Multidimensional scaling: I. theory and method

Tsividis PA, Loula J, Burga J, Foss N, Campero A, et al. 2021. Human-Level reinforcement learning through Theory-Based modeling, exploration, and planning

Tversky A. 1977. Features of similarity. *Psychol. Rev.* 84(4):327–352

Van den Bos W, Cohen MX, Kahnt T, Crone EA. 2012. Striatum–medial prefrontal cortex connectivity predicts developmental changes in reinforcement learning. *Cerebral cortex* 22(6):1247–1255

van Rooij I, Blokpoel M, Kwisthout J, Wareham T. 2019. Cognition and intractability: A guide to classical and parameterized complexity analysis. Cambridge University Press

Wang R, Salakhutdinov R, Yang LF. 2020. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension

Whittington JC, McCaffary D, Bakermans JJ, Behrens TE. 2022. How to build a cognitive map. *Nature neuroscience* 25(10):1257–1272

Whittington JCR, Muller TH, Mark S, Chen G, Barry C, et al. 2020. The Tolman-Eichenbaum machine: Unifying space and relational memory through generalization in the hippocampal formation. *Cell* 183(5):1249–1263.e23

Wilson RC, Geana A, White JM, Ludvig EA, Cohen JD. 2014. Humans use directed and random exploration to solve the explore-exploit dilemma. *J. Exp. Psychol. Gen.* 143(6):2074–2081

Wimmer GE, Elliott Wimmer G, Daw ND, Shohamy D. 2012. Generalization of value in reinforcement learning by humans

Wingate D, Diuk C, O'Donnell T, Tenenbaum J, Gershman S. 2013. Compositional policy priors. Tech. Rep. MIT-CSAIL-TR-2013-007, MIT

Wise T, Emery K, Radulescu A. 2024. Naturalistic reinforcement learning. *Trends in Cognitive Sciences* 28(2):144–158

Witt A, Toyokawa W, Lala K, Gaissmaier W, Wu CM. 2023. Social learning with a grain of salt, In *Proceedings of the 45th Annual Conference of the Cognitive Science Society*, eds. M Goldwater, F Anggoro, B Hayes, D Ong. Sydney, Australia: Cognitive Science Society

Wu CM, Dale R, Hawkins RD. 2023a. Group coordination catalyzes individual and cultural intelligence. *PsyArXiv*

Wu CM, Deffner D, Kahl B, Meder B, Ho MH, Kurvers RH. 2023b. Visual-spatial dynamics drive adaptive social learning in immersive environments. *bioRxiv*

Wu CM, Schulz E, Garvert MM, Meder B, Schuck NW. 2020. Similarities and differences in spatial and non-spatial cognitive maps. *PLoS Comput. Biol.* 16(9):e1008149

Wu CM, Schulz E, Gershman SJ. 2021. Inference and search on Graph-Structured spaces. *Computational Brain & Behavior* 4(2):125–147

Wu CM, Schulz E, Pleskac TJ, Speekenbrink M. 2022a. Time pressure changes how people explore and respond to uncertainty. *Scientific Reports* 12:1–14

Wu CM, Schulz E, Speekenbrink M, Nelson JD, Meder B. 2018. Generalization guides human exploration in vast decision spaces. *Nat Hum Behav* 2(12):915–924

Wu CM, Vélez N, Cushman FA. 2022b. Representational exchange in human social learning: Balancing efficiency and flexibility. In *The Drive for Knowledge: The Science of Human Information-Seeking*, eds. IC Dezza, E Schulz, CM Wu. Cambridge: Cambridge University Press

Xu F, Tenenbaum JB. 2007. Word learning as bayesian inference. *Psychol. Rev.* 114(2):245–272

Zajkowski WK, Kossut M, Wilson RC. 2017. A causal role for right frontopolar cortex in directed, but not random, exploration. *Elife* 6:e27430

Zhang C, Bengio S, Hardt M, Recht B, Vinyals O. 2016. Understanding deep learning requires rethinking generalization

Zhao B, Lucas CG, Bramley NR. 2023. A model of conceptual bootstrapping in human cognition. *Nature Human Behaviour* :1–12

Zhu X, Lafferty J, Ghahramani Z. 2003. Semi-Supervised learning: From gaussian fields to gaussian processes. Tech. Rep. CMU-CS-03-175, Carnegie Mellon University