

Apply rich psychological terms in AI with care

Henry Shevlin and Marta Halina

Many of the most promising approaches in cognitive science seek to explicate notions like perception, belief, and motivation in information processing terms (Simon 1980). A similar move from information processing to psychology is occurring in artificial intelligence (AI) research and machine learning. Currently, few people working in AI would literally attribute beliefs, thoughts, or feelings to machines. However, as new techniques extend the capabilities of artificial systems, it has become increasingly common to use psychological terms to describe processing architectures (Weber et al. 2017; Kim et al. 2018; Rabinowitz et al. 2018; Shylaja et al. 2017; Abate 2018).

Such use of psychological terms in AI research may be in many cases justified and unproblematic. Some terms with both narrow computational and psychological meanings, such as “memory” and “reinforcement”, have clearly distinct histories in psychology and artificial intelligence and their senses are unlikely to be confused. Others such as “learning” and “behavior” straddle psychology and artificial intelligence, but are broad enough that there is little reason to quibble with their use to describe machine capacities. However, we argue here that there is a third kind of more robustly psychological concepts—including notions like awareness, perception, agency, and theory of mind—that have rich and complex histories in cognitive science. We suggest that these terms – which we call rich psychological concepts – require greater caution when employed to describe the capabilities of machine intelligences.

(1) From psychology to AI and back: fostering joint understanding

Our first concern for the use of rich psychological terms in artificial intelligence comes from the fact that for many such concepts (for example, theory of mind, or perception) there is considerable controversy in cognitive science as to whether and when they can be applied even to biological organisms. In order to maintain consistency and clear communication across different branches of cognitive science, it is thus key for artificial intelligence researchers who wish to employ these concepts to do so with an awareness of the standards of evidence and proof applied to them by researchers concerned with biological intelligence. More positively, we would also

suggest that many in the machine learning community might gain insights into the development of sophisticated artificial intelligences by paying close attention to these debates.

Consider, for example, the debate concerning theory of mind (ToM) in nonhuman animals. In 1978, Premack and Woodruff asked whether chimpanzees have a ToM or the ability to attribute mental states to others—such as goals, knowledge, and beliefs (Premack and Woodruff 1978). The challenges involved in answering this question quickly became apparent, however, as researchers realised there were multiple available explanations for subjects' successful performance on ToM tasks. A subject could pass a false belief task, for example, by applying a behavioural rule, such as “an agent will choose the container that had the preferred reward placed into it while that agent was present” (Lurz 2011). In grappling with this problem, psychologists developed alternative paradigms and standards for testing ToM. New experimental approaches, such as experience–projection tasks were designed to dissociate the attribution of mental states from the application of behavioural rules (Karg et al. 2015; Halina 2015). Forty years and numerous studies later, a consensus is beginning to emerge that chimpanzees have the ability to attribute some mental states to others (Krupenye et al. 2016), although not everyone agrees (Heyes 2017). Crucially for our purposes, however, the above challenges have led to a more nuanced understanding of ToM and how to test for it.

Recently, AI researchers have been exploring the idea of building an artificial theory of mind (Görür et al. 2017; Winfield 2018; Rabinowitz et al. 2018). The results of these studies are impressive, with machines capable of predicting the behaviour of agents they have never seen before. Researchers working in this area draw heavily on contemporary work in cognitive science in order to guide the development and assessment of machine ToM (Winfield 2018; Rabinowitz et al. 2018). Engaging with cognitive science in this way provides AI developers with a rich database of theories, mechanisms, behaviours, and experimental tasks (Hassabis et al. 2017).

We think this is the right approach and would like to encourage more engagement of this kind. In the case of ToM research, for example, an artificial system like that developed by Rabinowitz and colleagues (2018) is an excellent example of how well an artificial system can learn to predict social behaviour. Within the context of a simple gridworld, their “Theory of Mind neural network” or ToMnet was able to pass traditional ToM tasks (including a false belief task) after being trained on no more than a population of behaving agents. The fact that psychologists would characterize ToMnet as engaging in behavior reading, rather than

attributing mental states, does not diminish the significance of this development. Quite the contrary: one of the extraordinary implications of this study is that it demonstrates that an agent can pass traditional ToM tests using behaviour reading alone. Psychologists have hypothesized that this is possible, but have not demonstrated it empirically, as they lack full knowledge of an organism's learning history. AI research such as this can provide greater insight into the learning and cognitive mechanisms potentially involved in ToM in humans and other animals. In this case, perhaps one need not represent mental states in order to achieve sophisticated social behaviour after all.

While we see this example as constituting a success story for the use of rich psychological terms by AI researchers, the model it provides has not yet been consistently adopted by the field. While we do not wish to single out individual publications for practices that are widespread, we would note a quick glimpse at almost any leading machine learning publication will reveal frequent reference made to rich psychological concepts like understanding, motivation, and even creativity, without the same care and attention exhibited by studies like those mentioned above. This risks miscommunication and cross-purpose in interdisciplinary attempts to understand the basis of psychological capacities and in comparing the capacities of artificial and biological intelligences.

A key point to acknowledge here, of course, is that those working in artificial intelligence frequently have a kind of privileged access to the inner workings of the cognitive systems they build, an advantage not enjoyed in the field of animal behavior. Thus, perhaps AI researchers have special insight into cognitive mechanisms. This advantage should not be overstated, however. For one, there are well-known issues with the interpretability of artificial neural networks, such that the precise cognitive structures of artificial systems may not be fully transparent even to their creators (Selbst and Barocas 2018). More fundamentally, however, there is considerable debate concerning the underlying cognitive structure of capacities such as theory of mind or perception even in adult humans, with extensive fine-grained behavioural measures being our only current reliable test for their presence (Schaafsma et al. 2015). Hence simply knowing *how* a machine arrives at a given output does not automatically warrant the conclusion that the relevant process can be understood under the same psychological concept as that applied to adult humans. We therefore suggest that machine learning experts continue to engage closely with the standards of behavioral testing and theoretical wrangling found in other

domains of cognitive science, and, in so doing, contribute their expertise and insights to these debates.

(2) Normative dimensions of rich psychological terms

A second reason for diligence, we would suggest, is that some – and perhaps all – rich psychological terms are ‘thick concepts’ (Hare 1952), bringing with them normative considerations and hence the potential for ethical and legal ramifications. This is not the case for more practical concepts such as memory or learning, of course, but when we speak of a system as having intrinsic motivations, being an agent, or displaying imagination, there is a risk of implicitly licensing inferences about the moral or intellectual status of the system in question. Thus when authors speak of “intrinsically motivated [artificial] agents” (Kulkarni et al. 2016), a natural but unwarranted (and doubtless unintended) assumption might be that the systems in question literally possess goals, desires, and perhaps even a form of moral responsibility. Similarly, when systems that learn with model-based reinforcement learning are described as having “imagination” (Weber et al. 2017), an incautious reader may leap to the conclusion that the system in question possesses a capacity for the kinds of intellectual insight we associate with some of the most dramatic feats of human intelligence.

The machine learning community should not be held responsible for media sensationalism or misinterpretations of its claims, of course. However, we would suggest a pragmatic interest for researchers in avoiding terminology easily associated with human-level intellectual capacities or moral or agential status. Given the highly technical nature of most artificial systems, policymakers and the scientific media are often ill-placed to evaluate when such terms are used literally and when they are employed in more restrictive senses. Such gaps in understanding between the technical communities and the general public thus risks inviting moral panics, ill-considered legislation, and, more prosaically, disillusionment and squeezes in funding if and when technology fails to live up to misguided expectations. On these grounds, then, we suggest that it is in the interests of AI researchers to take particular care when invoking such rich psychological concepts.

(3) Forging a new psychology for AI

Finally, and most speculatively, we suggest that by overeagerly attributing familiar psychological notions to machine intelligences we may miss new and potentially illuminating information processing structures developed by AI researchers. To illustrate this point, compare current models of AI object recognition with their biological equivalents. The ability to spontaneously acquire new categories and apply them to inputs from external sensors is common to humans, animals, and many artificial systems.

Underlying this shared functional capacity, however, are key differences at the level of mechanism. For example, in order to reliably recognize a given object, artificial systems typically require a much larger set of training examples than humans or animals, and the kind of one-shot learning that is common in nature is extremely challenging for current computational architectures (Lake et al. 2016). Similarly, contemporary artificial approaches to object recognition are easy to fool via the use of adversarial examples, while no clear parallel exists for biological perception (however, see Elsayed et al. 2018 for some parallels). Artificial systems are also in many cases poor at extrapolating from natural images to cartoons or other pictorial forms of representation, a striking difference from the human case.

Yet there are also advantages of current machine systems over their biological counterparts, notably in speed and accuracy at classification among real world examples. Indeed, there is no reason some possible novel biological or artificial system to which we would comfortably ascribe general intelligence should have perceptual systems more like those found in animals and humans than those in current machine vision systems.

In light of this, we should take seriously the possibility that *new* rich psychological terms might be better suited to application in artificial systems. While it might be tempting to think of our repertoire of everyday psychological concepts as static, even in the human case, science has regularly led us to innovate in this regard: notions like implicit bias, associative learning, and motivated reasoning have found their way into everyday explanations of behavior, but were initially coined by psychological researchers. Similar theoretical innovations might allow us to develop a set of new high-level psychological concepts able to capture the specific functional dynamics of artificial systems and allow for useful generalizations about their role in systems.

As a toy example, we might use the aforementioned functional differences in machine vision to distinguish between the *perceptual* capacities of biological systems and the ‘sensory-

recognition' capacities of machines, without dismissing the latter merely as an impoverished species of the former. By thus expanding psychology to recognize these new kinds of processing, we might even in turn be better able to categorise and understand the cognitive systems of biological creatures quite different from ourselves, perhaps including phylogenetically remote animals such as eusocial insects or cephalopods (Godfrey-Smith 2017).

Conclusion

In summary, we claim that care should be taken before we employ the kind of psychological vocabulary currently applied to humans and animals to artificial systems, on the grounds that it may (i) impair scientific communication and understanding, (ii) invite premature conclusions of ethical or legal significance, and (iii) lead us to miss opportunities to expand our understanding of the varieties of minds.

It is commonly noted that previous 'AI winters' have been produced not so much by scientific stagnation as mismatches in expectations between researchers and broader academic and social communities. Ensuring such expectations are realistic by, among other things, reining in current tendencies to invoke rich psychological concepts to describe artificial systems will, we believe, help foster a more sustainable dialogue between researchers, funding bodies, and society at large. This is particularly important given that recent scientific developments have rendered it conceivable that artificial systems may soon literally possess some of the rich psychological capacities currently ascribed to them in only a narrow technical sense, and many readers—including informed ones—may struggle to determine which of these senses is being employed on a given occasion.

Photo caption: A male Eurasian jay shares food with a female partner according to her specific satiety. The attribution of desire-states may be an evolutionary and developmental precursor to Theory of Mind (Ostojić et al. 2013). Photo provided by Ljerka Ostojić.

References

- Abate, Tom. 2018. "An Artificial Nerve System Gives Prosthetic Devices and Robots a Sense of Touch." *Stanford News*. May 31, 2018. <https://news.stanford.edu/2018/05/31/artificial-nerve-system-gives-prosthetic-devices-robots-sense-touch/>.
- Call, Josep, and Michael Tomasello. 2008. "Does the Chimpanzee Have a Theory of Mind? 30 Years Later." *Trends in Cognitive Sciences* 12 (5): 187–92. <https://doi.org/10.1016/j.tics.2008.02.010>.
- Elsayed, Gamaleldin F., Shreya Shankar, Brian Cheung, Nicolas Papernot, Alex Kurakin, Ian Goodfellow, and Jascha Sohl-Dickstein. 2018. "Adversarial Examples That Fool Both Computer Vision and Time-Limited Humans." *ArXiv:1802.08195 [Cs, q-Bio, Stat]*, February. <http://arxiv.org/abs/1802.08195>.
- Godfrey-Smith, Peter. 2017. *Other Minds: The Octopus and the Evolution of Intelligent Life*. HarperCollins UK.
- Görür, O. C., Benjamin S. Rosman, G. Hoffman, and S. Albayrak. 2017. "Toward Integrating Theory of Mind into Adaptive Decision-Making of Social Robots to Understand Human Intention," March. <https://researchspace.csir.co.za/dspace/handle/10204/9653>.
- Halina, Marta. 2015. "There Is No Special Problem of Mindreading in Nonhuman Animals." *Philosophy of Science* 82 (3): 473–490.
- Hare, Richard M. 1952. *The Language of Morals*. OUP Oxford.
- Hassabis, Demis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. 2017. "Neuroscience-Inspired Artificial Intelligence." *Neuron* 95 (2): 245–258.
- Heyes, Cecilia. 2017. "Apes Submentalise." *Trends in Cognitive Sciences* 21 (1): 1–2.
- Karg, Katja, Martin Schmelz, Josep Call, and Michael Tomasello. 2015. "The Goggles Experiment: Can Chimpanzees Use Self-Experience to Infer What a Competitor Can See?" *Animal Behaviour* 105: 211–221.
- Kim, Yeongin, Alex Chortos, Wentao Xu, Yuxin Liu, Jin Young Oh, Donghee Son, Jiheong Kang, et al. 2018. "A Bioinspired Flexible Organic Artificial Afferent Nerve." *Science* 360 (6392): 998–1003. <https://doi.org/10.1126/science.aao0098>.
- Krachun, Carla, Robert Lurz, Jamie L. Russell, and William D. Hopkins. 2016. "Smoke and Mirrors: Testing the Scope of Chimpanzees' Appearance-Reality Understanding." *Cognition* 150 (May): 53–67. <https://doi.org/10.1016/j.cognition.2016.01.012>.
- Krupenye, Christopher, Fumihiko Kano, Satoshi Hirata, Josep Call, and Michael Tomasello. 2016. "Great Apes Anticipate That Other Individuals Will Act According to False Beliefs." *Science* 354 (6308): 110–14. <https://doi.org/10.1126/science.aaf8110>.
- Kulkarni, Tejas D., Karthik R. Narasimhan, Ardavan Saeedi, and Joshua B. Tenenbaum. 2016. "Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation." *ArXiv:1604.06057 [Cs, Stat]*, April. <http://arxiv.org/abs/1604.06057>.
- Lake, Brenden M., Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman. 2016. "Building Machines That Learn and Think Like People." *ArXiv:1604.00289 [Cs, Stat]*, April. <http://arxiv.org/abs/1604.00289>.
- Lurz, Robert W. 2011. *Mindreading Animals: The Debate over What Animals Know about Other Minds*. MIT press.
- Ostojić, Ljerka, Rachael C Shaw, Lucy G Cheke, and Nicola S Clayton. 2013. "Evidence Suggesting That Desire-State Attribution May Govern Food Sharing in Eurasian Jays." *Proceedings of the National Academy of Sciences* 110 (10): 4123–4128.

- Premack, David, and Guy Woodruff. 1978. "Does the Chimpanzee Have a Theory of Mind?" *Behavioral and Brain Sciences* 1 (4): 515–26. <https://doi.org/10.1017/S0140525X00076512>.
- Rabinowitz, Neil C., Frank Perbet, H. Francis Song, Chiyuan Zhang, S. M. Ali Eslami, and Matthew Botvinick. 2018. "Machine Theory of Mind," February. <https://arxiv.org/abs/1802.07740>.
- Schaafsma, Sara M., Donald W. Pfaff, Robert P. Spunt, and Ralph Adolphs. 2015. "Deconstructing and Reconstructing Theory of Mind." *Trends in Cognitive Sciences* 19 (2): 65–72. <https://doi.org/10.1016/j.tics.2014.11.007>.
- Selbst, Andrew D., and Solon Barocas. 2018. "The Intuitive Appeal of Explainable Machines." SSRN Scholarly Paper ID 3126971. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=3126971>.
- Shylaja, K. R., M. V. Vijayakumar, E. Vani Prasad, and Darryl N. Davis. 2017. "Artificial Minds with Consciousness and Common Sense Aspects." *International Journal of Agent Technologies and Systems (IJATS)* 9 (1): 20–42. <https://doi.org/10.4018/IJATS.2017010102>.
- Simon, Herbert A. 1980. "Cognitive Science: The Newest Science of the Artificial." *Cognitive Science* 4 (1): 33–46. [https://doi.org/10.1016/S0364-0213\(81\)80003-1](https://doi.org/10.1016/S0364-0213(81)80003-1).
- Weber, Théophane, Sébastien Racanière, David P. Reichert, Lars Buesing, Arthur Guez, Danilo Jimenez Rezende, Adria Puigdomènech Badia, et al. 2017. "Imagination-Augmented Agents for Deep Reinforcement Learning." *ArXiv:1707.06203 [Cs, Stat]*, July. <http://arxiv.org/abs/1707.06203>.
- Winfield, Alan F. T. 2018. "Experiments in Artificial Theory of Mind: From Safety to Story-Telling." *Frontiers in Robotics and AI* 5. <https://doi.org/10.3389/frobt.2018.00075>.