Biological Cybernetics (2023) 117:407–409
https://doi.org/10.1007/s00422-023-00979-4

Biological Cybernetics

FORUM

Four attributes of intelligence, a thousand questions

Matthieu Bardal[1] · **Eric Chalmers**[1]

Received: 7 September 2023 / Accepted: 8 November 2023 / Published online: 7 December 2023

Abstract

Jeff Hawkins is one of those rare individuals who speaks the languages of both AI and neuroscience. In his recent book, "A Thousand Brains: A New Theory of Intelligence", Hawkins proposes that current learning algorithms lack four attributes which will be necessary for true machine intelligence. Here we demonstrate that a minimal learning system which satisfies all four points can be constructed using only simple, classical machine learning techniques. We illustrate that such a system falls short of biological intelligence in some important ways. We suggest that Hawkins' list is a useful model, but the "recipe" for true intelligence—if there is one—may not be so easily defined.

Keywords Artificial intelligence · Reinforcement learning · Machine intelligence

In his recent book, "A Thousand Brains: A New Theory of Intelligence", Jeff Hawkins proposes that four attributes of intelligence are key to creating truly intelligent machines [1]:

1. **Learning many models**. Various aspects of the world are modeled separately. One model might describe how to anticipate and catch a baseball, while another describes how to catch a frisbee. The most relevant models have the greatest influence on behavior in a given situation.
2. **Learning through movement and sensation**. Agents should learn by moving sensors around the thing being modeled, just as we might learn about a small object by looking at it from every side. The meaning of "movement" depends on how the agent is embodied and situated (a web-crawling AI "moves" through web pages).
3. **Using general-purpose reference frames**. Hippocampal circuitry (grid cells, place cells, etc.) that evolved to serve spatial navigation now likely supports all kinds of learning [2]. In Hawkins' view, most knowledge is stored—and most reasoning performed—using generic map-like structures.
4. **Continuous learning**—unlike today's artificial neural networks, which typically learn new things by modifying connection weights (essentially overwriting previous knowledge).

This list of attributes is an insightful and useful model of intelligence. But God is in the details, and these four points abstract away so much complexity, they may not be a sufficient "recipe" for intelligence. To illustrate this limitation, we construct a minimal learning agent which satisfies all four points at a fundamental level, using only classical machine learning techniques (challenging Hawkins' claim that the machine learning field lacks these four attributes). We then illustrate a few questions that the list leaves unanswered.

We envision our agent living in a "world" containing many tasks. Each task is described by a Markov Decision Process (MDP), following the usual reinforcement learning paradigm [3]. When immersed in a task, the agent can sense its current state, perform actions that move it to a new state, and sense again (thus implementing the principle of learning through movement and sensation). The agent must learn how to achieve a goal located in one of the states. MDPs like this are perhaps the simplest possible metaphor for the sequential decision making of life.

Our minimal learning agent is a reinforcement learner based on the Dyna algorithm [4]. This algorithm maintains tables that track state transition and reward events. After sufficient experience, these tables can estimate state transition probabilities and expected rewards for particular actions, and

✉ Eric Chalmers
echalmers@mtroyal.ca

Matthieu Bardal
mbard9211@mtroyal.ca

1 Department of Mathematics and Computing, Mount Royal University, 4825 Mt Royal Gate SW, Calgary, AB T3E6K6, Canada

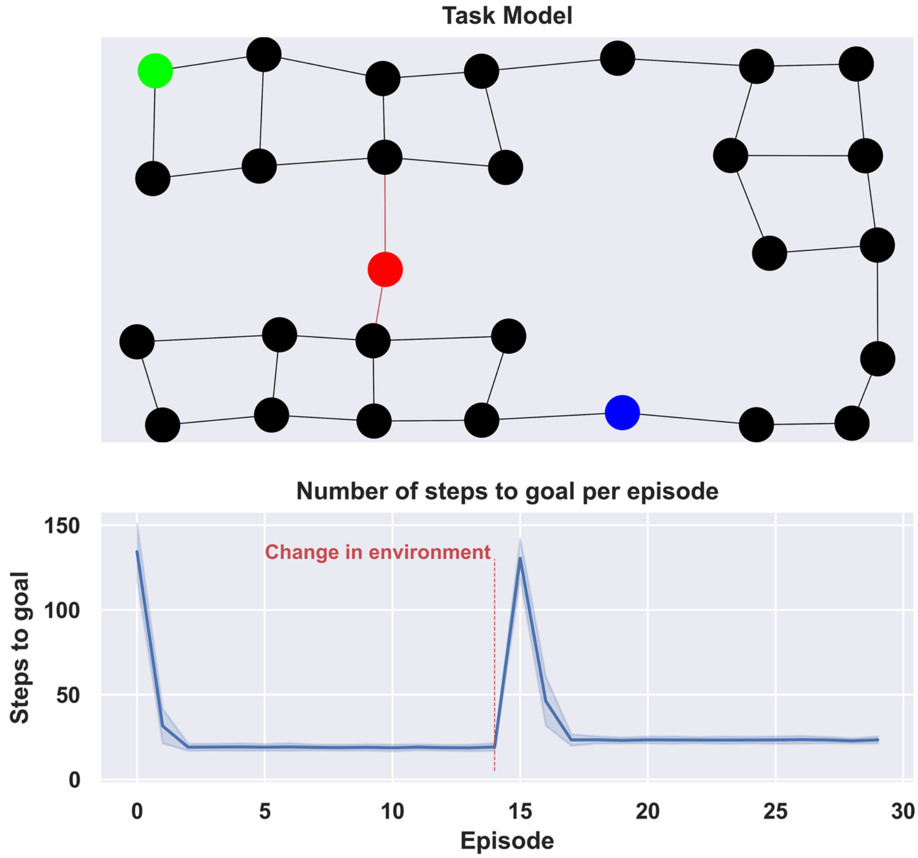

Fig. 1 The agent learns a model of a task (above) that includes the shortest route between blue start and green goal states. If the red state is deleted, the agent can learn a new model representing the altered task, but this takes time (below). Better would be to flexibly *adapt* the old model, but how?

so together they serve as a model of the underlying MDP. These tabular models are generic structures that can track reward and transition probabilities for any MDP in the agent's "world", and so they are an abstraction of the general-purpose reference frames Hawkins describes. We then use the prioritized sweeping algorithm [5] to estimate the value (expected discounted reward) of executing a particular action from a particular state, given the tabular model.

We extend this minimal learning system with a memory of recent state transitions, and the ability to learn and store multiple tabular models as in Chalmers et al. [6]. At each decision, the agent computes the likelihood of its recent experiences under each model, and uses the one that seems most relevant. If none of the existing models match the recent experiences, a new one is started. In this way the agent implements the principles of continuously learning multiple models.

This minimal learning system learns the shortest route to a task's goal state. If we then create a blockage by deleting one of the MDP states along this route, the agent can learn a new model to represent the new best route. But in our minimal learning system this requires learning a *completely* new model, which takes time and effort (Fig. 1). A more intelligent, more animal-like response to the blocking problem would be to flexibly *adapt* the contents of the existing model to the new situation—trying the other, counterclockwise route to the goal before resorting to learning a new model. This example is quite simple, but one can imagine more dramatic environmental changes (multiple obstacles, changes in reward locations, etc.) that would require more sophisticated adaptation of previously-learned models. A later paper co-authored by Hawkins does acknowledge the importance of this kind of flexibility [7], but how does it work?

Along with the "blocking" problem we could consider the related "shortcutting" problem. If a new transition appears in the MDP should the agent explore it, hoping it represents a shortcut to a desired state? Or exploit a tried-tested-and-true behavior? Animals can sometimes respond intelligently to the appearance of apparent shortcuts, using a cognitive map to predict whether the shortcut will lead in the right "direction". To be fair, Hawkins' point 3 surely envisions precisely such a sophisticated cognitive map [2, 8]. But how is egocentric sensory information transformed to build an allocentric cognitive-map-style representation, and how is the map used to modulate exploration vs exploitation? The list of four attributes does not seem to offer guidance on these details.

Perhaps the tabular models are too simple to support such intelligent behaviors, and function approximation through neural networks is required. But then the key to machine intelligence would lie not only in the four attributes, but also (or perhaps instead) in the neural architecture. And if the issues of flexible model adaptation, cognitive mapping, and exploration/exploitation are solved, how then is the agent to compose skills hierarchically [9], communicate with other agents [10], or use a cognitive map to make plans [11] as biological agents do? Of course, these phenomena are not complete mysteries: many computational accounts of each phenomena have been proposed. The point is that those proposals are part of a rich discussion about intelligence which seems impossible to contain in a list of four attributes.

Considering that this list of attributes may be insufficient as a recipe for intelligence, the question remains: what are the missing secret ingredients? It's been suggested that artificial intelligence falls short of biological intelligence in areas like hierarchical composition and transfer learning [9], meta learning [12], few-shot learning [13], and energy efficiency [14]. Would these features complete the recipe? Or should we focus instead on the more fundamental question: what *is* a truly intelligent machine? The evolution of the Turing test and its variants and alternatives over the years [15] suggests that this question is just as slippery—and sometimes "passing" such a test only reveals how far we have yet to go [16]. One begins to wonder; is writing a recipe for intelligence even possible?

Acknowledgements The authors would like to acknowledge financial support from Mount Royal University.

Author Contributions MB wrote and executed experiments, created the figures, and contributed to manuscript writing. EC conceived of the experiments and contributed to manuscript writing.

Declarations

Conflict of interest The authors have no competing interests to declare.

References

1. Hawkins J, Dawkins R (2021) A thousand brains: a new theory of intelligence. Basic Books, New York
2. Hawkins J, Lewis M, Klukas M, Purdy S, Ahmad S (2019) A framework for intelligence and cortical function based on grid cells in the neocortex. Front Neural Circuits 12:121
3. Sutton RS, Barto AG (2018) Reinforcement learning: an introduction, 2nd edn. MIT Press, Cambridge
4. Sutton RS (1991) Planning by incremental dynamic programming. In: Machine learning proceedings 1991. Elsevier, pp 353–357
5. Moore AW, Atkeson CG (1993) Prioritized sweeping: reinforcement learning with less data and less time. Mach Learn 13(1):103–130
6. Chalmers E , Luczak A, Gruber AJ (2016) Computational properties of the hippocampus increase the efficiency of goal-directed foraging through hierarchical reinforcement learning. Front Comput Neurosci 10:128
7. Anthony Z, Sean E, Blake R, Bence Ö, Yoshua B, Kwabena B, Matthew B, Dmitri C, Anne C, Claudia C, James DC, Surya G, Jeff H, Konrad K, Alexei K, Yann LC, Timothy L, Adam M, Bruno O, Alexandre P, Cristina S, Terrence S, Eero S, Sara S, David S, Tolias Andreas S, Doris T (2023) Catalyzing next-generation Artificial Intelligence through NeuroAI. Nat Commun 14(1):1597
8. Kjell JH, Ahmad S (2021) A thousand brains: toward biologically constrained AI. SN Appl Sci 3(8):743
9. Eppe M, Gumbsch C, Kerzel M, Nguyen PDH, Butz MV, Wermter S (2022) Intelligent problem-solving as integrated hierarchical reinforcement learning. Nat Mach Intell 4(1):11–20
10. Elio T (2009) An investigation of the evolutionary origin of reciprocal communication using simulated autonomous agents. Biol Cybern 101(3):183–199
11. Daw N (2012) Model-based reinforcement learning as cognitive search: neurocomputational theories
12. Matthew B, Sam R, Wang Jane X, Zeb K-N, Charles B, Demis H (2019) Reinforcement learning, fast and slow. Trends Cogn Sci 23(5):408–422
13. Henry S, Karina V, Matthew C, Marta H (2019) The limits of machine intelligence. EMBO Rep 20(10):e49177
14. Korteling JH, van de Boer-Visschedijk GC, Blankendaal RA, Boonekamp RC, Eikelboom AR (2021) Human-versus artificial intelligence. Front Artif Intell 4:622364
15. Shane L, Marcus H (2007) Universal intelligence: a definition of machine intelligence. Minds Mach 17(4):391–444
16. Sebastian L, Stephan W, Alexander B, Grégoire M, Wojciech S, Klaus-Robert M (2019) Unmasking Clever Hans predictors and assessing what machines really learn. Nat Commun 10(1):1096