



Behavioral and Brain Sciences

Article contents

- Abstract
- Financial support
- Competing interests

The reinforcement metalearner as a biologically plausible meta-learning framework

Published online by Cambridge University Press: **23 September 2024**

Tim Vriens , Mattias Horan , Jacqueline Gottlieb and Massimo Silvetti

Show author details

Commentary Related commentaries Metrics

Abstract

We argue that the type of meta-learning proposed by Binz et al. generates models with low interpretability and falsifiability that have limited usefulness for neuroscience research. An alternative

Access

Linked content

In response to: **Meta-learned models of cognition**

[Related commentaries \(22\)](#)

[Author response](#)

Related content


AI-generated results: by UNSILO

Article

Dopamine, schizophrenia, mania, and depression: Toward a unified hypothesis of cortico-striatopallido-thalamic function

References

approach to meta-learning based on hyperparameter optimization obviates these concerns and can generate empirically testable hypotheses of biological computations.

Type	Open Peer Commentary
Information	Behavioral and Brain Sciences, Volume 47, 2024, e168 DOI: https://doi.org/10.1017/S0140525X24000219 
Copyright	Copyright © The Author(s), 2024. Published by Cambridge University Press

Binz et al. describe four different meta-learning approaches and focus on the last one – methods for learning arbitrary new tasks without the need for a priori hypotheses about brain or cognitive architectures. They show that this approach can be implemented in recurrent neural networks (RNNs) that are universal approximators (Hornik, Stinchcombe, & White, 1989), and argue that it is powerful in producing Bayesian (near-optimal) learning in an arbitrarily large set of cognitive tasks. While acknowledging the power of the proposed framework for artificial intelligence (AI), we question its usefulness in cognitive and neuroscience research. We argue that an alternative approach of hyperparameter optimization (which

Neal R. Swerdlow and George F. Koob
[Behavioral and Brain Sciences](#)
Published online: 4 February 2010

Article**Don't leave the "psych" out of neuropsychology**

Jeffrey A. Gray and Ilan Baruch
[Behavioral and Brain Sciences](#)
Published online: 4 February 2010

Article**Where have all the peptides gone?**

Peter W. Kalivas
[Behavioral and Brain Sciences](#)
Published online: 4 February 2010

Article**Toward a neurological psychiatry**

Michel Le Moal
[Behavioral and Brain Sciences](#)
Published online: 4 February 2010

Article**Intracellular considerations in models of psychopathology**

was first proposed by Doya, 2002, and is mentioned but not discussed by Binz et al.) is far more powerful for this role.

To be valuable for empirical research, a computational framework should generate models that are interpretable in neurocognitive terms and make predictions that can be falsified or confirmed through empirical tests. The internal computations used by the models should be analogous to those of neurocognitive systems (e.g., attention, memory, valuation, etc.; e.g., Castelvechi, 2016), and predict activity patterns that can be empirically validated. The framework advocated by Binz et al. has neither property, and instead generates models that are governed by immense numbers of free parameters (up to billions) and are not interpretable in cognitive terms, amounting to a “black box” data-driven approach.

A hyperparameter optimization approach alleviates these concerns by constraining the models it generate to emulate biologically plausible architectures. This allows for formulating and testing mechanistic hypotheses that are based in established literature. The reinforcement meta-learner (RML) model is a good illustration of this framework in the context of executive function (Silvetti, Vassena, Abrahamse, & Verguts, 2018).

Douglas L. Chute

[Behavioral and Brain Sciences](#)

Published online: 4 February 2010

Article

Unified theories of psychoses and affective disorders: Are they feasible without accurate neural models of cognition and emotion?

Anthony Phillips

[Behavioral and Brain Sciences](#)

Published online: 4 February 2010

Article

Toward a unified neuropsychiatric hypothesis

Neal R. Swerdlow and George F.

Koob

[Behavioral and Brain Sciences](#)

Published online: 4 February 2010

Article

The neuropathology of schizophrenia, mania, and depression: Diseases of cognitive initiation and switching?

Eliot L. Gardner

[Behavioral and Brain Sciences](#)

Published online: 4 February 2010

Consistent with abundant empirical evidence on biological executive circuits (e.g., Shackman et al., 2011; Silvetti, Seurinck, van Bochove, & Verguts, 2013; Varazzani, San-Galli, Gilardeau, & Bouret, 2015; Yarkoni, Poldrack, Nichols, Van Essen, & Wager, 2011), the RML emulates interactions between the medial prefrontal cortex (MPFC) and two catecholamine nuclei – the ventral tegmental area, releasing dopamine (DA), and the locus coeruleus, releasing norepinephrine (NE). The MPFC module monitors reward rates conveyed by DA and, when detecting a “need for control” (e.g., a decrease in the rates), calls for the release of NE and DA. In turn, these neurotransmitters are broadcast to task-specific cognitive modules and enhance their efficiency, thereby restoring performance and reward rates. The MPFC registers a boost of neurotransmitter release as a cost and uses Bayesian and reinforcement-learning (RL) optimization to learn control settings that maximize rewards while minimizing costs. The RML thus uses traditional Bayesian/RL optimization frameworks to simultaneously regulate motor input and internal cognitive computations, thus modeling both first-order performance and its executive (meta-level) control.

Article**The ghost in the machine: What if the midbrain output is excitatory?**

James H. Fallon

[Behavioral and Brain Sciences](#)

Published online: 4 February 2010

Article**The relevance of feedforward loops**

A. R. Cools

[Behavioral and Brain Sciences](#)

Published online: 4 February 2010

Recent studies have shown that the RML explains empirical findings that have long stumped traditional frameworks, including nonstandard reward modulations in visual areas (Horan, Daddaoua, & Gottlieb, 2019; Silvetti, Lasaponara, Daddaoua, Horan, & Gottlieb, 2023) and curiosity – the intrinsic desire to obtain information in the absence of instrumental rewards (Daddaoua, Lopes, & Gottlieb, 2016; Horan et al., 2019; Silvetti et al., 2023). By monitoring the volatility of the environment, the RML provides a meta-learning-based explanation of the empirical finding of volatility-sensitive learning rates (Silvetti et al., 2013, 2018). Moreover, when coupled to modules emulating memory, motor output, decision making, or attention, the RML reproduces a wide array of behavioral and neural results related, respectively, to memory capacity, motor effort, adaptive regulation of learning rates, and instrumental or curiosity-driven information gathering (Silvetti et al., 2018, 2023). Thus, despite its biologically constrained architecture, the RML gains considerable flexibility and generalizability because it can control different task-specific cognitive computations.

Because the RML uses a biologically plausible architecture with a parsimonious parameter set, it generates a rich set of novel

predictions that can be tested against empirical data. These predictions involve possible relationships between behavior and neural activity, between neural activity and neurotransmitter release, and between activity in different brain structures. Existing versions of the RML make predictions about individual computations (e.g., how much memory effort to engage in a particular context) but future versions can be extended to probe how the brain arbitrates between computations (e.g., how it trades-off between relying on memory versus acquiring new sensory information when performing a task).

In conclusion, different meta-learning approaches can differ greatly in their comparative strengths. The entirely unconstrained approach discussed by Binz et al. may be desirable for AI applications where there is no need for biological constraints, for example, when developing an algorithm for a self-driving car, or optimizing planning in multiple tasks. In contrast, we believe that a biologically constrained meta-learning framework is vastly superior for advancing cognitive and neuroscience research (Marblestone, Wayne, & Kording, 2017). Such a biologically constrained framework is grounded in the neuroscientific literature, and can

generate testable and falsifiable hypotheses about neurobiological processes underlying cognitive function.

Acknowledgments

Tim Vriens is a PhD student enrolled in the National PhD in Artificial Intelligence, XXXVII cycle, course on Health and life sciences, hosted by Università Campus Bio-Medico di Roma, Italy.

Financial support

M. S. is funded by the Italian Ministry of University and Research, PRIN 2022 program, Grant No. 64.20227MPSEH. M. H. is supported via the Sainsbury Wellcome Centre PhD Programme and has received grants from Reinholdt W. Jorck og Hustrus Fond, Knud Højgaards Fond and Anglo-Danish Society

Competing interests

None.

References



Castelvecchi, D. (2016). Can we open the black box of AI? *Nature*, 538, 20–23.

<https://doi.org/10.1038/538020a> [CrossRef](#) [Google Scholar](#) [PubMed](#)



Daddaoua, N., Lopes, M., & Gottlieb, J. (2016). Intrinsically motivated oculomotor exploration guided by uncertainty reduction and conditioned reinforcement in non-human primates. *Scientific Reports*, 6(1), Article 1.

<https://doi.org/10.1038/srep20202> [CrossRef](#) [Google Scholar](#) [PubMed](#)



Doya, K. (2002). Metalearning and neuromodulation. *Neural Networks*, 15(4–6), 495–506.

[https://doi.org/10.1016/s0893-6080\(02\)00044-8](https://doi.org/10.1016/s0893-6080(02)00044-8) [CrossRef](#) [Google Scholar](#) [PubMed](#)



Horan, M., Daddaoua, N., & Gottlieb, J. (2019). Parietal neurons encode information sampling based on decision uncertainty. *Nature Neuroscience*, 22(8), 1327–1335.

<https://doi.org/10.1038/s41593-019-0440-1> [CrossRef](#) [Google Scholar](#) [PubMed](#)



Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5), 359–366.

[https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8) [CrossRef](#) [Google Scholar](#)



Marblestone, A. H., Wayne, G., & Kording, K. P. (2017). Understand the cogs to understand cognition. *Behavioral and Brain Sciences*, 40, e272.

<https://doi.org/10.1017/S0140525X17000218> [CrossRef](#) [Google Scholar](#) [PubMed](#)



Shackman, A. J., Salomons, T. V., Slagter, H. A., Fox, A. S., Winter, J. J., & Davidson, R. J. (2011). The integration of negative affect, pain, and cognitive control in the cingulate cortex. *Nature Reviews. Neuroscience*, 12(3), 154–167.

<https://doi.org/10.1038/nrn2994> [CrossRef](#) [Google Scholar](#) [PubMed](#)



Silvetti, M., Lasaponara, S., Daddaoua, N., Horan, M., & Gottlieb, J. (2023). A reinforcement meta-learning framework of executive function and information demand. *Neural Networks*, 157, 103–113.

<https://doi.org/10.1016/j.neunet.2022.10.004> [CrossRef](#) [Google Scholar](#) [PubMed](#)



Silvetti, M., Seurinck, R., van Bochove, M., & Verguts, T. (2013). The influence of the noradrenergic system on optimal control of neural plasticity. *Frontiers in Behavioral Neuroscience*, 7, 160.
<https://www.frontiersin.org/articles/10.3389/fnbeh.2013.00160> [CrossRef](#) [Google Scholar](#) [PubMed](#)



Silvetti, M., Vassena, E., Abrahamse, E., & Verguts, T. (2018). Dorsal anterior cingulate-brainstem ensemble as a reinforcement meta-learner. *PLoS Computational Biology*, 14(8), e1006370.
<https://doi.org/10.1371/journal.pcbi.1006370> [CrossRef](#) [Google Scholar](#) [PubMed](#)



Varazzani, C., San-Galli, A., Gilardeau, S., & Bouret, S. (2015). Noradrenaline and dopamine neurons in the reward/effort trade-off: A direct electrophysiological comparison in behaving monkeys. *The Journal of Neuroscience*, 35(20), 7866–7877.
<https://doi.org/10.1523/JNEUROSCI.0454-15.2015> [CrossRef](#) [Google Scholar](#) [PubMed](#)



Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., & Wager, T. D. (2011). Large-scale automated synthesis of

human functional neuroimaging data. Nature Methods, 8(8), 665–670.

<https://doi.org/10.1038/nmeth.1635> [CrossRef](#) [Google Scholar](#) [PubMed](#)



Our Site

[Accessibility](#)

[Contact & Help](#)

[Legal Notices](#)

Our Platforms

[Cambridge Core](#)

[Cambridge Open Engage](#)

[Cambridge Higher Education](#)

Our Products

[Journals](#)

[Books](#)

[Elements](#)

[Textbooks](#)

[Courseware](#)

Join us online

Location

[Rights & Permissions](#) [Copyright](#) [Privacy Notice](#) [Terms of Use](#) [Cookies Policy](#)

Cambridge University Press 2024