https://doi.org/10.1038/s41583-025-00946-x

Top-down and bottom-up neuroscience: overcoming the clash of research cultures

Fernando E. Rosas, Andrea I. Luppi, Pedro A. M. Mediano, Morten L. Kringelbach, Luiz Pessoa & Federico Turkheimer

As scientists, we want solid answers, but we also want to answer questions that matter. Yet, the brain's complexity forces trade-offs between these desiderata, bringing about two distinct research approaches in neuroscience that we describe as 'top-down' and 'bottom-up'. Recognizing the validity of both approaches dispels misunderstandings and unnecessary tension and promotes constructive interactions.

A passion for neuroscience is often sparked from an early exposure to big-picture questions — for example, how brains give rise to minds, or what consciousness is. The approaches used to tackle such grand questions tend to coalesce into two distinct categories, corresponding to different neuroscientific 'research cultures' that diverge in not only research methods but also expectations on what it means to make progress.

Bottom-up neuroscience

The bottom-up research culture cares primarily about building knowledge from firm foundations and drives progress through systematic accumulation of solid evidence and increasing detail. Accordingly, bottom-up approaches emphasize careful experimental design for studying neuroscientific scenarios by keeping everything as controlled as possible. When this cannot be achieved, researchers are willing to revise the definition of their object of study until it becomes amenable to full experimental control — thereby prioritizing a less direct but safer path to address the 'big' questions by breaking them down into more manageable parts.

Examples of bottom-up research include the tightly controlled experiments in the Nobel prize-winning work of Eric Kandel, whose team operationalized memory as habituation to repeated prodding in a sea slug, *Aplysia californica*[1]. By accepting this — arguably drastic — simplification, Kandel's team were able to leverage the exquisite experimental accessibility of *Aplysia*, characterising synaptic long-term potentiation with unprecedented detail.

Top-down neuroscience

In contrast to the bottom-up culture, the top-down research culture is driven by the desire to work on 'big' questions directly, with less willingness to swap the object of study for more tractable surrogates. However, the price for directly addressing 'big' questions rather than more tractable substitutes is that top-down researchers often have to tolerate more uncertainty in their investigations and their outcomes, as experimental and analytical tools are often inadequate to fully capture the richness of big-picture questions. Thus, top-down approaches usually start from approximate attempts at an answer that are then progressively refined, with the details added later — rather than thoroughly answering a simplified question and then expanding it. Accordingly, top-down approaches sometimes follow hard-to-anticipate intuitions that go beyond the next logical step.

A paradigmatic example is physicist Ludwig Boltzmann's work: while their ground-breaking results were correct, Boltzmann's often-incomplete arguments sometimes appeared as 'logical jumps'[2]. Network neuroscience also exemplifies successful work based on top-down principles[3]: by emphasizing fundamental properties that the brain shares with other networked systems (for example, cities and social networks), network neuroscience has become a flourishing neuroscientific subfield, bridging scales and promoting translational discovery.

How the divergence can hinder progress in neuroscience

Because of the brain's inherent complexity, from the microscale to the macroscale[4], research programmes face unavoidable compromises and must set priorities. The difference between top-down and bottom-up neuroscience can thus be understood as stemming from diverging priorities — big-picture questions and building on solid foundations, respectively. Drawing on an analogy from statistics, bottom-up approaches prioritize precision and make progress by reducing bias from the target, one step at a time, while top-down approaches prioritize accuracy and make progress by 'narrowing down' to the target (Fig. 1).

It is not difficult to see how each approach can be criticized. Most neuroscientists might accept the usefulness of studying neural circuits as networks, but some might disagree with scaling this up to analyse statistical correlations between functional MRI activity from centimetre-sized brain regions. Are top-down approaches simply too vague, lacking in rigour and insufficiently grounded in observable biology? Conversely, it is uncontroversial that habituation is a type of learning and that its study enabled great strides in our understanding of synaptic function. Yet, some may argue that habituation in slugs remains far from the richness of internal experience we tend to associate with memories. Progress has been made, but is this really progress towards the original goal, or has the goal been shifted by the bottom-up mindset to make it attainable?

This dichotomy is reminiscent of the 20th-century schism between 'analytic' and 'continental' philosophy[5]. Analytic philosophers emphasized breaking down complex issues through logical analysis to enable rigour and precision, at the risk of becoming bogged down in overly narrow questions. Continental philosophers, by contrast, devoted themselves to arguably 'existential' questions, at the risk of failing to tame the complexity of such big-picture questions. Although not as severe, the tension between top-down and bottom-up mindsets in neuroscience can become obvious in academic gatherings or during peer review, with many sterile debates arising from a misunderstanding

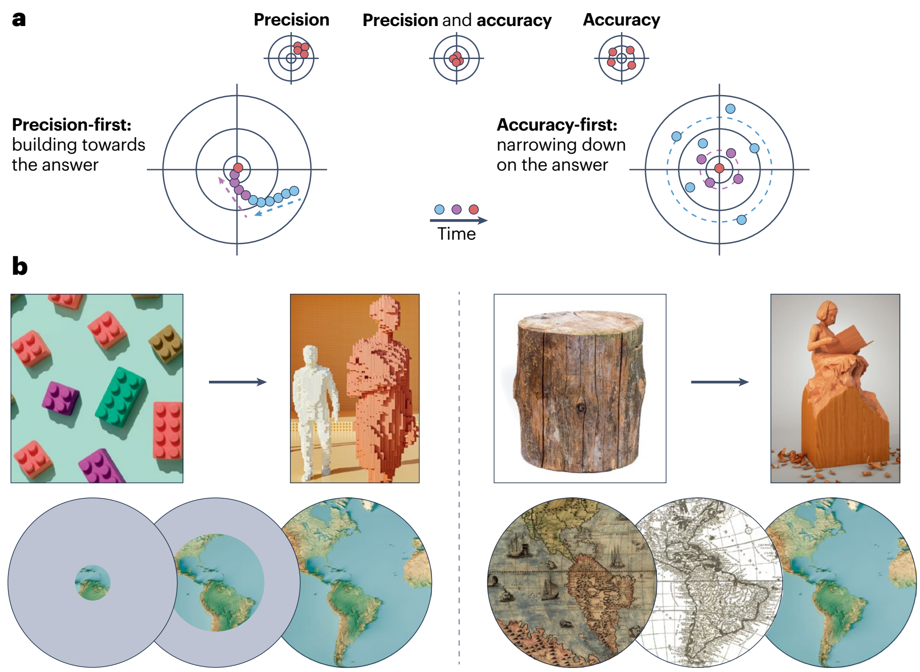

Fig. 1 | Illustrating the divergence between top-down and bottom-up approaches. a, Precision and accuracy (or bias and variance) from statistics. Precision means that hits are closely clustered. Accuracy means that hits are centred around the true target. For bottom-up neuroscience, 'hits' will be close to each other at every step of the way (high precision), but the starting point may be far from the ultimate target (low initial accuracy). Top-down neuroscience goes for the desired target right away, but its tools may initially be inadequate, resulting in wide dispersion (low precision) but more focus around the target itself (higher accuracy). **b**, Analogies from art and cartography. In art (top), bottom-up approaches build from the parts to the whole; top-down strategies start with a barely specified whole and develop it until the figure within becomes apparent. In cartography, bottom-up approaches can be seen as building maps by carefully charting one patch of land at a time, moving on to the next only when one is complete; top-down approaches resemble the ancient cartographers, providing a large-scale picture that is initially coarse, with the details filled in later. Credits: colourful toy blocks on green: Javier Zayas Photography/Moment/Getty. Generic business people made out of toy blocks: gremlin/E+/Getty. World map 3D render topographic map colour: FrankRamspott/E+/Getty. Childhood girl sculpture: timandtim/DigitalVision/Getty. Wood log isolated on a white background: R.Tsubin/Moment/Getty. Old map: katatonia82/iStock/Getty Images Plus. Eighteenth-century white map of America: comabu/iStock/Getty Images Plus.

of each other's priorities and motivations. Proponents of each culture mistakenly believe that others are operating under their own same conception of progress, but improperly so. Such a situation is analogous to judging a statistical model purely on the basis of precision or accuracy (Fig. 1). However, the well-known bias–variance trade-off from statistics states that precision and accuracy represent complementary – rather than competing – evaluation criteria for a model's quality. Different contexts (for example, the cost of making observations; whether false positives or false negatives are more costly) shape the precision–accuracy trade-off. We need to embrace this way of thinking ourselves, not only when modelling neuroscientific data but also in how we evaluate our own and others' neuroscientific work.

Consilience and synthesis: concrete actions for progress

Ultimately, science is about rational, evidence-based resolution of disagreement, and this process is more productive if we understand why we disagree. This is why we have devoted much of this piece to explaining the values and motivations of each research culture. We urge neuroscientists to acknowledge that others may have good-faith disagreements about what counts as 'making progress'.

Respecting the value of each research culture calls for more than mere lip-service. Concretely, reviewers, grant-makers and hiring committees should refrain from imposing their preferred meaning of scientific progress when evaluating others' work. Instead, a fair critique should acknowledge that there can be different but equally valid ways of contributing to progress. The understanding that there are different but equally valid meanings of scientific progress should be passed on to trainees. While it is fair to teach one's approach to one's students, it is a disservice – to those students and the broader neuroscience community – to act as though that approach is the only valid one. These are concrete actions that every individual neuroscientist can take to reduce unnecessary conflict arising from a simple but pervasive misunderstanding.

Furthermore, a synthesis between these different research cultures can bring about tangible progress. Concretely, computational modelling provides one fruitful avenue for achieving this synthesis (although certainly not the only avenue). Coming up with a viable model that can actually be implemented forces the top-down neuroscientist to ensure that theories remain concrete and testable[6,7]. Conversely, model-development forces the bottom-up neuroscientist to confront the bigger picture by drawing a line between (putatively) necessary and unnecessary details, whose relevance is then demonstrated or refuted by the model's performance. Models can then integrate data and knowledge across different scales and domains of neuroscience[7]: for example, modelling the effects of pharmacological interventions by enriching one's macroscopic model with biological detail about receptors[8]. It can be especially valuable to engage in this modelling exercise with colleagues from the 'other culture', as the model represents a tangible output that grounds discussions by demanding convergence – combining the virtues of both research cultures.

In summary, as scientists we want our answers to be correct, and we also want to answer questions that matter. When these two desiderata pull in different directions, we need to decide how to weigh them. We advocate for transitioning away from the current pervasive but unacknowledged dichotomy and towards a more constructive mindset in which different dimensions of progress in neuroscience are explicitly acknowledged and their value is recognized – to temper their respective limitations and build on each other's strengths.

Fernando E. Rosas [1,2,3,4,5,6,16] ✉, **Andrea I. Luppi** [5,7,8,9,16] ✉, **Pedro A. M. Mediano** [5,10,11], **Morten L. Kringelbach** [5,9,12], **Luiz Pessoa** [13] **& Federico Turkheimer** [14,15]

[1]Sussex Centre for Consciousness Science, Department of Informatics, University of Sussex, Brighton, UK. [2]Sussex AI, Department of Informatics, University of Sussex, Brighton, UK. [3]Department of Brain Science, Imperial College London, London, UK. [4]Centre for Complexity Science, Imperial College London, London, UK. [5]Centre for Eudaimonia and Human Flourishing, University of Oxford, Oxford, UK. [6]Principles of Intelligent Behavior in Biological and Social Systems (PIBBSS), Prague, Czech Republic. [7]St John's College, University of Cambridge, Cambridge, UK. [8]Montreal Neurological Institute, McGill University, Montreal, Quebec, Canada. [9]Department of Psychiatry, University of Oxford, Oxford, UK. [10]Department of Computing, Imperial College London, London, UK. [11]Division of Psychology and Language Sciences, University College London, London, UK. [12]Centre for Music in the Brain, Aarhus University, Aarhus, Denmark. [13]University of Maryland, College Park, MD, USA. [14]Institute for Human and Synthetic Minds, King's College London, London, UK. [15]Department of Neuroimaging, Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK. [16]These authors contributed equally: Fernando E. Rosas, Andrea I. Luppi.

✉e-mail: f.rosas@sussex.ac.uk; andrea.luppi@psych.ox.ac.uk

Published online: 21 July 2025

References

1. Kandel, E. R. *In Search of Memory: The Emergence of a New Science of Mind* (W. W. Norton & Company, 2006).
2. Uffink, J. Boltzmann's work in statistical physics. *Stanford Encyclopedia of Philosophy* (Winter 2024 Edition) (eds. Zalta. E. N. & Nodelman, U.) https://plato.stanford.edu/archives/win2024/entries/statphys-Boltzmann (2024).
3. Bassett, D. S. & Sporns, O. Network neuroscience. *Nat. Neurosci.* **20**, 353–364 (2017).
4. Pessoa, L. The entangled brain. *J. Cogn. Neurosci.* **35**, 349–360 (2023).
5. Levy, N. Analytic and continental philosophy: explaining the differences. *Metaphilosophy* **34**, 284–304 (2003).
6. D'Angelo, E. & Jirsa, V. The quest for multiscale brain modeling. *Trends Neurosci.* **45**, 777–790 (2022).
7. Shine, J. M. et al. Computational models link cellular mechanisms of neuromodulation to large-scale neural dynamics. *Nat. Neurosci.* **24**, 765–776 (2021).
8. Deco, G. et al. Whole-brain multimodal neuroimaging model using serotonin receptor maps explains non-linear functional effects of LSD. *Curr. Biol.* **28**, 3065–3074 (2018).

Acknowledgements

The authors thank S. Raza and H. Scott-Fordsmand for insightful discussion about the sociology of science and theories of scientific progress. F.E.R. is supported by the PIBBSS Affiliate Program. A.I.L. is supported by St John's College, Cambridge, and an Early Career Award from the Wellcome Trust (grant number 226924/Z/23/Z). M.L.K. is supported by the Centre for Eudaimonia and Human Flourishing (funded by the Pettit and Carlsberg Foundations) and the Center for Music in the Brain (funded by the Danish National Research Foundation, DNRF117). L.P. is supported by the National Institute of Mental Health (MH071589), NIH, USA. F.T. is funded by the National Institute for Health and Care Research (NIHR) Maudsley Biomedical Research Centre (BRC).

Competing interests

The authors declare no competing interests.