70 ideas from “Meta-Learned Models of Cognition” by Binz et al. 2023

1. Constructing computational models by learning them through repeated interactions with an environment…
2. … instead of requiring an a priori specification from a researcher.
3. (Mitchell, 1997) For a given task, performance measure, and training experience, an algorithm is said to learn if its performance at the task improves with experience.
4. Multiple arguments justify Bayesian inference as a normative procedure and thereby its use for rational analysis.
5. These arguments include dutch book arguments, free energy minimization, and performance-based justifications.
6. A meta-learning algorithm is any algorithm that uses its experience to adjust particular 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.
7. Decide on an inner loop or base learning algorithm and determine which of its aspects can be modified.
8. It is possible to meta-learn: (1) hyperparameters for a base learning algorithm such as the number of training epochs, batch sizes, or learning rates …
9. … (2) initial parameters of a neural network that is trained via stochastic gradient descent …
10. … (3) prior distributions for a probabilistic graphical model and (4) entire learning algorithms
11. The meta-learning framework can be used to study how people improve their learning abilities over time.
12. The meta-learning framework can be used as a methodological tool to construct learning algorithms with particular properties of interest.
13. At what time scale does meta-learning take place in humans?
14. To what extent is meta-learning due to task-specific adaptations?
15. To what extent is meta-learning based on developmental or evolutionary processes?
16. Both variational inference and meta-learning involve optimization and require one to define a functional form of the respective distribution
17. However, the optimization process in both approaches involves a different loss function and occurs on different time scales.
18. In meta-learning, the optimization process occurs in an outer-loop learning process, but not during the actual learning itself.
19. To update how a meta-learned model makes predictions in light of new data, we only have to perform a single forward pass through the network.
20. Meta-learned inference only demands samples from the data-generating distribution to meta-learn an approximately Bayes-optimal learning algorithm.
21. Many different notions of what constitutes a computational resource have been suggested such as…
22. … memory, thinking time, or physical effort.
23. Algorithmic complexity: the number of bits needed to implement the algorithm.
24. Computational complexity: the space, time, or effort required to execute it.

25. Limiting the complexity of network weights places a constraint on the algorithmic complexity.
26. Limiting the complexity of activations places a constraint on the computational complexity.
27. (Wang et al, 2018) Prefrontal circuits may constitute a meta-reinforcement learning system.
28. Meta-learning strives to learn a faster, inner-loop learning algorithm…
29. … via an adjustment of neural network weights in a slower, outer-loop learning process.
30. “Learning to Infer,” 2020: trained a neural network on a distribution of probabilistic inference problems while controlling for the number of hidden units.
31. This model – when restricted to just a single hidden unit – captured biases in human reasoning, …
32. … including conservatism bias and base rate neglect.
33. Traditional rational process models struggle to capture that human strategy selection is…
34. … context-dependent even before receiving any direct feedback signal.
35. From Lake 2019: “tackling a series of changing learning problems rather than iterating through a static data set”
36. Model-free learning algorithms directly adjust their strategies using observed outcomes.
37. Model-based learning algorithms learn about the transition and reward probabilities of an environment, which are then used for downstream reasoning tasks.
38. Having a model of the world acts as the basis for causal reasoning
39. Meta-learned inference shifts most of the compute burden from run-time to training time, …
40. … which is advantageous when training time is ample and fast answers are needed at run-time, …
41. … and may therefore explain how people can perform such intricate computations within a reasonable time frame.
42. The ability to perform temporally extended planning using imagined rollouts
43. More planning early-on and with an increased distance to the goal
44. Patterns of hippocampal replays resemble the imagined rollouts
45. To integrate observed information into existing knowledge
46. To actively determine what information to sample
47. Cognitive control is described as the processes behind the ability to adapt to task-specific demands
48. Gradient-based learning is at the heart of the Rescorla-Wagner model
49. Exemplar-based models categorize a new instance based on the estimated similarity between that instance and previously seen examples.
50. In the reinforcement learning literature, episodic memory is the ability to store and recollect states or trajectories for later use.
51. Ritter et al. (2018) builds on the Neural Episodic Control paper and …
52. … uses a differential neural dictionary for episodic recall in the context of meta-learning.
53. Integrating gradient-based learning, exemplar and prototype-based reasoning, and episodic memory into a meta-learned model

54. (Newell, 1992): “Unified theories of cognition are the only way to bring this wonderful, increasing fund of knowledge under intellectual control.”
55. Ideally, such a unified theory should manifest itself in a domain-general cognitive model that…
56. … cannot only solve prediction tasks but is also capable of human-like decision-making,...
57. … category learning, navigation, problem-solving, and so on.
58. We consider the meta-learning framework the ideal tool for accomplishing this goal as it allows us to…
59. … compile arbitrary assumptions about an agent’s beliefs of the world (arguments 1 and 2)...
60. … and its computational architecture (arguments 3 and 4) into a cognitive model.
61. How should a data-generating distribution that contains many different problems be constructed?
62. Instead of using Bayesian Inference to obtain the posterior predictive distribution, …
63. … we teach a general purpose function approximator to do this inference.
64. The meta-learned predictive distribution takes a predetermined functional form whose parameters are given by the network outputs.
65. Model-based reasoning capabilities can emerge internally in a meta-learned model – if they are beneficial for solving the encountered problem.
66. Taking contextual cues into consideration is a vital aspect of cognitive control.
67. Meta-learning can produce approximately optimal learning algorithms even if …
68. … it is not possible to phrase the corresponding inference problem in the first place.
69. Meta-learning makes it easy to manipulate a learning algorithm’s complexity …
70. … and can therefore be used to construct resource-rational models of learning.