

Towards a common philosophy of explanation for artificial and biological intelligence

Jessica A Thompson (j.thompson@umontreal.ca)

International Laboratory for Brain, Music and Sound (BRAMS)

Montreal Institute for Learning Algorithms (MILA)

University of Montreal

Montreal, Quebec, Canada

Abstract

Much of the confusion that occurs when working at the intersection of cognitive science, artificial intelligence, and neuroscience stems from disagreement about what it means to understand intelligence. I claim that to integrate these fields, we must reconcile their different theories of explanation. After a brief review of the philosophy of explanation, I recontextualize the stated views of several prominent cognitive computational neuroscientists in terms of the theories of explanation they espouse. Finally, I describe some of the challenges of forging a new theory of explanation that would apply equally to artificial and biological intelligence. As a first step towards an integration of research on biological and artificial intelligence, my goal in writing this paper is to equip scientists of intelligence to interrogate and justify the theories of explanation that underlie their definitions of scientific progress.

Keywords: philosophy of science; ; explanation; neuroscience; deep learning

Integration of AI and neuroscience

Much lip service is paid to the "integration of deep learning and neuroscience", with the goal of creating a feedback loop deep learning for neuroscience and neuroscience for deep learning. This turns out to be hard to do in practice. The confusion when trying to work at this intersection comes in large part from lack of agreement about what progress towards a common goal would look like. This topic came up at the inaugural Cognitive Computational Neuroscience (CCN) conference last year, which was assembled to unify the "disconnected communities of cognitive science, artificial intelligence, and neuroscience" towards the common goal of "understanding the computational principles that underlie complex behavior" (Naselaris et al., 2018). Jim DiCarlo, chairing a panel discussion, asked, "when people say they want to work together, usually there is some idea of a shared goal...some idea of what success would even look like...are we even after the same thing?" This question received a number of very different answers from the panel, demonstrating the challenge of even agreeing on a common goal. Panelist Yann LeCun stated the common goal to "explain intelligence" but this doesn't answer the question because we disagree about what is important to explain intelligence. LeCun wants to replicate animal intelligence in artificial systems. On the other hand, for neuroscientist Jackie

Gottlieb, "success means characterizing a system at a particular level of abstraction ... in a way that is reproducible and solid" (Kay, 2017). Cognitive scientist Josh Tenenbaum, stressed the importance of distinguishing between goals on different time scales and suggested that all the CCN attendees probably share some long term vision of success, even if they disagree about what to do to work towards that goal in the short term. I think that coming up with good answers to these questions is the most important hurdle to overcome right now. An integration of cognitive science, artificial intelligence and neuroscience will not be possible until we are able to motivate our research by reference to a shared definition of what it means to make progress towards the goal of understanding intelligence.

The same debate is happening in machine learning right now. The quest for "interpretable" AI is ultimately asking what explanations of AI systems will we accept? Are some systems more explainable than others? For example, are systems that are designed specifically to expose 'disentangled' representations more interpretable? Several events were dedicated to related topics at the Neural Information Processing Systems conference in 2017 (e.g. Interpretable ML Symposium, Learning Disentangled Representations: from Perception to Control). I think it is no coincidence that machine learning and neuroscience are both having these conversations now. Rather, it's precisely because artificial systems are looking more like biological ones and our models of biological intelligence are looking increasingly like AI that we are forced to question our standard conceptions of what makes a good explanation.

These questions are ultimately in the realm of epistemology and philosophy of science, yet philosophical theories are not often (explicitly) invoked in the discussion. I've realized that philosophers actually have a lot to say about computational explanation in neuroscience and cognitive science. As scientists, instead of reinventing the wheel, we would do well to look to our philosopher colleagues to help us wade through these difficult but crucially important questions about what constitutes an explanation. At the very least, I think our discussion would be simplified if we borrowed from the established language of philosophy of explanation. But I will make a stronger claim that what is needed is actually to create a new theory of explanation that applies equally to biological intelligence and artificial intelligence. The way we understand AI systems (the methods we use to study them, the nature of explanations that we accept) are very different from the way we traditionally un-

derstand and study biological systems. At present, this constitutes a challenge to the CCN goal, but in the long term, I see this as an opportunity to define a new science of intelligence that includes both artificial and biological intelligence. My central claim is that to achieve an integration of cognitive science, AI and neuroscience, we must reconcile their different theories of explanation. Background in the philosophy of scientific explanation of neural and computational systems will better equip us to express our views on questions like 'how ought intelligence be explained?', and ultimately to be able to design experiments that we can rigorously justify.

Theories of Explanation

To begin, we assume that the goal of science is to know and understand our world. A primary goal of science is to provide explanations of phenomena (be they natural, social, or otherwise). The role of a theory of scientific explanation is to characterize the structure of explanations in science. An account of scientific explanation must distinguish between explanations that are scientific and those that are not. It must also distinguish between explanations and non-explanations. Sometimes this second contrast is presented as the difference between explanation and description. For example, a set of claims about the appearance of a particular species may be true, accurate and supported by evidence without being explanatory in any way. They are "merely" descriptive (Woodward, 2017).

An important but neglected piece of scientific literacy is the ability to look at a set of claims or a model and decide whether it is descriptive or explanatory (or neither) according to your favorite theory of scientific explanation. At the same time, I don't want to suggest that description is somehow inherently less valuable than explanation. We must know that a phenomenon exists before we can ever hope to explain it. In calling for better literacy of theories of explanation, my wish is not to reduce the amount of descriptive science, but rather to reduce the misrepresentation of scientific activities. If we already think that we are explaining, then we will not spend time figuring out how to explain. Instead, if we acknowledge that what we are currently doing is descriptive, we can better see our role in a larger scientific enterprise, i.e. the role of a particular series of experiments may be to describe a specific set of phenomena that may later be explained by other experiments. When our perceived explanatory power is inflated, we're closed off from seeing how we might work together to ultimately better explain in the long term.

What do contemporary scientists say?

In this section I will discuss how some computational neuroscientists have recently answered the question "what makes a good model (or theory) of the brain?" In their answers, we will find their philosophical commitments.

In his talk *Playing Newton: Automatic Construction of Phenomenological, Data-Driven Theories and Models*, Ilya Nemenman refutes the claim that a good theory of brain must

be a large, multi-scale computational model with a quote from Wiener and Rosenblueth, "Theories must lose details and must be developed to explain limited sets of phenomena. Otherwise, the best material mode of a cat is another, or preferably the same, cat" (Philosophy of Science, 1945). In other words "Don't model bulldozers with quarks" (Goldenfeld and Kadanoff, 1999). A good theory is one that accurately explains a limited set of phenomena and throws away everything it wasn't designed for (Nemenman, 2018). But what does it mean to *explain* to Nemenman? He claims that philosophy has failed to define science, but maybe we don't need it anyway because we have Bayesian statistics. Bayesian statistics can already tell us what is falsifiable and what is falsified. If theory does not explain observed data then the theory is falsified. If theory can explain any dataset, then it is unfalsifiable. What makes a good model of the brain is determined by Bayesian model selection. The goodness of a theory is related to generalization and prediction, with little regard to whether it is true. Falsifiability is then real-valued and empirical. If it explains the data, we don't care if it is correct or not. Good models:

- are phenomenological
- predict your data and generalize to new data and experiments
- Only explain the specific question they were designed to answer with as few parameters as possible

This view reflects the well known heuristic of Occam's Razor: the best solution is the one that explains the data best with the simplest model.

Nemenman advocates for phenomenological models, which have been characterized by philosophers of science as "black box models" because they merely capture input-output relationship for the phenomena to be explained without positing intervening variables or mechanisms. In computational neuroscience, phenomenological models are called descriptive models because they "summarize data compactly" without addressing "the question of how nervous systems operate on the basis of known anatomy, physiology, and circuitry" (Dayan & Abbott, 2005). Nemenman stresses that the important test is whether the model generalizes to unseen data and new situations (Nemenman, 2018). Consider the Balmer formula as a counter example, which was constructed to describe the four visible lines in the emission spectrum of hydrogen. Not only does this model accurately capture the target phenomena, but it also successfully predicted the existence of previously unobserved spectral lines outside of the visual range. Despite its success, none of the model elements have any physical interpretation. Balmer arrived at this formula via trial and error curve fitting to find the best mathematical fit to the four visible spectral lines of hydrogen. It is generally agreed that the Balmer formula is not explanatory because it does not explain why the emission spectrum for hydrogen shows the pattern that it does (Kaplan, 2011). Phenomenological models can

certainly be very useful and are an important part of any scientific enterprise, but to claim that they are explanatory is a difficult claim to defend.

In *Principles for models of neural information processing* (2017), Kendrick Kay defines cognitive neuroscience to be the quest for explanations of the mind in terms of the brain. Yet, in his definition of explanation, he says "models posit that specific variables relate to neural activity. As such, models provide explanations of measurements of the brain." Is the phenomena to be explained by the model the behaviour/perception of the animal or the neural activity itself? The implication seems to be that if we can explain the activity, this will imply an explanation for the behaviour or perception. Consider the example given in the paper, why is a neuron highly active during a clip of rock music? We can compare different models (the neuron responds to sound intensity vs the neuron responds to guitar tones) to decide which model is more accurate. For example, we can vary the intensity of a variety of sounds and find that the activity of this neuron scales with the sound intensity, while not being selective to any other tested features. Have we then explained the activity of this neuron and at the same time, explained perception of sound intensity? Some philosophers claim that explanations of human brains proceed via this type of *functional analysis*, "according to the explanatory strategy of functional analysis, the overall behavioral capacities [are] explained by breaking down or decomposing the capacities into a number of "simpler" subcapacities and their functional organization" (Kaplan, 2011). This theory of explanation goes hand in hand with the information processing theory of mind. When you assume that the brain is best understood as an information processing machine, then once you've discovered the signals that it sends around and the operations that it performs, then you should be able to explain the resulting behaviour. Under a different theory of mind (e.g. embodied dynamicism), or a different theory of explanation (e.g. mechanistic), functional analysis is not explanatory. However, many mechanists would say that functional analysis is a necessary step en route to a mechanistic explanation.

Jonas Kubilius, also asks the question 'what does it mean to understand?' in his commentary *Predict, then simplify*. To Kubilius, predictive power is the first and foremost attribute on which to assess a model. He and the rest of Jim DiCarlo's lab appear to be staunch predictivists, taking an engineering approach to neuroscience. This has led their group to find that deep neural network models trained originally for visual object recognition, better predict activity in much of visual cortex better than previous 'scientist designed' models (Yamins & DiCarlo, 2016). Predictivism, refers to the view that phenomenological models are explanatory by virtue of their descriptive and predictive power. The problems facing predictivism are well known in philosophy of explanation. "Simple examples readily expose how accurate prediction is insufficient for explanation, and so how the predictive import of a given model can and often does vary independently of its ex-

planatory force. One can accurately predict a storm's occurrence from a dropping barometer, but this does not explain the occurrence of the storm. Rather, a common cause—a drop in atmospheric pressure—explains both the dropping barometer value and the approaching storm. Similarly, a p-model might be predictively adequate, and yet its variables might only represent factors that are mere behavioural correlates of some common or joint cause for the target phenomenon. Just as we reject the claim that p-models of this kind provide explanations, we should also resist the claim that p-models of this kind provide explanations" (Kaplan, 2011, p.351). A defense of predictivism must ultimately address these types of counter examples.

On first read of Kay and Kubilius, one might think that their disagreement about DNNs being explanatory models of the brain comes from placing different values on the predictive power of a model. Kubilius says predictive power comes first and foremost. Where as Kay says that the model must first be "understood" before it can explain. Yet both rely largely on prediction of observational data to validate their models. The larger disagreement is about the form of a "good" model. Both Nemenman and Kay ascribe to the heuristic of Occam's Razor, that the simplest model that explains the data is best. They want the number of components to be small because ultimately their modelling efforts are about finding which components are important. Kubilius points out that although historically, science's greatest successes have been the result of describing complex systems with relatively few parameters, there is no reason to assume that this will be the case for understanding the brain. Others have pointed out the absurdity of Occam's Razor, "how could a fixed bias toward simplicity indicate the possibly complex truth any better than a broken thermometer that always reads zero can indicate the temperature? You don't have to be a card-carrying skeptic to wonder what the tacit connection between simplicity and truthfinding could possibly be" (Kelly, 2007). This adherence to Occam's razor part of what separates Nemenman and Kubilius. Simpler things seem to be easier for us to understand, so the tendency of science to look for simplicity when it is available makes sense. Kubilius suggests that we may look for simple generative processes that generate complex systems. For example, training deep neural networks depends on a number of principles (e.g. optimization of cost functions via gradient-based methods, compositionality, distributed representations) that is likely smaller than the number of parameters in the model.

Towards a new theory of explanation for a new science of intelligence

There is a growing body of work trying to understand deep learning using empirical methods that look almost neuroscientific: ablation analyses, receptive field analysis, psychophysics. Nikolas Kriegeskorte calls it "synthetic neurophysiology" (Kriegeskorte, 2015), Jeff Clune calls it "artificial neuroscience" (Metz, 2018), and Maithra Raghu calls

is "deep learning science" (personal communication, Dec 7 2017), presumably because it involves applying the scientific method to understanding artificial systems. What I find most exciting about this line of inquiry is that it makes many connections to deep learning theory. If we can relate deep learning theory to empirical analyses of deep learning systems, then might there also be the potential to relate our empirical analyses of biological intelligence to a similar mathematical understanding? On the other hand, Eric Jonas and Konrad Kording asked *Could a neuroscientist understand a microprocessor?* in their work which applied common neuroscientific analyses to electrical measurements of a microprocessor. They show that although these analyses were able to make replicable/reliable descriptions of patterns of activity, they did not ultimately reveal the known and accepted explanation of how the microprocessor works (Jonas & Kording, 2016). Swapping methodologies, approaches and philosophies between deep learning and neuroscience has the potential to demonstrate the strengths and limitations of our scientific activities and perspectives, helping us to select those that will be most useful towards our common goal of understanding intelligence.

These approaches are still pre-theoretical in the sense that they do not fit well into existing theories of explanation. All contemporary theories of explanation in neuroscience focus on physical computation. The mechanistic account, which currently dominates philosophy of neuroscience, requires that explanations consist of physical components in causal relationships with one another. The explanations we have for deep learning are concerned with abstract computation. The same principles hold regardless of which type of GPU your model is trained on. Reconciliation of these views suggests a commitment to multiple realizability. However multiple realizability is often associated with functionalism and computational chauvinism. We want a version of multiple realizability for the 21st century aligned with a modern theory of mind.

Conclusion

If our ultimate goal is truly to explain intelligence, then we must eventually agree on a theory of explanation that accounts for successful explanations of both artificial and biological intelligence. Although the mechanistic framework is currently very popular in most of philosophy of science, it is difficult to apply to the explanation of cognitive phenomena. Hence, much of cognitive and computational neuroscience that claims to be explanatory relies on older problematic theories of explanation and outdated theories of mind. To forge a new path forward, we need to first acknowledge that we fundamentally do not know how to explain intelligence. Until we do, let us be explicit about our philosophical commitments and let us value our descriptions without mistaking them for explanations.

References

Dayan, P., & Abbott, L. (2005). *Theoretical neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press.

- Jonas, E., & Kording, K. P. (2016). Could a neuroscientist understand a microprocessor? *bioRxiv*, 055624. doi: 10.1101/055624
- Kaplan, D. M. (2011). Explanation and description in computational neuroscience. *Synthese*, 183(3). doi: 10.1007/s1229-001-9970-0
- Kay, K. N. (2017). Principles for models of neural information processing. *NeuroImage*(March), 1–9. doi: 10.1016/j.neuroimage.2017.08.016
- Kelly, K. T. (2007). Simplicity, Truth, and the Unending Game of Science. In S. Bold, Benedikt Lowe, T. Rasch, & J. van Benthem (Eds.), *Foundations of the formal sciences v: Infinite games*.
- Kriegeskorte, N. (2015). Deep neural networks: a new framework for modelling biological vision and brain information processing. *Annual Review of Vision Science*, 1, 417–446. doi: 10.1101/029876
- Metz, C. (2018, mar). *Google Researchers Are Learning How Machines Learn*. <https://www.nytimes.com/2018/03/06/technology/google-artificial-intelligence.html>.
- Naselaris, T., Bassett, D. S., Fletcher, A. K., Kording, K., Kriegeskorte, N., Nienborg, H., ... Kay, K. N. (2018). Cognitive Computational Neuroscience: A New Conference for an Emerging Discipline. *Trends in Cognitive Sciences*, 22(5), 365 – 367. doi: 10.1016/j.tics.2018.02.008
- Nemenman, I. (2018). *Playing Newton: Automatic Construction of Phenomenological, Data-Driven Theories and Models*. <https://simons.berkeley.edu/talks/ilya-nemenman-4-17-18>.
- Woodward, J. (2017). Scientific Explanation. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (<https://pl.ed>). Metaphysics Research Lab, Stanford University.
- Yamins, D. L. K., & DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3), 356–365. doi: 10.1038/nn.4244