

Meta-Learned Models of Cognition

Marcel Binz

Max Planck Institute for Biological Cybernetics

Ishita Dasgupta

DeepMind

Akshay Jagadish

Max Planck Institute for Biological Cybernetics

Matthew Botvinick

DeepMind

Jane X. Wang

DeepMind

Eric Schulz

Max Planck Institute for Biological Cybernetics

Author Note

Correspondence concerning this article should be addressed to Marcel Binz, MPRG Computational Principles of Intelligence, Max Planck Institute for Biological Cybernetics, Max Planck Ring 8, 72076 Tübingen, Germany. E-mail: marcel.binz@tue.mpg.de

Marcel Binz  <https://orcid.org/0000-0001-8872-8386>

Abstract

Short: Meta-learning has established itself as a promising tool for building models of human cognition in the recent years. Yet, a coherent research program around meta-learned models of cognition is still missing. The purpose of the present article is to develop such a research program. We accomplish this by pointing out that meta-learning can be used to construct Bayes-optimal learning algorithms, allowing us to draw strong connections to the rational analysis of cognition. We then discuss several advantages of the meta-learning framework over traditional Bayesian methods and reexamine prior work in the context of these new insights.

Long: Meta-learning is a framework for learning learning algorithms through repeated interactions with an environment as opposed to designing them by hand. In recent years, this framework has established itself as a promising tool for building models of human cognition. Yet, a coherent research program around meta-learned models of cognition is still missing. The purpose of this article is to synthesize previous work in this field and establish such a research program. We rely on three key pillars to accomplish this goal. We first point out that meta-learning can be used to construct Bayes-optimal learning algorithms. This result not only implies that any behavioral phenomenon that can be explained by a Bayesian model can also be explained by a meta-learned model but also allows us to draw strong connections to the rational analysis of cognition. We then discuss several advantages of the meta-learning framework over traditional Bayesian methods. In particular, we argue that meta-learning can be applied to situations where Bayesian inference is impossible and that it enables us to make rational models of cognition more realistic, either by incorporating limited computational resources or neuroscientific knowledge. Finally, we reexamine prior studies from psychology and neuroscience that have applied meta-learning and put them into the context of these new insights. In summary, our work highlights that meta-learning considerably extends the scope of rational analysis and thereby of cognitive theories more generally.

Keywords: meta-learning, rational analysis, Bayesian inference, cognitive modeling, neural networks

Meta-Learned Models of Cognition

It is hard to imagine cognitive psychology and neuroscience without computational models – they are invaluable tools to study, analyze, and understand the human mind. Traditionally, such computational models have been hand-designed by expert researchers. In a cognitive architecture, for instance, researchers provide a fixed set of structures and a definition of how these structures interact with each other (Anderson, 2013b). In a Bayesian model of cognition, researchers instead specify a prior and a likelihood function which – in combination with Bayes’ rule – fully determine the model’s behavior (L Griffiths, Kemp, & B Tenenbaum, 2008). The framework of meta-learning (Bengio, Bengio, & Cloutier, 1991; Schmidhuber, 1987; Thrun & Pratt, 1998) offers a radically different approach for constructing computational models by learning them through repeated interactions with an environment instead of requiring an a priori specification from a researcher.

Recently, psychologists have started to apply meta-learning to the study of human learning (Griffiths et al., 2019). It has been shown that meta-learned models can capture a wide range of empirically observed phenomena that could not be explained otherwise. They, amongst others, reproduce human biases in probabilistic reasoning (Dasgupta, Schulz, Tenenbaum, & Gershman, 2020), discover heuristic decision-making strategies used by people (Binz, Gershman, Schulz, & Endres, 2022), and generalize compositionally on complex language tasks in a human-like manner (Lake, 2019). The goal of the present article is to develop a research program around meta-learned models of cognition and, in doing so, offer a synthesis of previous work and outline new research directions.

To establish such a research program, we will make use of a recent result from the machine learning community showing that meta-learning can be used to construct Bayes-optimal learning algorithms (Mikulik, Delétang, et al., 2020; Ortega et al., 2019; Rabinowitz, 2019). This correspondence is interesting from a psychological perspective because it allows us to connect meta-learning to another already well-established

framework: the rational analysis of cognition (Anderson, 2013a; Chater & Oaksford, 1999). In a rational analysis, one first has to specify the goal of an agent along with a description of the environment the agent interacts with. The Bayes-optimal solution for the task at hand is then derived based on these assumptions and tested against empirical data. If needed, assumptions are modified and the whole process is repeated. This approach for constructing cognitive models has had a tremendous impact on psychology because it explains “why cognition works, by viewing it as an approximation to ideal statistical inference given the structure of natural tasks and environments” (Tenenbaum, 2021). The observation that meta-learned models can implement Bayesian inference implies that a meta-learned model can be used as a replacement for the corresponding Bayesian model in a rational analysis and thus suggests that any behavioral phenomenon that can be captured by a Bayesian model can also be captured by a meta-learned model.

We start our article by presenting a simplified version of an argument originally formulated by Ortega et al. (2019) and thereby make their result accessible to a broader audience. Having established that meta-learning produces models that can simulate Bayesian inference, we go on to discuss what additional explanatory power the meta-learning framework offers. After all, why should one not just stick to the tried-and-tested Bayesian approach? We answer this question by providing four original arguments in favor of the meta-learning framework (see Figure 1 for a visual synopsis):

1. Meta-learning can produce approximately optimal learning algorithms even if exact Bayesian inference is computationally intractable.
2. Meta-learning can produce approximately optimal learning algorithms even if it is not possible to phrase the corresponding inference problem in the first place.
3. Meta-learning makes it easy to manipulate a learning algorithm’s complexity and can therefore be used to construct resource-rational models of learning.
4. Meta-learning allows us to integrate neuroscientific insights into the rational analysis

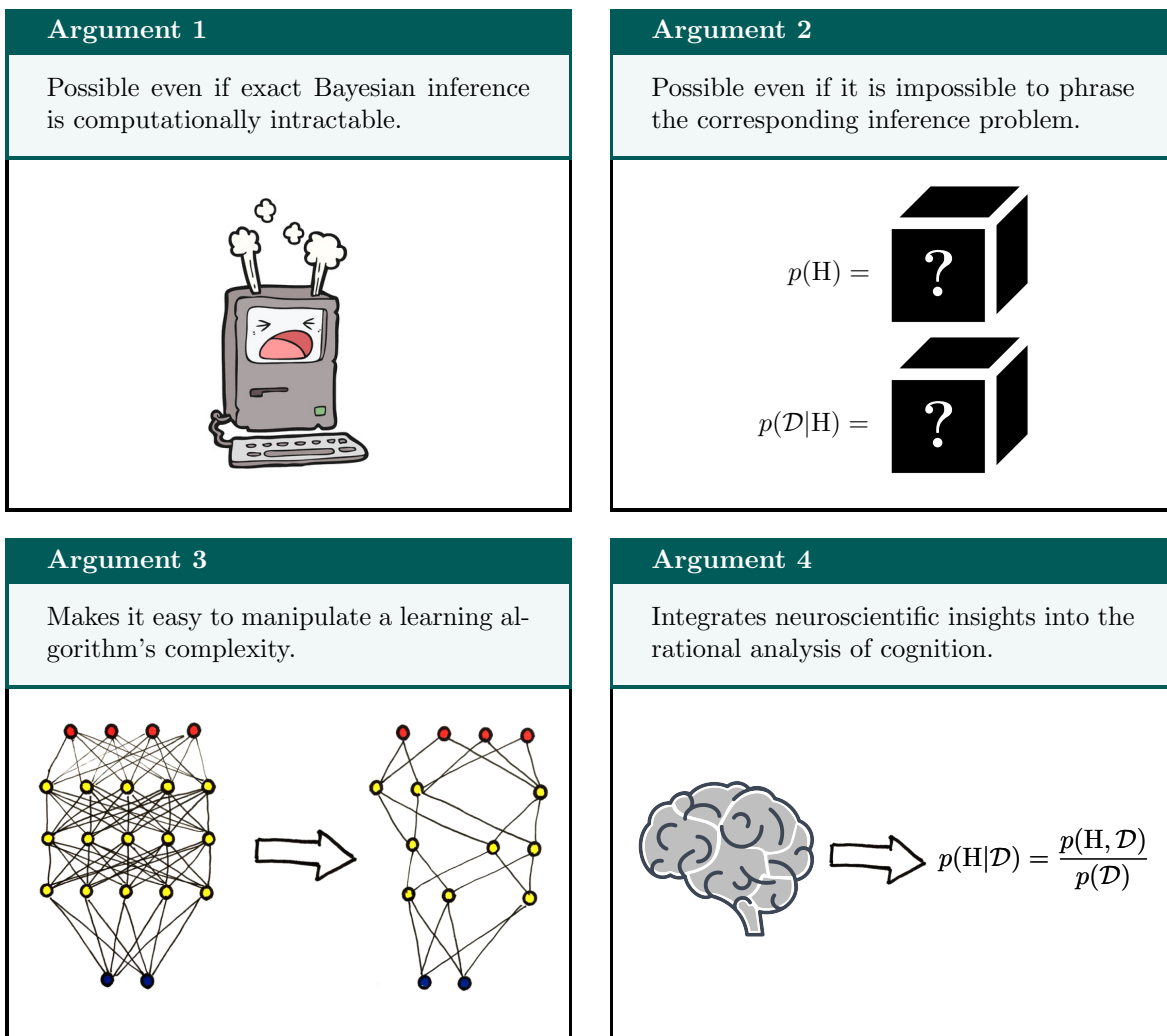


Figure 1
Visual synopsis of the four different arguments for meta-learning over Bayesian inference put forward in this article.

of cognition by incorporating these insights into model architectures.

The first two points highlight situations in which meta-learned models can be used for rational analysis but traditional Bayesian models cannot. The latter two points provide examples of how meta-learning enables us to make rational models of cognition more realistic, either by incorporating limited computational resources or neuroscientific insights. Taken together, these arguments showcase that meta-learning considerably extends the scope of rational analysis and thereby of cognitive theories more generally.

We will discuss each of these four points in detail and provide illustrations to highlight their relevance. We then reexamine prior studies from psychology and neuroscience that have applied meta-learning and put them into the context of our newly-acquired insights. For each of the reviewed studies, we highlight how it relates to the four presented arguments, and discuss why its findings could not have been obtained using a classical Bayesian model. Following that, we describe under which conditions traditional models are preferable to those obtained by meta-learning. We finish our article by speculating what the future holds for meta-learning. Therein, we focus on how meta-learning could be the key to building a domain-general model of human cognition.

1. Meta-Learned Rationality

The prefix *meta-* is generally used in a self-referential sense: a meta-rule is a rule about rules, a meta-discussion is a discussion about discussions, and so forth. Meta-learning, consequently, refers to learning about learning. We, therefore, need to first establish a common definition of *learning* before covering meta-learning in more detail. For the present article, we adopt the following definition from Mitchell (1997):

Definition: Learning

“For a given task, training experience, and performance measure, an algorithm is said to learn if its performance at the task improves with experience.”

To illustrate this definition, consider the following example which we will return to throughout the text: you are a biologist who has just discovered a new insect species and now set yourself the task of predicting how large members of this species are. You have already observed three exemplars in the wild with lengths of 16cm, 12cm, and 15cm respectively. This data amounts to your training experience. Ideally, you can use this experience to make better predictions about the length of the next individual you encounter. You are said to have learned something if your performance is better after seeing the data than it was before. Typical performance measures for this example problem

include the mean squared error or the (negative) log-likelihood.

1.1. Bayesian Inference for Rational Analyses

In a rational analysis of cognition, researchers are trying to compare human behavior to that of an optimal learning algorithm. However, it turns out that no learning algorithm is better than another when averaged over all possible problems (Wolpert, 1996; Wolpert & Macready, 1997), which means that we first have to make additional assumptions about the to-be-solved problem to obtain a well-defined notion of optimality. For our running example, one may make the following – somewhat unrealistic – assumptions:

1. Each observed insect length x_k is sampled from a normal distribution with mean μ and standard deviation σ .
2. An insect species' mean length μ cannot be observed directly, but the standard deviation σ is known to be 2cm.
3. Mean lengths across all insect species are distributed according to a normal distribution with a mean of 10cm and a standard deviation of 3cm.

An optimal way of making predictions about new observations under such assumptions is specified by Bayesian inference. Bayesian inference requires access to a prior distribution $p(\mu)$ that defines an agent's initial beliefs about possible parameter values before observing any data and a likelihood $p(x_{1:t}|\mu)$ that captures the agent's knowledge about how data is generated for a given set of parameters. In our running example, the prior and the likelihood can be identified as follows:

$$p(\mu) = \mathcal{N}(\mu; 10, 3) \tag{1}$$

$$p(x_{1:t}|\mu) = \prod_{k=1}^t p(x_k|\mu) = \prod_{k=1}^t \mathcal{N}(x_k; \mu, 2) \tag{2}$$

where $x_{1:t} = x_1, x_2, \dots, x_t$ denotes a sequence of observed insect lengths and the product in Equation 2 arises due to the additional assumption that observations are independent given the parameters.

The outcome of Bayesian inference is a posterior predictive distribution $p(x_{t+1}|x_{1:t})$, which the agent can use to make probabilistic predictions about a hypothetical future observation. To obtain this posterior predictive distribution, the agent first combines prior and likelihood into a posterior distribution over parameters by applying Bayes' theorem:

$$p(\mu|x_{1:t}) = \frac{p(x_{1:t}|\mu)p(\mu)}{\int p(x_{1:t}|\mu)p(\mu)d\mu} \quad (3)$$

In a subsequent step, the agent then averages over all possible parameter values weighted by their posterior probability to get the posterior predictive distribution:

$$p(x_{t+1}|x_{1:t}) = \int p(x_{t+1}|\mu)p(\mu|x_{1:t})d\mu \quad (4)$$

Multiple arguments justify Bayesian inference as a normative procedure, and thereby its use for rational analyses (Corner & Hahn, 2013). This includes dutch book arguments (Lewis, 1999; M. Rescorla, 2020), free energy minimization (Friston, 2010; Hinton & Van Camp, 1993), and performance-based justifications (Aitchison, 1975; Rosenkrantz, 1992). For this article, we are mainly interested in the latter class of performance-based justifications because these can be used – as we will demonstrate later on – to derive meta-learning algorithms that learn approximations to Bayesian inference.

Performance-based justifications are based on the notion of frequentist statistics. They assert that no learning algorithm can be better than Bayesian inference on a certain performance measure. Particularly relevant for this article is a theorem first proved by Aitchison (1975). It states that the posterior predictive distribution is the distribution (from the set of all possible distributions \mathcal{Q}) that maximizes the log-likelihood of hypothetical future observations when averaged over the data-generating distribution

$$p(\mu, x_{1:t+1}) = p(\mu)p(x_{1:t+1}|\mu):$$

$$p(x_{t+1}|x_{1:t}) = \arg \max_{q \in \mathcal{Q}} \mathbb{E}_{p(\mu, x_{1:t+1})} [\log q(x_{t+1}|x_{1:t})] \quad (5)$$

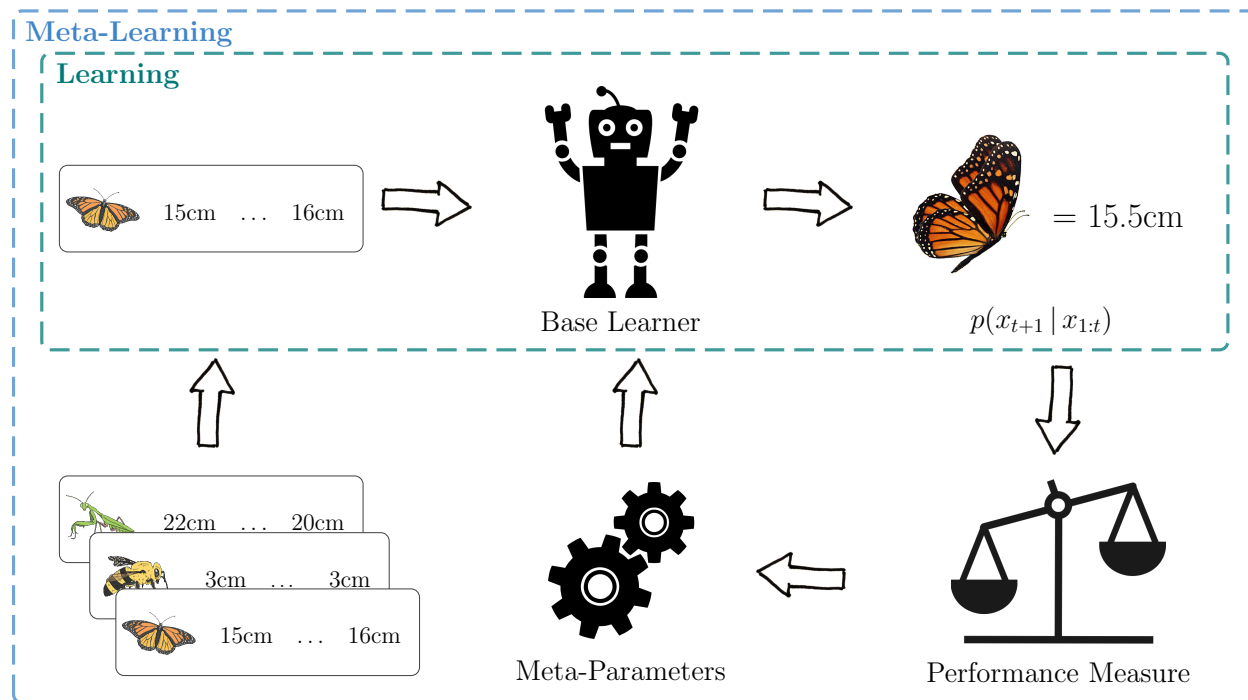
Equation 5 implies that if an agent wants to make a prediction about the length of a still unobserved exemplar from a particular insect species and measures its performance using the log-likelihood, then – averaged across all possible species that can be encountered – there is no better way of doing it than using the posterior predictive distribution. We decided to include a short proof of this theorem within Box 1 for the curious reader as it does not appear in popular textbooks on probabilistic machine learning (Bishop, 2006; Murphy, 2012) nor in survey articles on Bayesian models of cognition. Note that, while the theorem itself is central to our later argument, working through its proof is not required to follow the remainder of this article.

1.2. Meta-Learning

Having summarized the general concepts behind Bayes-optimal learning, we can now start to describe meta-learning in more detail. Formally speaking, a meta-learning algorithm is defined as any algorithm that “uses its experience to change certain aspects of a learning algorithm, or the learning method itself, such that the modified learner is better than the original learner at learning from additional experience” (Schaul & Schmidhuber, 2010).

To accomplish this, one first decides on an inner-loop (or base) learning algorithm and determines which of its aspects can be modified. We also refer to these modifiable aspects as meta-parameters. In an outer-loop (or meta-learning) process, the system is then trained on a series of learning problems such that the inner-loop learning algorithm gets better at solving the problems that it encounters. We provide a high-level overview of this framework in Figure 2.

The previous definition is quite broad and includes a variety of methods. It is, for

**Figure 2**

High-level overview of the meta-learning process. A base learner (green rectangle) receives data and performs some internal computations that improve its predictions on future data-points. A meta-learner (blue rectangle) encompasses a set of meta-parameters that can be adapted to create an improved learner. This is accomplished by training the learner on a distribution of related learning problems.

example, possible to meta-learn:

- Hyperparameters for a base learning algorithm, such as learning rates, batch sizes, or the number of training epochs (Doya, 2002; Feurer & Hutter, 2019; Li, Zhou, Chen, & Li, 2017).
- Initial parameters of a neural network that is trained via stochastic gradient descent (Finn, Abbeel, & Levine, 2017; Nichol, Achiam, & Schulman, 2018).
- Prior distributions in a probabilistic graphical model (Baxter, 1998; Grant, Finn, Levine, Darrell, & Griffiths, 2018).
- Entire learning algorithms (Hochreiter, Younger, & Conwell, 2001; Santoro,

Bartunov, Botvinick, Wierstra, & Lillicrap, 2016).

While all these methods have their own merits, we will be primarily concerned with the latter approach. Learning entire learning algorithms from scratch is arguably the most general and ambitious type of meta-learning, and it is the focus of this article because it is the only one among the aforementioned approaches leading to Bayes-optimal learning algorithms that can be utilized for rational analyses.

1.3. Meta-Learned Inference

It may seem like a daunting goal to learn an entire learning algorithm from scratch, but the core idea behind the approach we discuss in the following is surprisingly simple: instead of using Bayesian inference to obtain the posterior predictive distribution, we teach a general-purpose function approximator to do this inference. Previous work has mostly focused on using recurrent neural networks as function approximators in this setting and thus we will – without loss of generality – focus our upcoming exposition on this class of models.

Like the posterior predictive distribution, the recurrent neural network processes a sequence of observed length from a particular insect species and produces a predictive distribution over the lengths of potential future observations from the same species. More concretely, the meta-learned predictive distribution takes a predetermined functional form whose parameters are given by the network outputs. If we had, for example, decided to use a normal distribution as the functional form of the meta-learned predictive distribution, outputs of the network would correspond to an expected length m_{t+1} and its standard deviation s_{t+1} . Figure 3a illustrates this setup graphically.

Initially, the recurrent neural network implements a randomly initialized learning algorithm.¹ The goal of the meta-learning process is then to turn this system into an

¹ Based on our earlier definition, it is at this point strictly speaking not a learning algorithm at all as it does not improve with additional data.

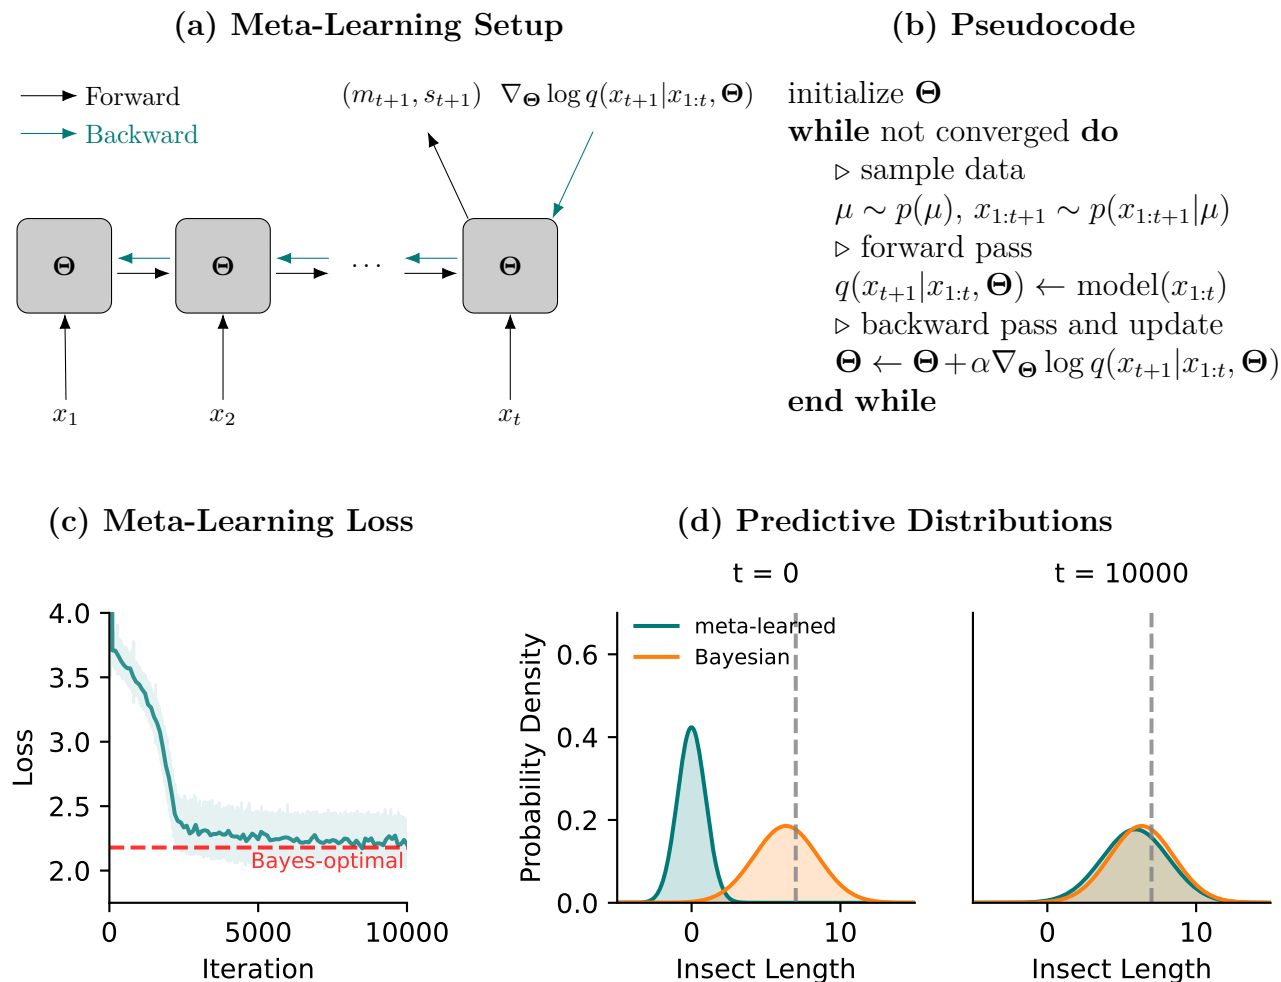


Figure 3
Meta-learning illustration. (a) A recurrent neural network processes a sequence of observations and produces a predictive distribution at the final time-step. (b) Pseudocode for a simple meta-learning algorithm. (c) Loss during meta-learning with shaded contours corresponding to the standard deviation across 30 runs. (d) Posterior and meta-learned predictive distributions for an example sequence at beginning and end of meta-learning. The dotted grey line denotes the (unobserved) mean length.

improved learning algorithm. The final result is a learning algorithm that is *learned* or trained rather than specified by a practitioner. To create a learning signal to do this training, we need a performance measure that can be used to optimize the network. Equation 5 suggests a straightforward strategy for designing such a measure by replacing the maximization over all possible distributions with a maximization over meta-parameters

Θ (in our case, the weights of the recurrent neural network):

$$\begin{aligned} & \arg \max_{q \in \mathcal{Q}} \mathbb{E}_{p(\mu, x_{1:t+1})} [\log q(x_{t+1} | x_{1:t})] \\ & \approx \arg \max_{\Theta} \mathbb{E}_{p(\mu, x_{1:t+1})} [\log q(x_{t+1} | x_{1:t}, \Theta)] \end{aligned} \quad (6)$$

To turn this expression into a practical meta-learning algorithm, we will – as common practice when training deep neural networks – maximize a sample-based version using stochastic gradient ascent:

$$\begin{aligned} & \arg \max_{\Theta} \mathbb{E}_{p(\mu, x_{1:t+1})} [\log q(x_{t+1} | x_{1:t}, \Theta)] \\ & \approx \arg \max_{\Theta} \frac{1}{N} \sum_{n=1}^N \log q(x_{t+1}^{(n)} | x_{1:t}^{(n)}, \Theta) \end{aligned} \quad (7)$$

Figure 3b presents pseudocode for a simple gradient-based procedure that maximizes Equation 7. The entire meta-learning algorithm can be implemented in just around 30 lines of self-contained `PyTorch` code (Paszke et al., 2019). We provide an annotated reference implementation on this article’s accompanying github repository.²

1.4. How Good Is a Meta-Learned Algorithm?

We have previously shown that the global optimum of Equation 7 is achieved by the posterior predictive distribution. Thus, by maximizing this performance measure, the network is actively encouraged to implement an approximation to exact Bayesian inference. Importantly, after the completion of meta-learning, producing an approximation to the posterior predictive distribution does not require any further updates to the network weights. To perform an inference, we simply have to query the network’s outputs after providing it with a particular sequence of observations.

If we want to use the fully optimized network for rational analyses, we have to ask

² <https://github.com/marcelbinz/meta-learned-models>

ourselves: how well does the resulting model approximate Bayesian inference? Two aspects have to be considered when answering this question. First, the network has to be sufficiently expressive to produce the exact posterior predictive distribution for all input sequences. Neural networks of sufficient width are universal function approximators (Hornik, Stinchcombe, & White, 1989), meaning that they can approximate any continuous function to arbitrary precision. Therefore, this aspect is not too problematic for the optimality argument. The second aspect is a bit more intricate: assuming that the network is powerful enough to represent the global optimum of Equation 7, the employed optimization procedure also has to find it. While we are not aware of any theorem that could provide such a guarantee, in practice, it has been observed that meta-learning procedures similar to the one discussed here often lead to networks that closely approximate Bayesian inference (Mikulik, Delétang, et al., 2020; Rabinowitz, 2019). We provide a visualization demonstrating that the predictions of a meta-learned model closely resemble those of exact Bayesian inference for our insect length example in Figure 3c-d.

While our exposition in this section focused on the supervised learning case, the same ideas can also be readily extended to the reinforcement learning setting (Duan et al., 2016; Wang et al., 2016). Box 2 outlines the general ideas behind the meta-reinforcement learning framework.

1.5. Tool or Theory?

It is often not so trivial to separate meta-learning from normal learning. We believe that part of this confusion arises from an underspecification regarding what is being studied. In particular, the meta-learning framework provides opportunities to address two distinct research questions:

1. It can be used to study how people improve their learning abilities over time.
2. It can be used as a methodological tool to construct learning algorithms with the properties of interest (and thereafter compare the emerging learning algorithms to

human behavior).

Historically, behavioral psychologists have been mainly interested in the former aspect (Doya, 2002; Harlow, 1949). In the 1940s, for example, Harlow (1949) already studied how learning in monkeys improves over time. He found that they adapted their learning strategies after sufficiently many interactions with tasks that shared a common structure, thereby showing a learning-to-learn effect. By now, examples of this phenomenon have been found in many different species – including humans – across nature (Wang, 2021).

More recently, psychologists have started to view meta-learning as a methodological tool to construct approximations to Bayes-optimal learning algorithms (Binz et al., 2022; Kumar, Dasgupta, Cohen, Daw, & Griffiths, 2020a), and subsequently utilize the resulting algorithms to study human cognition. The key difference from the former approach is that, in this setting, one abstracts away from the process of meta-learning and instead focuses on its outcome. From this perspective, only the fully converged model is of interest. Importantly, this approach allows us to investigate human learning from a rational perspective since we have demonstrated that meta-learning can be used to construct approximations to Bayes-optimal learning.

We place an emphasis on the second aspect in the present article and advocate for using fully converged meta-learned algorithms – as replacements for the corresponding Bayesian models – for rational analyses of cognition.³ In the next section, we will outline several arguments that support this approach. However, it is important to mention that we believe that meta-learning can also be a valuable tool to understand the process of learning-to-learn itself. In this context, several intriguing questions arise: at what time scale does meta-learning take place in humans? How much of it is due to task-specific

³ There has been a longstanding conceptual debate in cognitive psychology on whether to view Bayesian models as normative standards or descriptive tools. We believe that this debate is beyond the scope of the current article and thus refer the reader to earlier work for an in-depth discussion (Griffiths, Chater, Norris, & Pouget, 2012; Jones & Love, 2011; Tauber, Navarro, Perfors, & Steyvers, 2017; Zednik & Jäkel, 2016). We only want to add that the framework outlined here is agnostic to this issue – meta-learned models may serve as both normative standards and descriptive tools.

adaptations? How much of it is based on evolutionary or developmental processes? While we agree that these are important questions, they are not the focus of this article.

2. Why Not Bayesian Inference?

We have just argued that it is possible to meta-learn Bayes-optimal learning algorithms. What are the implications of this result? If one has access to two different theories that make identical predictions, which of them should be preferred? Bayesian inference has already established itself as a valuable tool for building cognitive models in the recent decades. Thus, the burden of proof is arguably on the meta-learning framework. In this section, we provide four different arguments that highlight the advantages of meta-learning for building models of cognition. Many of these arguments are novel and have not been put forward explicitly in previous literature. The first two arguments highlight situations in which meta-learned models can be used for rational analysis but traditional Bayesian models cannot. The latter two provide examples of how meta-learning enables us to make rational models of cognition more realistic, either by incorporating limited computational resources or neuroscientific insights.

2.1. Intractable Inference

Argument 1

Meta-learning can produce approximately optimal learning algorithms even if exact Bayesian inference is computationally intractable.

Bayesian inference becomes intractable very quickly because the complexity of computing the normalization constant that appears in the denominator grows exponentially with the number of unobserved parameters. In addition, it is only possible to find a closed-form expression of the posterior distribution for certain combinations of prior and likelihood. In our running example, we assumed that both prior and likelihood follow a normal distribution, which, in turn, leads to a normally-distributed posterior. However, if

one would instead assume that the prior over mean length follows an exponential distribution – which is arguably a more sensible assumption as it enforces lengths to be positive – it becomes already impossible to find an analytical expression for the posterior distribution.

Researchers across disciplines have recognized these challenges and have, in turn, developed approaches that can approximate Bayesian inference without running into computational difficulties. Prime examples of this are variational inference (Jordan, Ghahramani, Jaakkola, & Saul, 1999) and Markov chain Monte-Carlo (MCMC) methods (Geman & Geman, 1984). In variational inference, one phrases inference as an optimization problem by positing a variational approximation whose parameters are fitted to minimize a divergence measure to the true posterior distribution. MCMC methods, on the other hand, draw samples from a Markov chain that has the posterior distribution as its equilibrium distribution. Previous research in cognitive science indicates that human learning shows characteristics of such approximations (Courville & Daw, 2008; Dasgupta, Schulz, & Gershman, 2017; Daw, Courville, & Dayan, 2008; A. N. Sanborn, Griffiths, & Navarro, 2010; A. N. Sanborn & Silva, 2013).

Meta-learned inference also never requires an explicit calculation of the exact posterior or posterior predictive distribution. Instead, it performs approximately optimal inference via a single forward pass through the network. Inference, in this case, is approximate because we had to determine a functional form for the predictive distribution. The chosen form may deviate from the true form of the posterior predictive distribution, which, in turn, leads to approximation errors.⁴ In some sense, this type of approximation is similar to variational inference: both approaches involve optimization and require one to define a functional form of the respective distribution. However, the optimization process in both approaches uses a different loss function and happens at different time scales.

⁴ In principle, one could select arbitrarily flexible functional forms, such as mixtures of normal distributions or discretized distributions with small bin sizes, which would reduce the accompanying approximation error.

While variational inference employs the negative evidence lower bound as its loss function, meta-learning directly maximizes for models which can be expected to generalize well to unseen observations (using the performance-based measure from Equation 5). Furthermore, meta-learned inference only involves optimization during the outer-loop meta-learning process but not during the actual learning itself. To update how a meta-learned model makes predictions in the light of new data, we only have to perform a simple forward pass through the network. In contrast to this, standard variational inference requires us to rerun the whole optimization process from scratch every time a new data point is observed.⁵

In summary, it is possible to meta-learn an approximately Bayes-optimal learning algorithm. If exact Bayesian inference is not tractable, such models are our best option for performing rational analyses. Yet, many other methods for approximate inference, such as variational inference and MCMC methods, also share this feature, and it will thus ultimately be an empirical question which of these approximations provides a better description of human learning.

2.2. Unspecified Problems

Argument 2

Meta-learning can produce optimal learning algorithms even if it is not possible to phrase the corresponding inference problem in the first place.

Bayesian inference is hard, but posing the correct inference problem can be even harder. What exactly do we mean by that? To perform Bayesian inference, we need to specify a prior and a likelihood. Together, these two objects fully specify the assumed data-generating distribution, and thus the inference problem. Ideally, the specified data-generating distribution should match how the environment actually generates its data. It is fairly straightforward to fulfill this requirement in artificial scenarios, but for many

⁵ This only holds for standard variational inference but not for more advanced methods that involve amortization such as variational autoencoders (Kingma & Welling, 2013).

real-world problems, it is not. Take for instance our running example: Does the prior over mean length really follow a normal distribution? If yes, what are the mean and variance of this distribution? Are the underlying parameters actually time-invariant or do they, for example, change based on seasons? None of these questions can be answered with certainty.

In his seminal work on Bayesian decision theory, Savage (1972) made the distinction between small and large world problems. A small world problem is one “in which all relevant alternatives, their consequences, and probabilities are known” (Gigerenzer & Gaissmaier, 2011). A large world problem, on the other hand, is one in which the prior, the likelihood, or both cannot be identified. Savage’s distinction between small and large worlds is relevant for the rational analysis of human cognition as its critics have pointed out that Bayesian inference only provides a justification for optimal reasoning in small world problems (Binmore, 2007) and that “very few problems of interest to the cognitive, behavioral, and social sciences can be said to satisfy [this] condition” (Brighton & Gigerenzer, 2012).

Identifying the correct set of assumptions becomes especially challenging once we deal with more complex problems. To illustrate this, consider a study conducted by Lucas, Griffiths, Williams, and Kalish (2015) who attempted to construct a Bayesian model of human function learning. Doing so required them to specify a prior over functions that people expect to encounter. Without direct access to such a distribution, they instead opted for a heuristic solution: 98.8% of functions are expected to be linear, 1.1% are expected to be quadratic, and 0.1% are expected to be non-linear. Empirically, this choice led to good results, but it is hard to justify from a rational perspective. We simply do not know the frequency with which these functions appear in the real world, nor whether the given selection fully covers the set of functions expected by participants.

There are also inference problems in which it is not possible to specify or compute the likelihood function. These problems have been studied extensively in the machine learning community under the names of simulation-based or likelihood-free inference

(Cranmer, Brehmer, & Louppe, 2020; Lueckmann, Boelts, Greenberg, Goncalves, & Macke, 2021). In this setting, it is typically assumed that we can sample data from the likelihood for a given parameter setting but that computing the corresponding likelihood is impossible. Take, for instance, our insect length example. It should be clear that an insect’s length does not only depend on its species’ mean but on many other factors such as climate, genetics, and the individual’s age. Even if all these factors were known, mapping them to a likelihood function does seem close to impossible. But, we can generate samples easily by observing insects in the wild. If we had access large database of insect length measurements for different species, this could be directly used to meta-learn an approximately Bayes-optimal learning algorithm for predicting their length, while circumventing an explicit definition of a likelihood function.

In cases where we do not have access to a prior or a likelihood, we can neither apply exact Bayesian inference nor approximate inference schemes such as variational inference or MCMC methods. In contrast to this, meta-learned inference does not require us to define the prior or the likelihood explicitly. It only demands samples from the data-generating distribution to meta-learn an approximately Bayes-optimal learning algorithm – a much weaker requirement (Müller, Hollmann, Arango, Grabocka, & Hutter, 2021). The ability to construct Bayes-optimal learning algorithms for large worlds problems is a unique feature of the meta-learning framework, and we believe that it could open up totally new avenues for constructing rational models of human cognition. To highlight one concrete example, it would be possible to take a collection of real-world decision-making tasks – such as the ones presented by Czerlinski, Gigerenzer, Goldstein, et al. (1999) – and use them to obtain a meta-learned agent that is adapted to the decision-making problems that people actually encounter in their everyday lives. This algorithm could then serve as a normative standard against which we can compare human decision-making.

2.3. Resource Rationality

Argument 3

Meta-learning makes it easy to manipulate a learning algorithm’s complexity and can therefore be used to construct resource-rational models of learning.

Bayesian inference has been successfully applied to model human behavior across a number of domains, including perception (Knill & Richards, 1996), motor control (Körding & Wolpert, 2004), everyday judgments (Griffiths & Tenenbaum, 2006), and logical reasoning (Oaksford, Chater, et al., 2007). Notwithstanding these success stories, there are also well-documented deviations from the notion of optimality prescribed by Bayesian inference. People, for example, underreact to prior information (Kahneman & Tversky, 1973), ignore evidence (Benjamin, 2019), and rely on heuristic decision-making strategies (Gigerenzer & Gaissmaier, 2011).

The intractability of Bayesian inference – together with empirically observed deviations from it – has led researchers to conjecture that people only attempt to approximate Bayesian inference. Many different notions of what constitutes a computational resource have been suggested, such as memory (Dasgupta & Gershman, 2021), thinking time (Ratcliff & McKoon, 2008), or physical effort (Hoppe & Rothkopf, 2016).

Cover (1999) put forward a dichotomy that will be useful for our following discussion. He refers to the algorithmic complexity of an algorithm as the number of bits needed to *implement* it. In contrast, he refers to the computational complexity of an algorithm as the space, time, or effort required to *execute* it. It is possible to cast many approximate inference schemes as resource-rational algorithms (A. N. Sanborn, 2017). The resulting models typically consider some form of computational complexity. In MCMC methods, computational complexity can be measured in terms of the number of drawn samples: drawing fewer samples leads to faster inference at the cost of introducing a bias (Courville & Daw, 2008; A. N. Sanborn et al., 2010). In variational inference, on the other

hand, it is possible to introduce an additional parameter that allows to trade-off performance against the computational complexity of transforming the prior into the posterior distribution (Binz & Schulz, 2022b; Ortega, Braun, Dyer, Kim, & Tishby, 2015). Likewise, other frameworks for building resource-rational models, such as rational meta-reasoning (Lieder & Griffiths, 2017), also only target computational complexity.

The prevalence of resource-rational models based on computational complexity is likely due to the fact that building similar models based on algorithmic complexity is much harder. Measuring algorithmic complexity historically relies on the notion of Kolmogorov complexity, which is the size of the shortest computer program that produces a particular data sequence (Chaitin, 1969; Kolmogorov, 1965; Solomonoff, 1964). Kolmogorov complexity is in general uncomputable, and, therefore, of limited practical interest.⁶

Meta-learning provides us with a straightforward way to manipulate both algorithmic and computational complexity in a common framework by adapting the size of the underlying neural network model. Limiting the complexity of network weights places a constraint on algorithmic complexity (as reducing the number of weights decreases the amount of bits needed to store them, and hence also the amount of bits needed to store the learning algorithm). Limiting the complexity of activations, on the other hand, places a constraint on computational complexity (reducing the number of hidden units, for example, decreases the memory needed for executing the meta-learned model).

Previously, both forms of complexity constraints have been realized in meta-learned models. Dasgupta et al. (2020) decreased the number of hidden units of a meta-learned inference algorithm, effectively reducing its computational complexity. In contrast, Binz et al. (2022) placed a constraint on the description length of neural network weights, which implements a form of algorithmic complexity. To the best of our knowledge, no other class

⁶ Having said that, it is possible to approximate it under certain circumstances and different authors have applied such approximations to study both human and animal cognition (Chater & Vitányi, 2003; Gauvrit, Zenil, Delahaye, & Soler-Toscano, 2014; Gauvrit, Zenil, & Tegnér, 2017; Griffiths, Daniels, Austerweil, & Tenenbaum, 2018; Zenil, Marshall, & Tegnér, 2015).

of resource-rational models exists that allows us to take both algorithmic and computational complexity into account, making this ability a unique feature of the meta-learning framework.

2.4. Neuroscience

Argument 4

Meta-learning allows us to integrate neuroscientific insights into the rational analysis of cognition by incorporating these insights into model architectures.

In addition to providing a framework for understanding many aspects of behavior, meta-learning offers a powerful lens through which to view brain structure and function. For instance, Wang et al. (2018) presented observations supporting the hypothesis that prefrontal circuits may constitute a meta-reinforcement learning system. From a computational perspective, meta-learning strives to learn a faster inner-loop learning algorithm via an adjustment of neural network weights in a slower outer-loop learning process. Within the brain, an analogous process plausibly occurs when slow, dopamine-driven synaptic change gives rise to reinforcement learning processes that occur within the activity dynamics of the prefrontal network, allowing for adaptation on much faster timescales. This perspective recontextualized the role of dopamine function in reward-based learning and was able to account for a range of previously puzzling neuroscientific findings. To highlight one example, Bromberg-Martin, Matsumoto, Hong, and Hikosaka (2010) found that dopamine signaling reflected updates in not only *experienced* but also *inferred* values of targets. Notably, a meta-reinforcement learning agent trained on the same task also recovered this pattern. Having a mapping of meta-reinforcement learning components onto existing brain regions furthermore allows us to apply experimental manipulations that directly perturb neural activity, for example by using optogenetic techniques. Wang et al. (2018) used this idea to modify their original meta-reinforcement learning architecture to mimic the blocking or enhancement of

dopaminergic reward prediction error signals, in direct analogy with optogenetic stimulation delivered to rats performing a two-armed bandit task (Stopper, Maric, Montes, Wiedman, & Floresco, 2014).

Importantly, the direction of exchange can also work in the other direction, with neuroscientific findings constraining and inspiring new forms of meta-learning architectures. Bellec, Salaj, Subramoney, Legenstein, and Maass (2018), for example, showed that recurrent networks of spiking neurons are able to display convincing learning-to-learn behavior, including in the realm of reinforcement learning. Episodic meta-reinforcement learning (Ritter et al., 2018) architectures are also heavily inspired by neuroscientific accounts of complementary learning systems in the brain (McClelland, McNaughton, & O'Reilly, 1995). Both of these examples demonstrate that meta-learning can be used to build more biologically plausible learning algorithms, and thereby highlight that it can act as a bridge between Marr's computational and implementational level (Marr, 2010).

Finally, the meta-learning perspective not only allows us to connect machine learning and neuroscience via architectural design choices but also via the kinds of tasks that are of interest. Dobs, Martinez, Kell, and Kanwisher (2022), for instance, suggested that functional specialization in neural circuits, which has been widely observed in biological brains, arises spontaneously as a consequence of task demands. In particular, they found that convolutional neural networks trained on both face and object recognition depicted emergent segregation on the basis of these tasks. Likewise, G. R. Yang, Joglekar, Song, Newsome, and Wang (2019) found that training a single recurrent neural network to perform a wide range of cognitive tasks yielded units that were clustered along different functional cognitive processes. Put another way, it seems plausible that functional specialization emerges by training neural networks on multiple tasks. While this has not been tested so far, we speculate that this also holds in the meta-learning setting, as it involves training on multiple tasks by design. If this were true, we could look at the emerging areas inside a meta-learned model, and use the resulting insights to generate

novel predictions about the processes happening in individual brain areas (Kanwisher, Khosla, & Dobs, 2023).

3. Previous Research

Meta-learned models are already starting to transform the cognitive sciences today. They allow us to model things that are hard to capture with traditional models such as compositional generalization, language understanding, and model-based reasoning. In this section, we provide an overview of what has been achieved with the help of meta-learning in previous work. We arranged this review into various thematic subcategories. For each of them, we summarize which key findings have been obtained by meta-learning and discuss why these results would have been difficult to obtain using traditional models of learning by appealing to the insights from the previous section.

3.1. Heuristics and Cognitive Biases

Meta-learning has been previously used to discover algorithms with a limited computational budget that show human-like cognitive biases as we have already alluded to earlier. Dasgupta et al. (2020) trained a neural network on a distribution of probabilistic inference problems while controlling for the number of its hidden units. They found that their model – when restricted to just a single hidden unit – captured many biases in human reasoning, including a conservatism bias and base rate neglect. Likewise, Binz et al. (2022) trained a neural network on a distribution of decision-making problems while controlling for the number of bits needed to represent the network. Their model discovered two previously suggested heuristics in specific environments and made precise prognoses about when these heuristics should be applied. In particular, knowing the correct ranking of features led to one reason decision-making, knowing the directions of features led to an equal weighting heuristic, and not knowing about either of them led to strategies that use weighted combinations of features (also see Figure 4a-b).

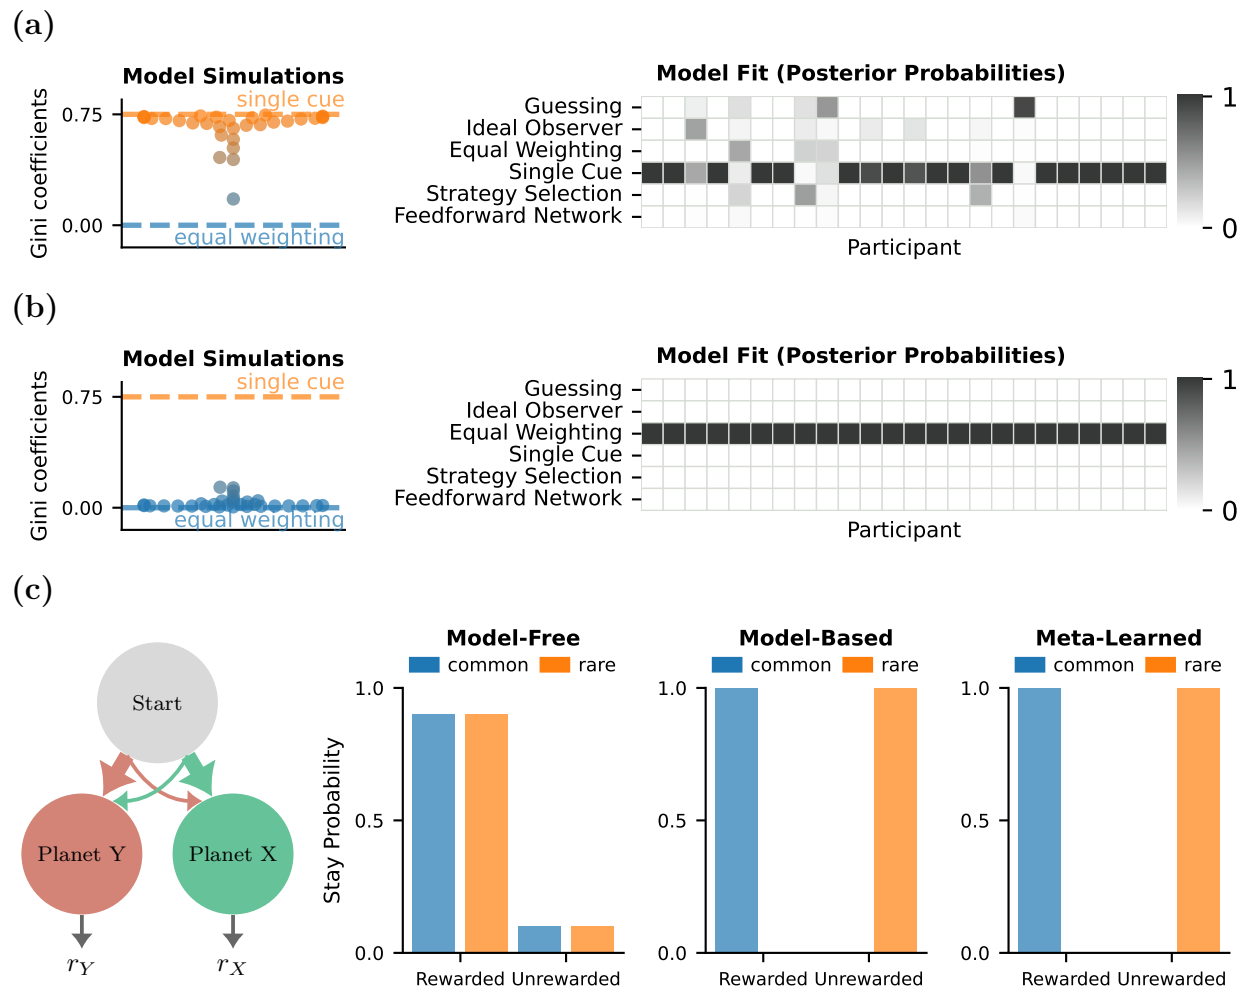


Figure 4
Example results obtained using meta-learned models. (a) In a paired comparison task, a meta-learned model identified a single-cue heuristic as the resource-rational solution when information about the feature ranking was available. Follow-up experiments revealed that people indeed apply this heuristic under the given circumstances. (b) If information about feature directions was available, the same meta-learned model identified an equal weighting heuristic as the resource-rational solution. People also applied this heuristic in the given context (Binz et al., 2022). (c) Wang et al. (2016) showed that meta-learned models can exhibit model-based learning characteristics in the two-step task (Daw, Gershman, Seymour, Dayan, & Dolan, 2011) even when they were purely trained through model-free approaches. The plots on the right illustrate the probability of repeating the previous action for different agents (model-free, model-based, meta-learned) after a common or uncommon transition and after a received or omitted reward.

In both of these studies, meta-learned models offered a novel perspective on results that were previously viewed as contradictory. This was in part possible because

meta-learning enabled us to easily manipulate the complexity of the underlying learning algorithm. While doing so is, at least in theory, also possible within the Bayesian framework, no Bayesian model that captures the full set of findings from Dasgupta et al. (2020) and Binz et al. (2022) has been discovered so far. We hypothesize that this could be because traditional rational process models struggle to capture that human strategy selection is context-dependent even before receiving any direct feedback signal (Mercier & Sperber, 2017). The meta-learned models of Dasgupta et al. (2020) and Binz et al. (2022), on the other hand, were able to readily show context-specific biases when trained on an appropriate task distribution.

3.2. Language Understanding

Meta-learning may also help us to answer questions regarding how people process, understand, and produce language. Whether the inductive biases needed to acquire a language are learned from experience or are inherited is one of these questions (Y. Yang & Piantadosi, 2022). McCoy, Grant, Smolensky, Griffiths, and Linzen (2020) investigated how to equip a model with a set of linguistic inductive biases that are relevant to human cognition. Their solution to this problem builds upon the idea of model-agnostic meta-learning (Finn et al., 2017). In particular, they meta-learned the initial weights of a neural network such that the network can adapt itself quickly to new languages using standard gradient-based learning. When being trained on a distribution over languages, these initial weights can be interpreted as universal factors that are shared across all languages. They showed that this approach identifies inductive biases (e.g. a bias for treating certain phonemes as vowels) that are useful for acquiring a language’s syllable structure. While their work focused on various modeling aspects, they suggested that their framework “could [...] be used to empirically investigate the effects that those inductive biases have.” They additionally argued that a Bayesian modeling approach would only be able to consider a restrictive set of inductive biases as it needs to commit to a particular

representation and inference algorithm. In contrast, the meta-learning framework made it easy to implement the intended inductive biases by simply manipulating the distribution of encountered languages.

The ability to compose simple elements into complex entities is at the heart of human language. The property of languages to “make infinite use of finite means” (Chomsky, 2014) is what allows us to make strong generalizations from limited data. For example, people readily understand what it means to “dax twice” or to “dax slowly” after learning about the meaning of the verb “dax.” How to build models with a similar proficiency, however, remains an open research question. Lake (2019) showed that a transformer-like neural network can be trained to make such compositional generalizations through meta-learning. Importantly, during meta-learning, his models were adapted to problems that required compositional generalization, and could thereby acquire the skills needed to solve entirely new problems. Lake (2019) argued that meta-learning “has implications for understanding how people generalize compositionally.” In particular, it highlights the importance of “tackling a series of changing learning problems rather than iterating through a static data-set”, as it is done in traditional neural network training schemes.

3.3. Inductive Biases

Human cognition comes with many useful inductive biases beyond the ability to reason compositionally. The preference for simplicity is one of these biases (Chater & Vitányi, 2003; Feldman, 2016). We readily extract abstract low-dimensional rules that allow us to generalize entirely new situations. Meta-learning is an ideal tool to build models with similar preferences because we can easily generate tasks based on simple rules and use them for meta-learning, thereby enabling an agent to acquire the desired inductive bias from data.

Towards this end, Kumar, Dasgupta, Cohen, Daw, and Griffiths (2020b) tested

humans and meta-reinforcement agents on a family of structured tasks generated by a grammar and compared their performance to those trained on a non-structured version of the same task with matched statistics. They expanded these results to a larger suite of tasks generated from simple abstract rules in Kumar, Dasgupta, et al. (2022). Humans found it easier to learn in structured tasks, confirming that they have strong priors towards simple abstract rules (Schulz, Tenenbaum, Duvenaud, Speekenbrink, & Gershman, 2017). However, their analysis also indicated that meta-learning is easier on non-structured tasks than on structured tasks. In follow-up work, they found that this result also holds for agents that were trained purely on the structured version of their task but evaluated on both versions (Kumar, Correa, et al., 2022) – a quite astonishing finding considering that one would expect an agent to perform better on the task distribution it was trained on. The authors addressed this mismatch between humans and meta-learned agents by guiding agents during training to reproduce natural language descriptions that people provided to describe a given task. They found that grounding meta-learned agents in natural language descriptions not only improved their performance but also led to more human-like inductive biases, demonstrating that natural language can serve as a source for abstractions within human cognition.

Their line of work utilizes another interesting technique for training meta-learning agents (Kumar, Correa, et al., 2022; Kumar, Dasgupta, et al., 2022). It does not rely on a hand-designed task distribution but instead involves sampling tasks from the prior distribution of human participants using a technique known as Gibbs sampling with people (Harrison et al., 2020; A. Sanborn & Griffiths, 2007). While doing so provides them with a data-set of tasks, no expression of the corresponding prior distribution over them is accessible and, hence, it is non-trivial to define a Bayesian model for the given setting. A meta-learned agent, on the other hand, was readily obtained by training on the collected samples.

3.4. Model-Based Reasoning

Many realistic scenarios afford two distinct types of learning: model-free and model-based. Model-free learning algorithms directly adjust their strategies using observed outcomes. Model-based learning algorithms, on the other hand, learn about the transition and reward probabilities of an environment, which are then used for downstream reasoning tasks. People are generally thought to be able to perform model-based learning, at least to some extent, and assuming that the problem at hand calls for it (Daw et al., 2011; Kool, Cushman, & Gershman, 2016). Wang et al. (2016) showed that a meta-learned algorithm can display model-based behavior, even if it was trained through a pure model-free reinforcement learning algorithm (see Figure 4c).

Having a model of the world also acts as the basis for causal reasoning. Traditionally, making causal inferences relies on the notion of Pearl’s do-calculus (Pearl, 2009). Dasgupta et al. (2019), however, showed that meta-learning can be used to create models that draw causal inferences from observational data, select informative interventions, and make counterfactual predictions. While they have not related their model to human data directly, it could in future work serve as the basis to study how people make causal judgments in complex domains and explain why and when they deviate from normative causal theories (Bramley, Dayan, Griffiths, & Lagnado, 2017; Gerstenberg, Goodman, Lagnado, & Tenenbaum, 2021).

Together, these two examples highlight that model-based reasoning capabilities can emerge internally in a meta-learned model if they are beneficial for solving the encountered problem. While there are already many traditional models that can perform such tasks, these models are often slow at run-time as they typically involve Bayesian inference, planning, or both. Meta-learning, on the other hand, “shifts most of the compute burden from inference time to training time [which] is advantageous when training time is ample but fast answers are needed at run-time” (Dasgupta et al., 2019), and may therefore explain how people can perform such intricate computations within a reasonable time-frame.

While model-based reasoning is an emerging property of meta-learned models, it may also be integrated explicitly into such models should it be desired. Jensen, Hennequin, and Mattar (2023) have taken this route, and augmented a standard meta-reinforcement learning agent with the ability to perform temporally extended planning using imagined rollouts. In each time-step, their agent can decide to perform a planning operation instead of directly interacting with the environment (in this case, a spatial navigation task). Their meta-learned agents opted to perform this planning operation consistently after training. Importantly, the model showed striking similarities to patterns of human deliberation by performing more planning early on and with an increased distance to the goal. Furthermore, they found that patterns of hippocampal replays resembled the rollouts of their model.

3.5. Exploration

People do not only have to integrate observed information into their existing knowledge, but they also have to actively determine what information to sample. They constantly face situations that require them to decide whether they should explore something new or whether they should rather exploit what they already know. Previous research suggests that people solve this exploration-exploitation dilemma using a combination of directed and random exploration strategies (Gershman, 2018; Schulz & Gershman, 2019; Wilson, Geana, White, Ludvig, & Cohen, 2014; Wu, Schulz, Speekenbrink, Nelson, & Meder, 2018). Why do people use these particular strategies and not others? Binz and Schulz (2022a) hypothesized that they do so because human exploration follows resource-rational principles. To test this claim, they devised a family of resource-rational reinforcement learning algorithms by combining ideas from meta-learning and information theory. Their meta-learned model discovered a diverse set of exploration strategies, including random and directed exploration, that captured human exploration better than alternative approaches. In this domain, meta-learning offered a direct path

towards investigating the hypothesis that people try to explore optimally but are subject to limited computational resources, whereas designing hand-crafted models for studying the same question would have been more intricate.

It is not only important to decide how to explore, but also to decide whether exploration is worthwhile in the first place. Lange and Sprekeler (2020) studied this question using the meta-learning framework. Their meta-learned agents are able to flexibly interpolate between implementing exploratory learning behaviors and hard-coded, non-learning strategies. Importantly, which behavior was realized crucially depended on environmental properties, such as the diversity of the task distribution, the task complexity, and the agent’s lifetime. They showed, for instance, that agents with a short lifetime should opt for small rewards that are easy to find, while agents with an extended lifetime should spend their time exploring the environment. The study of Lange and Sprekeler (2020) clearly demonstrates that meta-learning makes it conceptually easy to iterate over different environmental assumptions inside a rational analysis of cognition. They only had to modify the environment as desired, followed by rerunning their meta-learning procedure. In contrast, traditional modeling approaches would require hand-designing a new optimal agent each time an environmental change occurs.

3.6. Cognitive Control

Humans are remarkable at adapting to task-specific demands. The processes behind this ability are collectively referred to as cognitive control (Botvinick, Braver, Barch, Carter, & Cohen, 2001). Cohen (2017) even argues that “the capacity for cognitive control is perhaps the most distinguishing characteristic of human behavior.” It should therefore come as no surprise that cognitive control has received a significant amount of attention from a computational perspective (Botvinick & Cohen, 2014; Collins & Frank, 2013). Recently, some of these computational investigations have been extended to the meta-learning framework.

The ability to adjust computational resources as needed is one hallmark of cognitive control. Moskowitz, Miller, Sahani, and Botvinick (2022) proposed a meta-learned model with such characteristics. Their model learns a simple default policy – similar to the model of Binz and Schulz (2022a) – that can be overwritten by a more complex one if necessary. They demonstrate that this model is not only able to capture behavioral phenomena from the cognitive control literature but also known effects in decision-making and reinforcement learning tasks, thereby linking the three domains. Importantly, their study highlights that the meta-learning framework offers the means to account for multiple computational costs instead of just a single one – in this case, a cost for implementing the default policy and one for deviating from it.

Taking contextual cues into consideration is another vital aspect of cognitive control. Dubey, Grant, Luo, Narasimhan, and Griffiths (2020) implemented this idea in the meta-learning framework. In their model, contextual cues determine the initialization of a task-specific neural network that is then trained using model-agnostic meta-learning. They showed that such a model captures “the context-sensitivity of human behavior in a simple but well-studied cognitive control task.” Furthermore, they demonstrated that it scales well to more complex domains (including tasks from the MuJoCo (Todorov, Erez, & Tassa, 2012), CelebA Liu, Luo, Wang, and Tang (2015) and MetaWorld (Yu et al., 2020) benchmarks), thereby opening up new opportunities for modeling human behavior in naturalistic scenarios.

4. Why Is Not Everything Meta-Learned?

We have laid out different arguments that make meta-learning a useful tool for constructing cognitive models, but it is important to note that we do not claim that meta-learning is the ultimate solution to every modeling problem. Instead, it is essential to understand when meta-learning is the right tool for the job and when not.

4.1. Lack of Interpretability

Perhaps its most significant detriment is that meta-learning never provides us with analytical solutions that we can inspect, analyze and reason about. In contrast to this, some Bayesian models have analytical solutions. Take as an example the data-generating distribution that we encountered earlier (Equations 1-2). For these assumptions, a closed-form expression of the posterior predictive distribution is available. By looking at this closed-form expression, researchers have generated new predictions and subsequently tested them empirically (Daw et al., 2008; Dayan & Kakade, 2000; Gershman, 2015). Performing the same kind of analysis with a meta-learned model is not as straightforward. We do not have access to an underlying mathematical expression, which makes a structured exploration of theories much harder.

That being said, there are still ways to analyze a meta-learned model's behavior. For one, it is possible to use model architectures that facilitate interpretability. Binz et al. (2022) relied on this approach and designed a neural network architecture that produced weights of a probit regression model which were then used to cluster applied strategies into different categories. Doing so enabled them to identify which strategy was used by their meta-learned model in a particular situation.

Recently, researchers have also started to use tools from cognitive psychology to analyze the behavior of black-box models (Rich & Gureckis, 2019; Ritter, Barrett, Santoro, & Botvinick, 2017; Schulz & Dayan, 2020). For example, it is possible to treat such models just like participants in a psychological experiment and use the collected data to analyze their behavior similar to how psychologists would analyze human behavior (Binz & Schulz, 2023; Dasgupta et al., 2022; Rahwan et al., 2019; Schramowski, Turan, Andersen, Rothkopf, & Kersting, 2022). We believe that this approach has great potential for analyzing increasingly capable and opaque artificial agents, including those obtained via meta-learning.

4.2. Intricate Training Processes

When using the meta-learning framework, one also has to deal with the fact that training neural networks is complex and takes time. Neural network models contain many moving parts, like weight initializations or the used optimizer, that have to be chosen appropriately such that training can take off in the first place, and training itself may take hours or days until it is finished. When we want to modify assumptions in the data-generating distribution, we have to retrain the whole system from scratch altogether. Thus, although the process of iterating over different environmental assumptions is conceptually straightforward in the meta-learning framework, it may be time-consuming. Bayesian models can, in comparison, sometimes be more quickly adapted to changes in environmental assumptions. To illustrate this, let us assume that you wanted to explain human behavior through a meta-learned model that was trained on the data-generating distribution from Equations 1-2, but found that the resulting model does not fit the observed data well. Next, you want to consider the alternative hypothesis that people assume a non-stationary environment. While this modification could be done quickly in the corresponding Bayesian model, the meta-learning framework requires retraining on newly generated data.

There is, furthermore, no guarantee that a fully converged meta-learned model actually implements a Bayes-optimal learning algorithm. While we were able to compare to analytical solutions for simple cases like our insect length example, it is in general impossible to verify that a meta-learned algorithm is optimal. Indeed, there are reported cases in which meta-learning failed to find the Bayes-optimal solution (Wang et al., 2021). We believe that this issue can be somewhat mitigated by validating meta-learned models in various different ways. But, ultimately future work should come up with techniques to verify meta-learned models.

4.3. Meta-Learned or Bayesian Inference?

In summary, both frameworks — meta-learning and Bayesian inference — have their unique strengths and weaknesses. The meta-learning framework does and will not replace Bayesian inference but complement it. It broadens our available toolkit and enables researchers to study questions that were previously out of reach. However, there are certainly situations in which traditional Bayesian inference is a more appropriate modeling choice as we have outlined in this section.

5. The Role of Neural Networks

Most of the points we have discussed so far are agnostic regarding the function approximator implementing the meta-learned algorithm. However, at the same time, we have appealed to neural networks at various points throughout the text. When one looks at prior work, it can also be observed that neural networks are the predominant model class in the meta-learning setting. Why is that the case? In addition to their universality, neural networks offer one big opportunity: they provide a flexible framework for engineering different types of inductive biases into a computational model (Goyal & Bengio, 2022). In the following section, we will highlight three examples of how previous work has accomplished this. For each of these examples, we take a concept from psychology, and show how it can be readily accommodated in a meta-learned model.

Perhaps one of the most persuasive idea in cognitive modeling is that of gradient-based learning. It is not only at the heart of one of the most influential models — the Rescorla-Wagner model (Gershman, 2015; R. A. Rescorla, 1972) — but also features prominently in many other theories of human learning, such as connectionist models (Rumelhart, McClelland, Group, et al., 1988). Even though the earlier outlined meta-learning procedure relies on gradient-based learning in the outer loop, the resulting inner-loop dynamics must bear no resemblance to gradient descent. However, it is possible to construct meta-learned models whose inner-loop updates rely on gradient-based

learning. Finn et al. (2017) proposed a meta-learning technique known as model-agnostic meta-learning that finds optimal initial parameters of a feedforward neural network that is subsequently trained via gradient descent. The idea is that these optimal initial parameters allow the feedforward network to generalize to multiple tasks in a minimal number of gradient steps. While their general setup is similar to the one we discussed, it leads to models that learn via gradient descent instead of models that implement a learning algorithm inside the dynamics of a recurrent neural network. Kirsch and Schmidhuber (2021) recently brought these two approaches together into a single model. Their proposed architecture consists of multiple recurrent neural networks that interact with each other. Importantly, they showed that one particular configuration of these networks could implement backpropagation in the forward pass, thereby being able to perform gradient-based learning in a memory-based system.

Exemplar-based models – like the generalized category model (Nosofsky, 2011) – are one of the most prominent approaches for modeling how people categorize items into different classes (Kruschke, 1990; Shepard, 1987). They categorize a new instance based on the estimated similarity between that instance and previously seen examples. Recently, meta-learned models with exemplar-based reasoning abilities have been proposed for the task of few-shot classification, in which a classifier must generalize based on a training set containing only a few examples. Matching networks (Vinyals, Blundell, Lillicrap, Wierstra, et al., 2016) accomplish this by classifying a new data-point using a similarity-weighted combination of categories in the training set. Importantly, similarity is computed over a learned embedding space, thereby ensuring that the model can scale to high-dimensional stimuli. Follow-up work has taken inspiration from another hugely influential model of human category learning and replaced the exemplar-based mechanism used in matching networks with one based on category prototypes (Snell, Swersky, & Zemel, 2017).

Finally, making inferences using similarities to previous experiences is not only useful for supervised learning but also in the reinforcement learning setting. In the

reinforcement learning literature, the ability to store and recollect states or trajectories for later use is studied under the name of episodic memory (Lengyel & Dayan, 2007). It has been argued that episodic memory could be the key to explaining human performance in naturalistic environments (Gershman & Daw, 2017). Episodic memory also plays a crucial role in neuroscience, with studies showing that highly rewarding instances are stored in the hippocampus and made available for recall as and when required Blundell et al. (2016). Ritter et al. (2018) build upon the neural episodic control idea from Pritzel et al. (2017) and utilize a differential neural dictionary for episodic recall in the context of meta-learning. Their dictionary stores encodings from previously experienced tasks, which can then be later queried as needed. With this addition, their meta-learned model is able to recall previously discovered policies, retrieve memories using category examples, handle compositional tasks, re-instate memories while traversing the environment, and recover a learning strategy people use in a neuroscience-inspired task.

In summary, human cognition comes with a variety of inductive biases and neural networks provide flexible ways to easily incorporate them into meta-learned models of cognition. We have outlined three such examples in the section, demonstrating how to integrate gradient-based learning, exemplar- and prototype-based reasoning, and episodic memory into a meta-learned model. There are, furthermore, many other inductive biases for neural network architectures that could be utilized in the context of meta-learning but have not been yet. Examples include networks that perform differentiable planning (Farquhar, Rocktäschel, Igl, & Whiteson, 2017; Tamar, Wu, Thomas, Levine, & Abbeel, 2016), extract object-based representations (Piloto, Weinstein, Battaglia, & Botvinick, 2022; Sancaktar, Blaes, & Martius, 2022), or modify their own connections through synaptic plasticity (Miconi, Rawal, Clune, & Stanley, 2020; Schlag, Irie, & Schmidhuber, 2021).

6. Towards a Domain-General Model of Human Learning

What does the future hold for meta-learning? The current generation of meta-learned models of cognition is almost exclusively trained on the data-generating distribution of a specific problem family. While this training process enables them to generalize to new tasks inside this problem family, they are unlikely to generalize to completely different domains. We would, for example, not expect a meta-learned algorithm to perform a challenging maze navigation task if it was only trained to predict the lengths of insect species.

While domain-specific models have (and will continue to) provide answers to important research questions, we agree with Newell (1992) that “unified theories of cognition are the only way to bring this wonderful, increasing fund of knowledge under intellectual control.” Ideally, such a unified theory should manifest itself in a domain-general cognitive model that cannot only solve prediction tasks but is also capable of human-like decision-making (Gigerenzer & Gaissmaier, 2011), category learning (Ashby, Maddox, et al., 2005), navigation (Montello, 2005), problem-solving (Newell, Simon, et al., 1972) and so on. We consider the meta-learning framework the ideal tool for accomplishing this goal as it allows us to compile arbitrary assumptions about an agent’s beliefs of the world (arguments 1 and 2) and its computational architecture (arguments 3 and 4) into a cognitive model.

To obtain such a domain-general cognitive model via meta-learning, however, a few challenges need to be tackled. First of all, there is the looming question of how a data-generating distribution that contains many different problems should be constructed. Here, we may take inspiration from the machine learning community, where researchers have devised generalist agents by training neural networks on a large set of problems (Reed et al., 2022). (A. A. Team et al., 2023) have recently shown that this is a promising path for scaling up meta-learning models. They trained a meta-reinforcement learning agent on a vast open-ended world with over 10^{40} possible tasks. The resulting agent can adapt to

held-out problems as quickly as humans, and “displays on-the-fly hypothesis-driven exploration, efficient exploitation of acquired knowledge, and can successfully be prompted with first-person demonstrations.” In the same vein, we may come up with a large collection of tasks that are more commonly used to study human behavior (Miconi, 2023; Molano-Mazon et al., 2022; G. R. Yang et al., 2019), and use them to train a meta-learned model of cognition.

Language will likely play an important role in future meta-learning systems. We do not want a system that learns every task from scratch via trial and error but one that can be provided with a set of instructions similar to how a human subject would be instructed in a psychological experiment. Having agents capable of language will not only enable them to understand new tasks in a zero-shot manner but may also facilitate their cognitive abilities. It, for example, allows them to decompose tasks into sub-tasks, learn from other agents, or generate explanations (Colas, Karch, Moulin-Frier, & Oudeyer, 2022). Fortunately, our current best language models (Brown et al., 2020; Chowdhery et al., 2022) and meta-learning systems are both based on neural networks. Thus, integrating language capabilities into a meta-learned model of cognition should – at least conceptually – be fairly straightforward. Doing so would furthermore enable such models to harvest the compositional nature of language to make strong generalizations to tasks outside of the meta-learning distribution. The potential for this was highlighted in a recent study of (Riveland & Pouget, 2022) which found that language-conditioned recurrent neural network models can perform entirely novel psychophysical tasks with high accuracy.

Moreover, a sufficiently general model of human cognition must not only be able to select amongst several given options, like in a decision-making or category learning setting, but it also needs to maneuver within a complex world. For this, it needs to perceive and process high-dimensional visual stimuli, it needs to control a body with many degrees of freedom, and it needs to actively engage with other agents. Many of these problems have been at the heart of the deep learning community (Hill et al., 2020; McClelland, Hill,

Rudolph, Baldridge, & Schütze, 2020; Strouse, McKee, Botvinick, Hughes, & Everett, 2021; O. E. L. Team et al., 2021), and it will be interesting to see whether the solutions developed there can be integrated into a meta-learned model of cognition.

Finally, there are also some challenges on the algorithmic side that need to be taken into account. In particular, it is unclear how far currently used model architectures and outer-loop learning algorithms scale. While contemporary meta-learning algorithms are able to find approximately Bayes-optimal solutions to simple problems, they sometimes struggle to do so on more complex ones (e.g. as in the earlier discussed work of Wang et al. (2021)). Therefore, it seems likely that simply increasing the complexity of the meta-learning distribution will not be sufficient – we will also need model architectures and outer-loop learning algorithms that can handle increasingly complex data-generating distributions. The transformer architecture (Vaswani et al., 2017), which has been very successful at training large language models (Brown et al., 2020; Chowdhery et al., 2022), provides one promising candidate, but there could be countless other (so far undiscovered) alternatives.

Thus, taken together, there are still substantial challenges involved in creating a domain-general meta-learned model of cognition. In particular, we have argued in this section that we need to (1) meta-learn on more diverse task distributions, (2) develop agents that can parse instructions in form of natural language, (3) embody these agents in realistic problem settings, and (4) find model architectures that scale to these complex problems. Figure 5 summarizes these points graphically.

7. Conclusion

Most computational models of human learning are hand-designed, meaning that at some point a researcher sat down and defined how they behave. Meta-learning starts with an entirely different premise. Instead of designing learning algorithms by hand, one trains a system to achieve its goals by repeatedly letting it interact with an environment. While

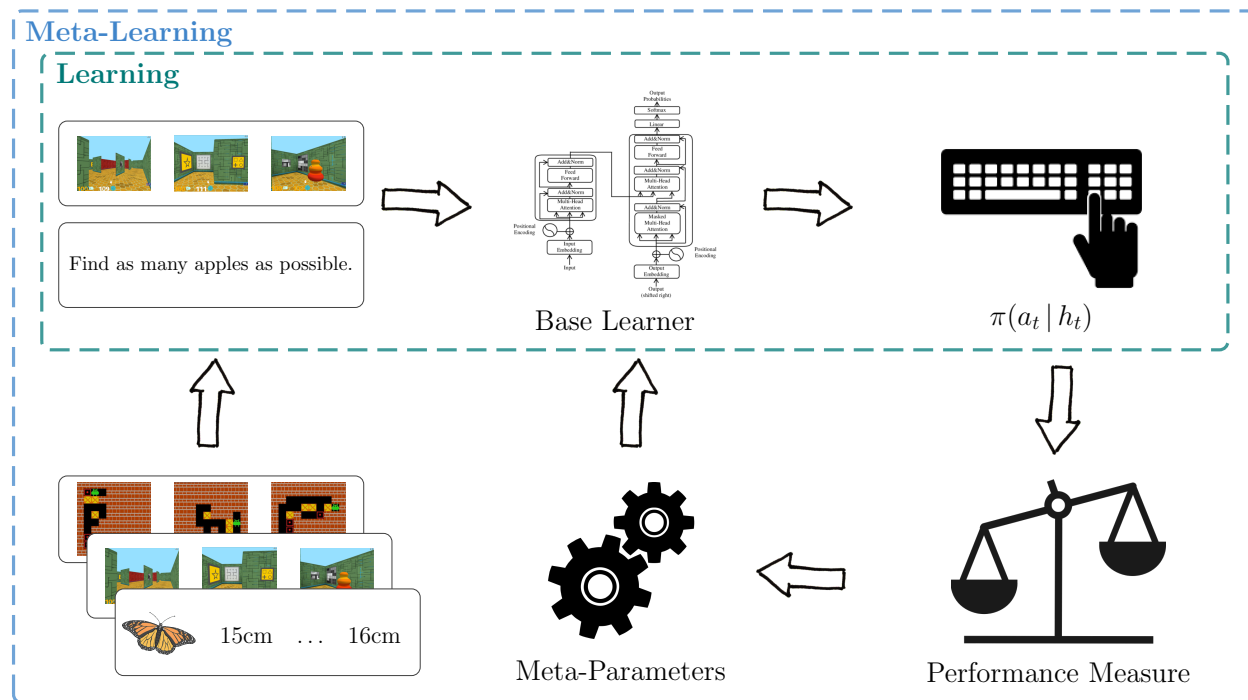


Figure 5
Illustration of how a domain-general meta-learned model of cognition could look like. Modifications include training on more diverse task distributions, providing natural language instructions as additional inputs, and relying on scalable model architectures.

this seems quite different from traditional models of learning on the surface, we have highlighted that the meta-learning framework actually has a deep connection to the idea of Bayesian inference, and thereby to the rational analysis of cognition. Using this connection as a starting point, we have highlighted several advantages of the meta-learning framework for constructing rational models of cognition. Together, our arguments demonstrate that meta-learning cannot only be applied in situations where Bayesian inference is impossible but also facilitates the inclusion of computational constraints and neuroscientific insights into rational models of human cognition. Earlier criticisms of the rational analysis of cognition have repeatedly pointed out that “rational Bayesian models are significantly unconstrained” and that they should be “developed in conjunction with mechanistic considerations to offer substantive explanations of cognition” (Jones & Love, 2011). We believe that the meta-learning framework provides the ideal opportunity to do so as it

allows for a painless integration of flexible computational mechanisms.

It is worth pointing out that meta-learning can be also motivated by taking neural networks as a starting point. From this perspective, it bridges two of the most popular theories of cognition – Bayesian models and connectionism – by bringing the scalability of neural network models into the rational analysis of cognition. We, therefore, believe that meta-learning provides a powerful tool to scale up psychological theories to more complex settings. However, at the same time, meta-learning has not delivered on this promise yet. Existing meta-learned models of cognition are typically applied to classical scenarios where established models already exist. Thus, we have to ask: what prevents the application to more complex and general paradigms? First, such paradigms themselves have to be developed. Fortunately, there is currently a trend toward measuring human behavior on more naturalistic tasks (Brändle, Stocks, Tenenbaum, Gershman, & Schulz, 2022; Brändle, Binz, & Schulz, 2022; Schulz et al., 2019), and it will be interesting to see what role meta-learning will play in modeling behavior in such settings. Furthermore, meta-learning can be intricate and time-consuming. We hope that the present article – together with the accompanying code examples – makes this technique less opaque and more accessible to a wider audience. Future advances in hardware will likely make the meta-learning process quicker and we are therefore hopeful that meta-learning can ultimately fulfill its promise of identifying plausible models of human cognition in situations that are out of reach for hand-designed algorithms.

Box 1. We proof that the posterior predictive distribution $p(x_{t+1}|x_{1:t})$ maximizes the log-likelihood of future observations averaged over the data-generating distribution:

$$p(x_{t+1}|x_{1:t}) = \arg \max_q \mathbb{E}_{p(\mu, x_{1:t+1})} [\log q(x_{t+1}|x_{1:t})] \quad (8)$$

The essence of this proof is to show that the posterior predictive distribution is superior to any other reference distribution $r(x_{t+1}|x_{1:t})$ in terms of log-likelihood:

$$\mathbb{E}_{p(\mu, x_{1:t})} [\log p(x_{t+1}|x_{1:t})] \geq \mathbb{E}_{p(\mu, x_{1:t})} [\log r(x_{t+1}|x_{1:t})]$$

or equivalently that:

$$\mathbb{E}_{p(\mu, x_{1:t})} \left[\log \frac{p(x_{t+1}|x_{1:t})}{r(x_{t+1}|x_{1:t})} \right] \geq 0$$

Proofing this conjecture is straight-forward (Aitchison, 1975):

$$\begin{aligned} & \mathbb{E}_{p(\mu, x_{1:t})} \left[\log \frac{p(x_{t+1}|x_{1:t})}{r(x_{t+1}|x_{1:t})} \right] \\ = & \sum_{\mu} \sum_{x_{1:t}} \sum_{x_{t+1}} \log \frac{p(x_{t+1}|x_{1:t})}{r(x_{t+1}|x_{1:t})} p(x_{t+1}|\mu) p(x_{1:t}|\mu) p(\mu) && \text{definition of expectation} \\ = & \sum_{x_{1:t}} \sum_{\mu} \sum_{x_{t+1}} \log \frac{p(x_{t+1}|x_{1:t})}{r(x_{t+1}|x_{1:t})} p(x_{t+1}|\mu) p(x_{1:t}|\mu) p(\mu) && \text{change order of summation} \\ = & \sum_{x_{1:t}} \sum_{\mu} \sum_{x_{t+1}} \log \frac{p(x_{t+1}|x_{1:t})}{r(x_{t+1}|x_{1:t})} p(x_{t+1}|\mu) p(\mu|x_{1:t}) p(x_{1:t}) && \text{Bayes' rule} \\ = & \sum_{x_{1:t}} \left[\sum_{\mu} \sum_{x_{t+1}} \log \frac{p(x_{t+1}|x_{1:t})}{r(x_{t+1}|x_{1:t})} p(x_{t+1}|\mu) p(\mu|x_{1:t}) \right] p(x_{1:t}) && \text{factor out } p(x_{1:t}) \\ = & \sum_{x_{1:t}} \left[\sum_{x_{t+1}} \sum_{\mu} \log \frac{p(x_{t+1}|x_{1:t})}{r(x_{t+1}|x_{1:t})} p(x_{t+1}|\mu) p(\mu|x_{1:t}) \right] p(x_{1:t}) && \text{change order of summation} \\ = & \sum_{x_{1:t}} \left[\sum_{x_{t+1}} \log \frac{p(x_{t+1}|x_{1:t})}{r(x_{t+1}|x_{1:t})} \left[\sum_{\mu} p(x_{t+1}|\mu) p(\mu|x_{1:t}) \right] \right] p(x_{1:t}) && \text{factor out log-term} \\ = & \sum_{x_{1:t}} \left[\sum_{x_{t+1}} \log \frac{p(x_{t+1}|x_{1:t})}{r(x_{t+1}|x_{1:t})} p(x_{t+1}|x_{1:t}) \right] p(x_{1:t}) && \text{Equation 4} \\ = & \sum_{x_{1:t}} \text{KL} [p(x_{t+1}|x_{1:t}) || r(x_{t+1}|x_{1:t})] p(x_{1:t}) && \text{definition of KL divergence} \\ \geq & 0 && \blacksquare \end{aligned}$$

Note that while we used sums in our proof, thereby assuming that relevant quantities take discrete values, the same ideas can be readily applied to continuous-valued quantities by replacing sums with integrals.

Box 2. The main text has focused on tasks in which an agent receives direct feedback about which response would have been correct. In the real world, however, people do not always receive such explicit feedback. They, instead, often have to deal with partial information – taking the form of rewards, utilities, or costs – that merely informs them about the quality of their response.

Problems that fall into this category are often modeled as Markov decision processes (MDPs). In an MDP, an agent repeatedly interacts with an environment. In each time-step, it observes the state of the environment s_t and can take an action a_t that leads to a reward signal r_t sampled from a reward distribution $p(r_t|s_t, a_t, \mu_r)$. Executing an action furthermore influences the environment state at the next time-step according to a transition distribution $p(s_{t+1}|s_t, a_t, \mu_s)$.

The goal of a Bayes-optimal reinforcement learning agent is to find a policy, which is a mapping from a history of observations $h_t = s_1, a_1, r_1, \dots, s_{t-1}, a_{t-1}, r_{t-1}, s_t$ to a probability distribution over actions $\pi^*(a_t|h_t)$, that maximizes the total amount of obtained rewards across a finite horizon H averaged over all problems that may be encountered:

$$\pi^*(a_t|h_t) = \arg \max_{\pi} \mathbb{E}_{p(\mu_r, \mu_s)} \prod p(r_t|s_t, a_t, \mu_r) p(s_{t+1}|s_t, a_t, \mu_s) \pi(a_t|h_t) \left[\sum_{t=1}^H r_t \right] \quad (9)$$

MDPs with unknown reward and transition distributions are substantially more challenging to solve optimally compared to supervised problems as there is no teacher informing the agent about which actions are right or wrong. Instead, the agent has to figure out the most rewarding course of action solely through trial and error. Finding an analytical solution to Equation 9 is extremely challenging and indeed only possible for a few special cases (Duff, 2003; Gittins, 1979), which made it historically near impossible to investigate such problems within the framework of rational analysis.

Even though finding an analytical expression of the Bayes-optimal policy is often impossible, it is straightforward to meta-learn an approximation to it (Duan et al., 2016; Wang et al., 2016). The general concept is almost identical to the supervised learning case: parameterize the to-be-learned policy with a recurrent neural network and replace the maximization over the set of all possible policies from Equation 9 with a sample-based approximation that maximizes over neural network parameters. The resulting problem can then be solved using any standard deep reinforcement learning algorithm.

Like in the supervised learning case, the resulting recurrent neural network implements a free-standing reinforcement learning algorithm after meta-learning is completed. Learning is once again implemented via a simple forward pass through the network, i.e., by conditioning the model on an additional data-point. The meta-learned reinforcement learning algorithm approximates the Bayes-optimal policy under the same conditions as in the supervised learning case: a sufficiently expressive model and an optimization procedure that is able to find the global optimum.

Acknowledgements

The authors would like to thank Sreejan Kumar, Tobias Ludwig, Dominik Endres, and Adam Santoro for their valuable feedback on an earlier draft.

Funding statement

This work was funded by the Max Planck Society, the Volkswagen Foundation, as well as the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy–EXC2064/1–390727645.

Conflicts of Interest statement

The author(s) declare none.

8. References

- Aitchison, J. (1975). Goodness of prediction fit. *Biometrika*, 62(3), 547–554.
- Anderson, J. R. (2013a). *The adaptive character of thought*. Psychology Press.
- Anderson, J. R. (2013b). *The architecture of cognition*. Psychology Press.
- Ashby, F. G., Maddox, W. T., et al. (2005). Human category learning. *Annual review of psychology*, 56(1), 149–178.
- Baxter, J. (1998). Theoretical models of learning to learn. In *Learning to learn* (pp. 71–94). Springer.
- Bellec, G., Salaj, D., Subramoney, A., Legenstein, R., & Maass, W. (2018). Long short-term memory and learning-to-learn in networks of spiking neurons. *Advances in neural information processing systems*, 31.
- Bengio, Y., Bengio, S., & Cloutier, J. (1991). Learning a synaptic learning rule. In *Ijcn-91-seattle international joint conference on neural networks* (Vol. 2, pp. 969–vol).
- Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics: Applications and Foundations 1, 2*, 69–186.
- Binmore, K. (2007). Rational decisions in large worlds. *Annales d’Economie et de Statistique*, 25–41.
- Binz, M., Gershman, S. J., Schulz, E., & Endres, D. (2022). Heuristics from bounded meta-learned inference. *Psychological review*.
- Binz, M., & Schulz, E. (2022a). Modeling human exploration through resource-rational reinforcement learning. In A. H. Oh, A. Agarwal, D. Belgrave, & K. Cho (Eds.), *Advances in neural information processing systems*. Retrieved from <https://openreview.net/forum?id=W1MUJv5zaXP>
- Binz, M., & Schulz, E. (2022b). Reconstructing the einstellung effect. *Computational Brain & Behavior*, 1–17.
- Binz, M., & Schulz, E. (2023). Using cognitive psychology to understand gpt-3.

- Proceedings of the National Academy of Sciences*, 120(6), e2218523120.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.
- Blundell, C., Uria, B., Pritzel, A., Li, Y., Ruderman, A., Leibo, J. Z., . . . Hassabis, D. (2016). Model-free episodic control. *arXiv preprint arXiv:1606.04460*.
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological review*, 108(3), 624.
- Botvinick, M. M., & Cohen, J. D. (2014). The computational and neural basis of cognitive control: charted territory and new frontiers. *Cognitive science*, 38(6), 1249–1285.
- Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2017). Formalizing neurath’s ship: Approximate algorithms for online causal learning. *Psychological review*, 124(3), 301.
- Brändle, F., Stocks, L. J., Tenenbaum, J. B., Gershman, S. J., & Schulz, E. (2022). Intrinsically motivated exploration as empowerment. *PsyArXiv*. January, 14.
- Brighton, H., & Gigerenzer, G. (2012). Are rational actor models “rational” outside small worlds. *Evolution and Rationality: Decisions, Co-operation, and Strategic Behavior*, 84–109.
- Bromberg-Martin, E. S., Matsumoto, M., Hong, S., & Hikosaka, O. (2010). A pallidus-habenula-dopamine pathway signals inferred stimulus values. *Journal of neurophysiology*, 104(2), 1068–1076.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., . . . others (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877–1901.
- Brändle, F., Binz, M., & Schulz, E. (2022). Exploration beyond bandits. In I. Cogliati Dezza, E. Schulz, & C. M. Wu (Eds.), *The drive for knowledge: The science of human information seeking* (p. 147–168). Cambridge University Press. doi: 10.1017/9781009026949.008
- Chaitin, G. J. (1969). On the simplicity and speed of programs for computing infinite sets

- of natural numbers. *Journal of the ACM (JACM)*, 16(3), 407–422.
- Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in cognitive sciences*, 3(2), 57–65.
- Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in cognitive sciences*, 7(1), 19–22.
- Chomsky, N. (2014). *Aspects of the theory of syntax* (Vol. 11). MIT press.
- Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., . . . others (2022). Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Cohen, J. D. (2017). Cognitive control: Core constructs and current considerations. *The Wiley handbook of cognitive control*, 1–28.
- Colas, C., Karch, T., Moulin-Frier, C., & Oudeyer, P.-Y. (2022). Language and culture internalization for human-like autotelic ai. *Nature Machine Intelligence*, 4(12), 1068–1076.
- Collins, A. G., & Frank, M. J. (2013). Cognitive control over learning: creating, clustering, and generalizing task-set structure. *Psychological review*, 120(1), 190.
- Corner, A., & Hahn, U. (2013). Normative theories of argumentation: are some norms better than others? *Synthese*, 190(16), 3579–3610.
- Courville, A. C., & Daw, N. D. (2008). The rat as particle filter. In *Advances in neural information processing systems* (pp. 369–376).
- Cover, T. M. (1999). *Elements of information theory*. John Wiley & Sons.
- Cranmer, K., Brehmer, J., & Louppe, G. (2020). The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48), 30055–30062.
- Czerlinski, J., Gigerenzer, G., Goldstein, D. G., et al. (1999). How good are simple heuristics. *Simple heuristics that make us smart*, 97–118.
- Dasgupta, I., & Gershman, S. J. (2021). Memory as a computational resource. *Trends in Cognitive Sciences*, 25(3), 240–251.

- Dasgupta, I., Lampinen, A. K., Chan, S. C., Creswell, A., Kumaran, D., McClelland, J. L., & Hill, F. (2022). Language models show human-like content effects on reasoning. *arXiv preprint arXiv:2207.07051*.
- Dasgupta, I., Schulz, E., & Gershman, S. J. (2017). Where do hypotheses come from? *Cognitive psychology, 96*, 1–25.
- Dasgupta, I., Schulz, E., Tenenbaum, J. B., & Gershman, S. J. (2020). A theory of learning to infer. *Psychological review, 127*(3), 412.
- Dasgupta, I., Wang, J., Chiappa, S., Mitrovic, J., Ortega, P., Raposo, D., ... Kurth-Nelson, Z. (2019). Causal reasoning from meta-reinforcement learning. *arXiv preprint arXiv:1901.08162*.
- Daw, N. D., Courville, A. C., & Dayan, P. (2008). Semi-rational models of conditioning: The case of trial order. *The probabilistic mind*, 431–452.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron, 69*(6), 1204–1215.
- Dayan, P., & Kakade, S. (2000). Explaining away in weight space. *Advances in neural information processing systems, 13*.
- Dobs, K., Martinez, J., Kell, A. J., & Kanwisher, N. (2022). Brain-like functional specialization emerges spontaneously in deep neural networks. *Science advances, 8*(11), eabl8913.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural networks, 15*(4-6), 495–506.
- Duan, Y., Schulman, J., Chen, X., Bartlett, P. L., Sutskever, I., & Abbeel, P. (2016). RL²: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*.
- Dubey, R., Grant, E., Luo, M., Narasimhan, K., & Griffiths, T. (2020). Connecting context-specific adaptation in humans to meta-learning. *arXiv preprint arXiv:2011.13782*.

- Duff, M. O. (2003). Optimal learning: Computational procedures for bayes-adaptive markov decision processes.
- Farquhar, G., Rocktäschel, T., Igl, M., & Whiteson, S. (2017). Treeqn and atreec: Differentiable tree-structured models for deep reinforcement learning. *arXiv preprint arXiv:1710.11417*.
- Feldman, J. (2016). The simplicity principle in perception and cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 7(5), 330–340.
- Feurer, M., & Hutter, F. (2019). Hyperparameter optimization. In *Automated machine learning* (pp. 3–33). Springer, Cham.
- Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning* (pp. 1126–1135).
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2), 127–138.
- Gauvrit, N., Zenil, H., Delahaye, J.-P., & Soler-Toscano, F. (2014). Algorithmic complexity for short binary strings applied to psychology: a primer. *Behavior research methods*, 46(3), 732–744.
- Gauvrit, N., Zenil, H., & Tegnér, J. (2017). The information-theoretic and algorithmic approach to human, animal, and artificial cognition. In *Representation and reality in humans, other living organisms and intelligent machines* (pp. 117–139). Springer.
- Geman, S., & Geman, D. (1984). Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*(6), 721–741.
- Gershman, S. J. (2015). A unifying probabilistic view of associative learning. *PLoS computational biology*, 11(11), e1004567.
- Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, 173, 34–42.
- Gershman, S. J., & Daw, N. D. (2017). Reinforcement learning and episodic memory in

- humans and animals: an integrative framework. *Annual review of psychology*, *68*, 101.
- Gerstenberg, T., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2021). A counterfactual simulation model of causal judgments for physical events. *Psychological review*, *128*(5), 936.
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, *62*, 451–482.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, 148–177.
- Goyal, A., & Bengio, Y. (2022). Inductive biases for deep learning of higher-level cognition. *Proceedings of the Royal Society A*, *478*(2266), 20210068.
- Grant, E., Finn, C., Levine, S., Darrell, T., & Griffiths, T. (2018). Recasting gradient-based meta-learning as hierarchical bayes. In *6th international conference on learning representations, iclr 2018*.
- Griffiths, T. L., Callaway, F., Chang, M. B., Grant, E., Krueger, P. M., & Lieder, F. (2019). Doing more with less: meta-reasoning and meta-learning in humans and machines. *Current Opinion in Behavioral Sciences*, *29*, 24–30.
- Griffiths, T. L., Chater, N., Norris, D., & Pouget, A. (2012). How the bayesians got their beliefs (and what those beliefs actually are): comment on bowers and davis (2012).
- Griffiths, T. L., Daniels, D., Austerweil, J. L., & Tenenbaum, J. B. (2018). Subjective randomness as statistical inference. *Cognitive psychology*, *103*, 85–109.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological science*, *17*(9), 767–773.
- Harlow, H. F. (1949). The formation of learning sets. *Psychological review*, *56*(1), 51.
- Harrison, P., Marjeh, R., Adolphi, F., van Rijn, P., Anglada-Tort, M., Tchernichovski, O., ... Jacoby, N. (2020). Gibbs sampling with people. *Advances in Neural Information Processing Systems*, *33*, 10659–10671.

- Hill, F., Lampinen, A., Schneider, R., Clark, S., Botvinick, M., McClelland, J. L., & Santoro, A. (2020). Environmental drivers of systematicity and generalization in a situated agent. In *International conference on learning representations*. Retrieved from <https://openreview.net/forum?id=Sk1GryBtwr>
- Hinton, G. E., & Van Camp, D. (1993). Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on computational learning theory* (pp. 5–13).
- Hochreiter, S., Younger, A. S., & Conwell, P. R. (2001). Learning to learn using gradient descent. In *International conference on artificial neural networks* (pp. 87–94).
- Hoppe, D., & Rothkopf, C. A. (2016). Learning rational temporal eye movement strategies. *Proceedings of the National Academy of Sciences*, *113*(29), 8332–8337.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural networks*, *2*(5), 359–366.
- Jensen, K. T., Hennequin, G., & Mattar, M. G. (2023). A recurrent network model of planning explains hippocampal replay and human behavior. *bioRxiv*, 2023–01.
- Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? on the explanatory status and theoretical contributions of bayesian models of cognition. *Behavioral and brain sciences*, *34*(4), 169.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., & Saul, L. K. (1999). An introduction to variational methods for graphical models. *Machine learning*, *37*(2), 183–233.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological review*, *80*(4), 237.
- Kanwisher, N., Khosla, M., & Dobs, K. (2023). Using artificial neural networks to ask ‘why’ questions of minds and brains. *Trends in Neurosciences*.
- Kingma, D. P., & Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kirsch, L., & Schmidhuber, J. (2021). Meta learning backpropagation and improving it.

- Advances in Neural Information Processing Systems*, 34.
- Knill, D. C., & Richards, W. (1996). *Perception as bayesian inference*. Cambridge University Press.
- Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information'. *Problems of information transmission*, 1(1), 1–7.
- Kool, W., Cushman, F. A., & Gershman, S. J. (2016). When does model-based control pay off? *PLoS computational biology*, 12(8), e1005090.
- Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247.
- Kruschke, J. (1990). Alcové: A connectionist model of human category learning. *Advances in Neural Information Processing Systems*, 3.
- Kumar, S., Correa, C. G., Dasgupta, I., Marjeh, R., Hu, M., Hawkins, R. D., . . . Griffiths, T. L. (2022). Using natural language and program abstractions to instill human inductive biases in machines. In A. H. Oh, A. Agarwal, D. Belgrave, & K. Cho (Eds.), *Advances in neural information processing systems*. Retrieved from <https://openreview.net/forum?id=buXZ7nIqiwE>
- Kumar, S., Dasgupta, I., Cohen, J., Daw, N., & Griffiths, T. (2020b). Meta-learning of structured task distributions in humans and machines. In *International conference on learning representations*.
- Kumar, S., Dasgupta, I., Cohen, J. D., Daw, N. D., & Griffiths, T. L. (2020a). Meta-learning of structured task distributions in humans and machines. *arXiv preprint arXiv:2010.02317*.
- Kumar, S., Dasgupta, I., Marjeh, R., Daw, N. D., Cohen, J. D., & Griffiths, T. L. (2022). Disentangling abstraction from statistical pattern matching in human and machine learning. *arXiv preprint arXiv:2204.01437*.
- Lake, B. M. (2019). Compositional generalization through meta sequence-to-sequence learning. *Advances in Neural Information Processing Systems*, 32, 9791–9801.

- Lange, R. T., & Sprekeler, H. (2020). Learning not to learn: Nature versus nurture in silico. *arXiv preprint arXiv:2010.04466*.
- Lengyel, M., & Dayan, P. (2007). Hippocampal contributions to control: the third way. *Advances in neural information processing systems*, 20.
- Lewis, D. (1999). Why conditionalize? In *Papers in metaphysics and epistemology* (Vol. 2, p. 403–407). Cambridge University Press. doi: 10.1017/CBO9780511625343.024
- L Griffiths, T., Kemp, C., & B Tenenbaum, J. (2008). Bayesian models of cognition.
- Li, Z., Zhou, F., Chen, F., & Li, H. (2017). Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*.
- Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological review*, 124(6), 762.
- Liu, Z., Luo, P., Wang, X., & Tang, X. (2015, December). Deep learning face attributes in the wild. In *Proceedings of international conference on computer vision (iccv)*.
- Lucas, C. G., Griffiths, T. L., Williams, J. J., & Kalish, M. L. (2015). A rational model of function learning. *Psychonomic bulletin & review*, 22(5), 1193–1215.
- Lueckmann, J.-M., Boelts, J., Greenberg, D., Goncalves, P., & Macke, J. (2021). Benchmarking simulation-based inference. In *International conference on artificial intelligence and statistics* (pp. 343–351).
- Marr, D. (2010). *Vision: A computational investigation into the human representation and processing of visual information*. MIT press.
- McClelland, J. L., Hill, F., Rudolph, M., Baldridge, J., & Schütze, H. (2020). Placing language in an integrated understanding system: Next steps toward human-level performance in neural language models. *Proceedings of the National Academy of Sciences*, 117(42), 25966–25974.
- McClelland, J. L., McNaughton, B. L., & O’Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory.

- Psychological review*, 102(3), 419.
- McCoy, R. T., Grant, E., Smolensky, P., Griffiths, T. L., & Linzen, T. (2020). Universal linguistic inductive biases via meta-learning. *arXiv preprint arXiv:2006.16324*.
- Mercier, H., & Sperber, D. (2017). *The enigma of reason*. Harvard University Press.
- Miconi, T. (2023). A large parametrized space of meta-reinforcement learning tasks. *arXiv preprint arXiv:2302.05583*.
- Miconi, T., Rawal, A., Clune, J., & Stanley, K. O. (2020). Backpropamine: training self-modifying neural networks with differentiable neuromodulated plasticity. *arXiv preprint arXiv:2002.10585*.
- Mikulik, V., Delétang, G., McGrath, T., Genewein, T., Martic, M., Legg, S., & Ortega, P. A. (2020). Meta-trained agents implement bayes-optimal agents. *arXiv preprint arXiv:2010.11223*.
- Mikulik, V., Delétang, G., McGrath, T., Genewein, T., Martic, M., Legg, S., & Ortega, P. (2020). Meta-trained agents implement bayes-optimal agents. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems* (Vol. 33, pp. 18691–18703). Curran Associates, Inc. Retrieved from <https://proceedings.neurips.cc/paper/2020/file/d902c3ce47124c66ce615d5ad9ba304f-Paper.pdf>
- Mitchell, T. M. (1997). *Machine learning* (Vol. 1) (No. 9).
- Molano-Mazon, M., Barbosa, J., Pastor-Ciurana, J., Fradera, M., Zhang, R.-Y., Forest, J., ... others (2022). Neurogym: An open resource for developing and sharing neuroscience tasks.
- Montello, D. R. (2005). *Navigation*. Cambridge University Press.
- Moskovitz, T., Miller, K., Sahani, M., & Botvinick, M. M. (2022). A unified theory of dual-process control. *arXiv preprint arXiv:2211.07036*.
- Müller, S., Hollmann, N., Arango, S. P., Grabocka, J., & Hutter, F. (2021). Transformers can do bayesian inference. *arXiv preprint arXiv:2112.10510*.

- Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.
- Newell, A. (1992). Unified theories of cognition and the role of soar. In *Soar: A cognitive architecture in perspective* (pp. 25–79). Springer.
- Newell, A., Simon, H. A., et al. (1972). *Human problem solving* (Vol. 104) (No. 9). Prentice-hall Englewood Cliffs, NJ.
- Nichol, A., Achiam, J., & Schulman, J. (2018). On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*.
- Nosofsky, R. M. (2011). The generalized context model: An exemplar model of classification. *Formal approaches in categorization*, 18–39.
- Oaksford, M., Chater, N., et al. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press.
- Ortega, P. A., Braun, D. A., Dyer, J., Kim, K.-E., & Tishby, N. (2015). Information-theoretic bounded rationality. *arXiv preprint arXiv:1512.06789*.
- Ortega, P. A., Wang, J. X., Rowland, M., Genewein, T., Kurth-Nelson, Z., Pascanu, R., . . . others (2019). Meta-learning of sequential strategies. *arXiv preprint arXiv:1905.03030*.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., . . . Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems 32* (pp. 8024–8035). Curran Associates, Inc. Retrieved from <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- Pearl, J. (2009). *Causality*. Cambridge university press.
- Piloto, L. S., Weinstein, A., Battaglia, P., & Botvinick, M. (2022). Intuitive physics learning in a deep-learning model inspired by developmental psychology. *Nature human behaviour*, 6(9), 1257–1267.

- Pritzel, A., Uria, B., Srinivasan, S., Badia, A. P., Vinyals, O., Hassabis, D., . . . Blundell, C. (2017). Neural episodic control. In *International conference on machine learning* (pp. 2827–2836).
- Rabinowitz, N. C. (2019). Meta-learners’ learning dynamics are unlike learners?. *arXiv preprint arXiv:1905.01320*.
- Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., . . . others (2019). Machine behaviour. *Nature*, *568*(7753), 477–486.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*, *20*(4), 873–922.
- Reed, S., Zolna, K., Parisotto, E., Colmenarejo, S. G., Novikov, A., Barth-Maron, G., . . . others (2022). A generalist agent. *arXiv preprint arXiv:2205.06175*.
- Rescorla, M. (2020). An improved dutch book theorem for conditionalization. *Erkenntnis*, 1–29.
- Rescorla, R. A. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Current research and theory*, 64–99.
- Rich, A. S., & Gureckis, T. M. (2019). Lessons for artificial intelligence from the study of natural stupidity. *Nature Machine Intelligence*, *1*(4), 174–180.
- Ritter, S., Barrett, D. G., Santoro, A., & Botvinick, M. M. (2017). Cognitive psychology for deep neural networks: A shape bias case study. In *International conference on machine learning* (pp. 2940–2949).
- Ritter, S., Wang, J., Kurth-Nelson, Z., Jayakumar, S., Blundell, C., Pascanu, R., & Botvinick, M. (2018). Been there, done that: Meta-learning with episodic recall. In *International conference on machine learning* (pp. 4354–4363).
- Riveland, R., & Pouget, A. (2022). Generalization in sensorimotor networks configured with natural language instructions. *bioRxiv*, 2022–02.
- Rosenkrantz, R. D. (1992). The justification of induction. *Philosophy of Science*, *59*(4), 527–539.

- Rumelhart, D. E., McClelland, J. L., Group, P. R., et al. (1988). *Parallel distributed processing* (Vol. 1). IEEE New York.
- Sanborn, A., & Griffiths, T. (2007). Markov chain monte carlo with people. *Advances in neural information processing systems*, 20.
- Sanborn, A. N. (2017). Types of approximation for probabilistic cognition: Sampling and variational. *Brain and cognition*, 112, 98–101.
- Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to rational models: alternative algorithms for category learning. *Psychological review*, 117(4), 1144.
- Sanborn, A. N., & Silva, R. (2013). Constraining bridges between levels of analysis: A computational justification for locally bayesian learning. *Journal of Mathematical Psychology*, 57(3-4), 94–106.
- Sancaktar, C., Blaes, S., & Martius, G. (2022). Curious exploration via structured world models yields zero-shot object manipulation. In A. H. Oh, A. Agarwal, D. Belgrave, & K. Cho (Eds.), *Advances in neural information processing systems*. Retrieved from <https://openreview.net/forum?id=NnuYZ1e124C>
- Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., & Lillicrap, T. (2016). Meta-learning with memory-augmented neural networks. In *International conference on machine learning* (pp. 1842–1850).
- Savage, L. J. (1972). *The foundations of statistics*. Courier Corporation.
- Schaul, T., & Schmidhuber, J. (2010). Metalearning. *Scholarpedia*, 5(6), 4650. (revision #91489) doi: 10.4249/scholarpedia.4650
- Schlag, I., Irie, K., & Schmidhuber, J. (2021). Linear transformers are secretly fast weight programmers. In *International conference on machine learning* (pp. 9355–9366).
- Schmidhuber, J. (1987). *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook* (Unpublished doctoral dissertation). Technische Universität München.

- Schramowski, P., Turan, C., Andersen, N., Rothkopf, C. A., & Kersting, K. (2022). Large pre-trained language models contain human-like biases of what is right and wrong to do. *Nature Machine Intelligence*, *4*(3), 258–268.
- Schulz, E., Bhui, R., Love, B. C., Brier, B., Todd, M. T., & Gershman, S. J. (2019). Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences*, *116*(28), 13903–13908.
- Schulz, E., & Dayan, P. (2020). Computational psychiatry for computers. *Isience*, *23*(12), 101772.
- Schulz, E., & Gershman, S. J. (2019). The algorithmic architecture of exploration in the human brain. *Current opinion in neurobiology*, *55*, 7–14.
- Schulz, E., Tenenbaum, J. B., Duvenaud, D., Speekenbrink, M., & Gershman, S. J. (2017, December). Compositional inductive biases in function learning. *Cogn. Psychol.*, *99*, 44–79. doi: 10.1016/j.cogpsych.2017.11.002
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*(4820), 1317–1323.
- Snell, J., Swersky, K., & Zemel, R. (2017). Prototypical networks for few-shot learning. *Advances in neural information processing systems*, *30*.
- Solomonoff, R. J. (1964). A formal theory of inductive inference. part i. *Information and control*, *7*(1), 1–22.
- Stopper, C. M., Maric, T., Montes, D. R., Wiedman, C. R., & Floresco, S. B. (2014). Overriding phasic dopamine signals redirects action selection during risk/reward decision making. *Neuron*, *84*(1), 177–189.
- Strouse, D., McKee, K., Botvinick, M., Hughes, E., & Everett, R. (2021). Collaborating with humans without human data. *Advances in Neural Information Processing Systems*, *34*, 14502–14515.
- Tamar, A., Wu, Y., Thomas, G., Levine, S., & Abbeel, P. (2016). Value iteration networks. *Advances in neural information processing systems*, *29*.

- Tauber, S., Navarro, D. J., Perfors, A., & Steyvers, M. (2017). Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological review*, *124*(4), 410.
- Team, A. A., Bauer, J., Baumli, K., Baveja, S., Behbahani, F., Bhoopchand, A., ... others (2023). Human-timescale adaptation in an open-ended task space. *arXiv preprint arXiv:2301.07608*.
- Team, O. E. L., Stooke, A., Mahajan, A., Barros, C., Deck, C., Bauer, J., ... others (2021). Open-ended learning leads to generally capable agents. *arXiv preprint arXiv:2107.12808*.
- Tenenbaum. (2021). *Joshua Tenenbaum's homepage*.
<http://web.mit.edu/cocosci/josh.html>. ([Online; accessed 9-November-2021])
- Thrun, S., & Pratt, L. (1998). Learning to learn: Introduction and overview. In *Learning to learn* (pp. 3–17). Springer.
- Todorov, E., Erez, T., & Tassa, Y. (2012). Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 5026–5033).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. (2016). Matching networks for one shot learning. *Advances in neural information processing systems*, *29*.
- Wang, J. X. (2021). Meta-learning in natural and artificial intelligence. *Current Opinion in Behavioral Sciences*, *38*, 90–95.
- Wang, J. X., King, M., Porcel, N., Kurth-Nelson, Z., Zhu, T., Deck, C., ... Botvinick, M. (2021). Alchemy: A structured task distribution for meta-reinforcement learning. *CoRR*, *abs/2102.02926*. Retrieved from <https://arxiv.org/abs/2102.02926>
- Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., ...

- Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature neuroscience*, *21*(6), 860–868.
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., . . . Botvinick, M. (2016). Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*.
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, *143*(6), 2074.
- Wolpert, D. H. (1996). The lack of a priori distinctions between learning algorithms. *Neural computation*, *8*(7), 1341–1390.
- Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, *1*(1), 67–82.
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature human behaviour*, *2*(12), 915–924.
- Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T., & Wang, X.-J. (2019). Task representations in neural networks trained to perform many cognitive tasks. *Nature neuroscience*, *22*(2), 297–306.
- Yang, Y., & Piantadosi, S. T. (2022). One model for the learning of language. *Proceedings of the National Academy of Sciences*, *119*(5).
- Yu, T., Quillen, D., He, Z., Julian, R., Hausman, K., Finn, C., & Levine, S. (2020). Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on robot learning* (pp. 1094–1100).
- Zednik, C., & Jäkel, F. (2016). Bayesian reverse-engineering considered as a research strategy for cognitive science. *Synthese*, *193*(12), 3951–3985.
- Zenil, H., Marshall, J. A., & Tegnér, J. (2015). Approximations of algorithmic and structural complexity validate cognitive-behavioural experimental results. *arXiv*

preprint arXiv:1509.06338.