

Commentary title: The hard problem of meta-learning is what-to-learn

1. **Authors:**

Yosef Prat

Ehud Lamm

2. **Institution:** The Cohn Institute for History and Philosophy of Science and Ideas, Tel Aviv University

3. **Institutional address:** The Cohn Institute for History and Philosophy of Science and Ideas, Tel Aviv University, Tel Aviv 6997801, Israel.

Yosef Prat: yosefprat@gmail.com

Ehud Lamm: ehudlamm@post.tau.ac.il

Abstract:

Binz et al. highlight the potential of meta-learning to greatly enhance the flexibility of AI algorithms, as well as to approximate human behavior more accurately than traditional learning methods. We wish to emphasize a basic problem that lies underneath these two objectives, and in turn suggest another perspective of the required notion of ‘meta’ in meta-learning: knowing what to learn.

We postulate that the hard problem in (natural or artificial) intelligence, is the question of “what to learn?”. At the fundamental level, this meta-question is resolved in nature by the evolutionary process. The question of “how to learn?”, which is the focus of the meta-learning framework that is presented in the target article, is not especially easy as well, but it can be captured by devising specific training structures and relevant optimization tasks. In general, it requires to specify the “search space” of possible learning strategies. The hard problem of learning, however, is the identification of the learning task itself (i.e., what, and if, to learn) (Niv, 2019). For instance, a real-life learner observing several specimens of some unknown insect species (following the example in the target article) must first somehow realize that she is required to evaluate the average length of that species, before she begins to tune her evaluation strategies. This is indeed a different meta-task than presented by Binz et al., but its solution is mandatory for any artificial (somewhat-)general intelligence, and it is regularly handled by the brain (Roli, Jaeger, & Kauffman, 2022).

In the quest to devise domain-general learning models, Binz et al. correctly identify the need for diverse (and maybe realistic) training sets. Training a model on many different tasks can achieve high performance in all of them, and maybe even in unrelated, but similar, tasks. Yet, the model will always be constrained by the task-space spanned by its training sets. The major challenge does not lie in amplifying the dimensionality, or variability, of the learned problem, but rather in determining the appropriate objective function. Here, we may be inspired by the observation that biological brains have in general not evolved for their ability to solve a specific task, but, rather, are shaped by the overall success of the organism. On the one hand, evolutionary success obscures the objective of each specific task, since it depends on long-term benefits, that are not always clearly related to short-term behavior. On the other hand, evolutionary success is a broader optimization challenge. A generic model that can both solve a maze and evaluate the average length of a newly identified insect, without being trained specifically on these tasks, must solve the hard problem of what-to-learn in a given context. To build a model that addresses this challenge we cannot handcraft the utility function (or error measurements) of each task separately. The meta-learning requirement thus becomes to learn how to identify the utility in learning, or in performing, each of the given tasks, and more broadly, to identify the task itself. Thus, it is constraining to use training sets and error functions that provide the learner with “correct” answers or feedback for each task separately, as is typically done in supervised, semi-supervised, and reinforcement machine learning. The biological brain is overall domain-general since it is not guided by a task specific “utility function”. Domain-specific processes, such as, maybe, those suggested to process language (Fedorenko & Blank, 2020), demonstrate cases in which natural selection narrowed or “optimized” the task of finding what-to-learn. Other indications may include modularity (Ellefsen, Mouret, & Clune, 2015; Sporns & Betzel, 2016), alongside sensory adaptations (Warrant, 2016), attention biases (Niv et al., 2015), and data acquisition mechanisms (Lotem & Halpern, 2012). Furthermore, in humans, cultural evolution may also adjust task specificity (Heyes, 2018).

The evolutionary process may also explain the limitations of treating cognition as rational, or optimal. Binz et al. suggest that unrealistic aspects of Bayesian models can be mitigated using resource constraints, for which the offered meta-learning framework is suitable. The problem, however, is that human (and other animal) behaviour is not straightforwardly rational, and often appears to defy Bayesian optimization (Tversky & Kahneman, 1981). Moreover, this may not be due to limited resources but because the success of living creatures is determined evolutionarily, rather than by immediate outcomes (Houston, McNamara, & Steer, 2007). When behavioural objectives are considered on an evolutionary scale, it may be revealed that they are (locally) optimal (Kacelnik, 2006), and this includes behaviours that depend upon learning, as is generally assumed in behavioural

ecology. When tasks for which learning is evolutionarily beneficial end up being learned (i.e., when those individuals who learn have higher fitness), natural selection resolves the meta-learning hard problem of what-to-learn (Dunlap & Stephens, 2016). This may bias the things that animals are able to learn to learn, by shaping the parameter search-space (Prat, Bshary, & Lotem, 2022), maybe of the outer learning loop described by Binz et al. These biases are often addressed in the biological learning literature as sub-problems of the what-to-learn problem, and include when-to-learn or from whom-to-learn (Laland, 2004).

We suggest that further advancements in meta-learning thinking require addressing the hard problem of learning as one of their aims. Inspired by (human and nonhuman) biological brains, this should be done by devising overarching objectives for learning algorithms that will enable them to learn what are the learning tasks. In nature, evolution provides some of the solution. Yet, it is not necessary to mimic the evolutionary process per se, but only to acknowledge the generality of evolutionary optimization in the natural world. To this end, it may be better to aspire to simulate nonhuman-animal behavioral studies, rather than psychological assays, since nonhuman animals are trained with no description of the boundaries of their task – they need to realize it by themselves (e.g., when a sparrow learns to relate sand color to food (Ben-Oren, Truskanov, & Lotem, 2022)). Thus, these studies usually contain a direct meta-learning challenge that requires solving the problem of what-to-learn.

Acknowledgements

We thank Yoav Ram for insightful comments on a previous version of the manuscript.

References

- Ben-Oren, Y., Truskanov, N., & Lotem, A. (2022). House sparrows use learned information selectively based on whether reward is hidden or visible. *Animal Cognition*, 25(6), 1545–1555.
<https://doi.org/10.1007/s10071-022-01637-1>
- Dunlap, A. S., & Stephens, D. W. (2016). Reliability, uncertainty, and costs in the evolution of animal learning. *Current Opinion in Behavioral Sciences*, 12, 73–79.
<https://doi.org/https://doi.org/10.1016/j.cobeha.2016.09.010>

- Ellefsen, K. O., Mouret, J.-B., & Clune, J. (2015). Neural modularity helps organisms evolve to learn new skills without forgetting old skills. *PLOS Computational Biology*, 11(4), e1004128. <https://doi.org/doi.org/10.1371/journal.pcbi.1004128>
- Fedorenko, E., & Blank, I. A. (2020). Broca's area is not a natural kind. *Trends in Cognitive Sciences*, 24(4), 270–284. <https://doi.org/10.1016/j.tics.2020.01.001>
- Heyes, C. (2018). *Cognitive gadgets: the cultural evolution of thinking*. Harvard University Press. <https://doi.org/10.2307/j.ctv24trbqx>
- Houston, A. I., McNamara, J. M., & Steer, M. D. (2007). Do we expect natural selection to produce rational behaviour? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485), 1531–1543. <https://doi.org/10.1098/rstb.2007.2051>
- Kacelnik, A. (2006). Meanings of rationality. In S. Hurley & M. Nudds (Eds.), *Rational Animals?* Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198528272.003.0002>
- Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, 32(1), 4–14. <https://doi.org/10.3758/BF03196002>
- Lotem, A., & Halpern, J. Y. (2012). Coevolution of learning and data-acquisition mechanisms: a model for cognitive evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1603), 2686–2694. <https://doi.org/10.1098/rstb.2012.0213>
- Niv, Y. (2019). Learning task-state representations. *Nature Neuroscience*, 22(10), 1544–1553. <https://doi.org/10.1038/s41593-019-0470-8>
- Niv, Y., Daniel, R., Geana, A., Gershman, S. J., Leong, Y. C., Radulescu, A., & Wilson, R. C. (2015). Reinforcement learning in multidimensional environments relies on attention mechanisms. *The Journal of Neuroscience*, 35(21), 8145 LP – 8157. <https://doi.org/10.1523/JNEUROSCI.2978-14.2015>
- Prat, Y., Bshary, R., & Lotem, A. (2022). Modelling how cleaner fish approach an ephemeral reward task demonstrates a role for ecologically tuned chunking in the evolution of advanced cognition. *PLOS Biology*, 20(1), e3001519. <https://doi.org/10.1371/journal.pbio.3001519>
- Roli, A., Jaeger, J., & Kauffman, S. A. (2022). How organisms come to know the world: fundamental limits on artificial general intelligence. *Frontiers in Ecology and Evolution*. <https://doi.org/10.3389/fevo.2021.806283>
- Sporns, O., & Betzel, R. F. (2016). Modular brain networks. *Annual Review of Psychology*, 67(1), 613–

640. <https://doi.org/10.1146/annurev-psych-122414-033634>

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453–458. <https://doi.org/10.1126/science.7455683>

Warrant, E. J. (2016). Sensory matched filters. *Current Biology*, 26(20), R976–R980. <https://doi.org/https://doi.org/10.1016/j.cub.2016.05.042>