PhiMiSci
Philosophy and the Mind Sciences | Vol. 7 | 2026

Do all brains have neurons?

Metaphysics and neutrality in *The Brain Abstracted*

Alexandros Constantinou[1]
John Donaldson[1]
Will Hornett[2]
Andrea Blomkvist
Derek H. Brown[1]
Jack C. Lyons[1]

Book symposium on *The Brain Abstracted* by Mazviita Chirimuuta, edited by Philipp Haueis

[1] University of Glasgow [2] University of Cambridge
* Primary contact: andrea.blomkvist@glasgow.ac.uk

Constantinou, A., Donaldson, J., Hornett, W., Blomkvist, A., Brown, D. H., & Lyons, J. C. (2026). Do all brains have neurons? Metaphysics and neutrality in *The Brain Abstracted*. *Philosophy and the Mind Sciences*, *7*.
https://doi.org/10.33735/phimisci.2026.12279

Abstract

Mazviita Chirimuuta's *The Brain Abstracted* (2024) is a fascinating intervention into the philosophy of mind and neuroscience, containing deeply interesting ideas and arguments. Our aim is to critically probe whether Haptic Realism is neutral on some substantive issues which Chirimuuta would like it to be neutral on. Firstly, it is unclear whether Haptic Realism is compatible with Chirimuuta's metaphysical neutrality. Causal notions feature heavily in Haptic Realism, including construction and interaction, without which it is unclear what the haptic component of Haptic Realism amounts to. We argue that the interpretation of these causal notions forces a choice that jeopardises metaphysical neutrality, since it is doubtful that there are any relevantly neutral theories of causation. Secondly, we ask the question 'Do all brains have neurons?' Haptic Realism trades on the idea of ideal patterns, patterns that are coaxed into existence by the interaction between practitioners and the world. If ideal patterns only exist within the confines of the lab, and neurons are idealizations, then brains out in the wild do not have neurons. Worryingly, this commitment of Haptic Realism undermines our ability to project neurological knowledge that we gain from studying particular brains to what is happening inside the everyday human beings.

Keywords: Causation • Haptic realism • Metaphysical neutrality • Perspectivism • Philosophy of mind • Projectability • Scientific realism

Mazviita Chirimuuta's *The Brain Abstracted* (2024) is a fascinating intervention into the philosophy of mind and neuroscience, containing deeply interesting ideas and arguments. Our aim is to critically probe the relation between the book's philosophy of science, epistemology, and metaphysics. Our central question is whether Haptic Realism is neutral on some substantive issues which Chirimuuta would like it to be neutral on.

As a variety of perspectivism, Haptic Realism is advertised as an alternative to traditional realism and anti-realism: there is scientific knowledge, but it is not knowledge of mind-independent reality, as such.[1] Haptic Realism centres on the idea of an ideal pattern: representational idealisations created by the "process of interaction between scientist, lab equipment, target system, data, and modelling" (p. 131). Ideal patterns are the objects of study and knowledge (pp. 45-46) which practitioners construct. They are, by their nature, creations of the scientist, not things found there already in the natural world. Haptic Realism includes (i) metaphysical, (ii) semantic, and (iii) epistemic components:

> (i) Ideal patterns are inherently more simple than reality, which is intractably complex. Thus, ideal patterns are not homomorphic with reality (Chirimuuta, 2023: 9-10; 2024: 36-45).
>
> (ii) Ideal patterns are constructed analogies of reality. Thus, the language expressing them is at least largely figurative (Chirimuuta, 2023: 10-11; 2024: 52-56).
>
> (iii) Ideal patterns give knowledge only of the human-relative properties of reality via our interactions with it. Thus, they do not deliver knowledge of mind-independent reality (Chirimuuta, 2023: 12-13; 2024: 46-52).

It is important to note two things about Haptic Realism. Firstly, it is a kind of transcendental idealism about scientific knowledge (p. 40). According to it, we can never know about reality as it is in itself, but only as it is presented to us given our interactions with it. This kind of transcendental idealism is not a big metaphysical thesis, but rather a call for epistemic humility about metaphysical claims (Langton, 1998). This is why Chirimuuta feels able to take no stance on, for example, "the question of whether nature fundamentally is or is not a densely connected 'causal nexus' [...] which is to say a 'causal structure' revealed in true causal explanations" (2024: 169, 178). Secondly, the way we construct ideal patterns is not by passively mirroring the world, but by *prodding*, *poking*, and *processing* it. These interactions are essential for making, upholding, and moulding our representations.

This is sufficient to frame our two concerns.

The first is the compatibility of Haptic Realism with Chirimuuta's metaphysical neutrality. Notice that (i-iii) include causal notions, particularly of *construction* and *interaction*. Clearly, much hangs on how to interpret such notions, since otherwise it is unclear what the *haptic* component of Haptic Realism amounts to. At one point, Chirimuuta endorses the Epistemic Theory of Causation (pp. 59-60), citing Williamson (2004), Falkenburg (2012), and Taylor (2021). On Williamson's version, this is the view that *A causes B* iff believing that A causes B is sufficiently inferentially useful.[2] So, there are objective facts about causality, but those facts depend on our epistemic framework. This provides an interpretation of the causal notions: a scientist constructs a model, measures a neuron, or anaesthetises a cat (and so on) iff believing that she does is sufficiently inferentially useful.

But now it is unclear how Haptic Realism, the Epistemic Theory of Causation, and metaphysical neutrality hang together. On its face, the Epistemic Theory is not metaphysically neutral, since it denies that there are mind-independent causal relations. And if we interpret Haptic Realism's causal notions in terms of the Epistemic Theory, then Haptic Realism isn't neutral either. This jeopardises Haptic Realism's advertised Kantian humility about metaphysical claims. So a difficult choice is forced. On the one hand, if Chirimuuta prioritises commitment to neutrality, then she must

[1] In the special sciences, at least. Chirimuuta is less sure about some parts of physics (2023: 6n5).

[2] Williamson claims that, for *A* to cause *B*, it is necessary that *believing that A causes B is sufficiently inferentially useful*. Williamson's view has undergone development (Williamson, 2013; 2021), but is explicit that in possible worlds where causal beliefs are not inferentially useful, there is no causation (Williamson, 2024). Taylor follows Williamson, but attempts to be less ontologically committal (2021: 9150n15).

jettison the Epistemic Theory in favour of some alternative interpretation of Haptic Realism's causal notions that is consistent with its aspiration to neutrality. However, it is not clear whether any theory of causation likely to shed much light on Chirimuuta's causal talk is completely metaphysically neutral, since it is doubtful that there are any neutral theories of causation. This means it is not clear whether there is any neutral interpretation of the causal notions. On the other hand, if Chirimuuta prioritises the Epistemic Theory, then the epistemic humility of neutrality is the cost. Either way, Chirimutta's promise of an alternative to traditional realism and anti-realism is potentially at stake.

Our second worry concerns neutrality in a different way. Put provocatively, our question is this: 'Do *all* brains have neurons?'

Let's start with ideal patterns again. Chirimuuta argues that *real* patterns are inaccessible because of their incomprehensible complexity. Hence, practitioners coax ideal patterns into existence so that they can investigate an accessibly simplified representation of the world, inextricably bound to the practitioner and their discipline-dependent modes of investigation.

One such ideal pattern is the model of the simple cell, a type of cell found in the early visual cortex (V1) "whose activity reflects an essentially linear sum of light falling within its receptive field" (p. 124). Chirimuuta argues that "the simple cell is an ideal pattern", a "creation" which "would not come about" but for the adaptiveness of neurons to the context of quite particular "experimental constraints" (p. 131). Chirimuuta seems to be saying the following. When you study brains in a certain way, you create a thing called 'the simple cell', which is an idealised pattern of neural behaviour that only exists because of, and in, that context. The simple cell arises as a result of all the features of your interaction with the brain and its neurons.

Yet this is puzzling. It initially implies that there *are* simple cells, but only because of, and in the context of, these experimental situations. Put bluntly, studied brains have simple cells, but unstudied brains do not. If this is true, then the ability of neuroscience to project its generalisations onto unstudied samples, a key component of all science, is undermined. If most brains don't have simple cells (because most brains have not been studied), then how can we apply knowledge of them to such brains? For projection, one needs the unstudied samples to resemble the studied samples in important ways. But this seemingly blocks projective generalisations: unstudied brains have no entities onto which generalisations can latch because simple cells are brought into existence by the process of studying the studied brains, thereby leaving projectability apparently impossible.

Here is a response to the initial puzzle. Chirimuuta says that "the simple cell is an ideal pattern" (p. 131), but that "there were cells in the primary visual cortex already before the experiment" (p. 127). She seems to think the claim that there are neurons in unstudied brains is unproblematic. So perhaps we should read her claim that "the simple cell is an ideal pattern" as the claim, not that particular neurons in studied brains are created by scientists, but rather as the claim that a certain *category* – the category of the simple cell – is created in experimental settings. And, in those settings, scientists discovered that some already-existing neurons fit that constructed category. This could permit projective generalisations. Assuming rough homomorphism between studied and unstudied brains, we can predict that, if an unstudied brain were studied in the ways required for some of its neurons to meet "the profile definitive of the simple cell" (p. 127), then they would (probably, *ceteris paribus*) meet that profile. That projection doesn't yield an invariant capacity of the cell across all contexts, but that's an empirical limitation of the theory of the simple cell (p. 137). At least this move permits Chirimuuta's view to licence *some* projections.

But why does Chirimuuta find it unproblematic that all of our heads are filled with neurons? The neuron, on this view, should clearly also be an ideal pattern; a category created through the studying of particular brains. The neuron is supposed to be "a device that gathers inputs at the dendrites, calculates a function, and delivers an output" (p. 99). It is supposed to be a biological logic-gate and is taken to be the brain's fundamental building block. But by Chirimuuta's own lights, the neuron "should not be taken to exist independently" of the scientific process by which an intelligible model is created (p. 46). Indeed, the criticisms she makes of the theory of the simple cell have something in common with her criticism of the theory that neurons are the building blocks of brains (pp. 99-101): both theories

rely on over-simplification, abstraction, and experimental constraints.

If the neuron is a constructed ideal pattern, we can ask what there is in unstudied brains that permits any projective generalisation from 'Studied brains with neurons do *X*' to 'Unstudied brains will (probably, *ceteris paribus*) do *X* too'? Chirimuuta cannot appeal to already-existing neurons, since this appeals to things which meet the profile definitive of neurons described by an ideal pattern that only exists because of, and in, the contexts in which that pattern is generated by experimental conditions. Now, Chirimuuta clearly thinks that the world is not homogeneous, structureless gunk which we project ideal patterns onto (p. 45). It must have some structure, or else it is hard to understand why different things respond to the same intervention differently. But it isn't clear what this can be, and whether her theory is consistent with any particular possible suggestions. So, do all brains have neurons, or is that just a tale neuroscientists tell us, a simplifying idealisation of the truth, that only some do? For if it is the latter, then once again we have a deep crisis of projectibility: we might understand, in some simplified way, why studied brains do what they do, but this wouldn't help us even a little bit in understanding why anyone who hasn't been studied displays any of the behaviour that they do. None of the neurological knowledge we have gained from studying particular brains can help us know what is happening inside the vast majority of human beings; there is no way to project from those with neurons to those without. And this crisis of projectibility undermines the idea that neuroscience is much of a science at all.

Acknowledgements

This work grew from a reading group on *The Brain Abstracted*. While all authors contributed to the overall process of producing this work, the order of authors as listed reflects the overall contributions to the final work, starting with the highest contribution. Constantinou and Donaldson are joint first authors and Hornett is sole second author.

Funding

This work was generously supported by multiple funders. Derek H. Brown and Jack C. Lyons were supported by a joint Arts and Humanities Research Council-Deutsche Forschungsgemeinschaft grant for the Sensory Engineering project (grant number: RF-2021-453; AHRC AH/Y007638/1). Jack Lyons was additionally supported by the Leverhulme Trust [grant number MRF-2023-191]. Andrea Blomkvist was supported by a British Academy Postdoctoral Fellowship (grant ref: PFSS23\230115) and Lord Kelvin/Adam Smith Fellowship from the University of Glasgow.

References

Chirimuuta, M. (2023). Haptic realism for neuroscience. *Synthese*, *202*(3), 1–16. https://doi.org/10.1007/s11229-023-04295-z

Chirimuuta, M. (2024). *The brain abstracted: Simplification in the history and philosophy of neuroscience*. The MIT Press.

Falkenburg, B. (2012). Wieviel erklärt uns die hirnforschung? mythos determinismus: Wieviel erklärt uns die hirnforschung?, 327–385.

Langton, R. (1998). *Kantian humility: Our ignorance of things in themselves*. Oxford University Press.

Taylor, S. D. (2021). Causation and cognition: An epistemic approach. *Synthese*, *199*(3), 9133–9160. https://doi.org/10.1007/s11229-021-03197-2

Williamson, J. (2004). Epistemic causality. In J. Williamson (Ed.), *Bayesian nets and causality: Philosophical and computational foundations* (pp. 130–151). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198530794.003.0009

Williamson, J. (2013). How can causal explanations explain? *Erkenntnis*, *78*, 257–275. https://doi.org/10.1007/s10670-013-9512-x

Williamson, J. (2021). Calibration for epistemic causality. *Erkenntnis*, *86*(4), 941–960. https://doi.org/10.1007/s10670-019-00139-w

Williamson, J. (2024). Bayesian nets and causality errata. https://jonwilliamson.uk/publications/BNAC-errata.html