

# Top-down and bottom-up neuroscience: Overcoming the clash of research cultures in neuroscience

Fernando E. Rosas<sup>1,2,3,4,\*</sup>, Andrea I Luppi<sup>3,5,\*</sup>, Pedro A.M. Mediano<sup>2,5,6</sup>,  
Morten L. Kringelbach<sup>3,7</sup>, Luiz Pessoa<sup>8</sup>, and Federico Turkheimer<sup>9</sup>

<sup>1</sup> University of Sussex, Brighton, UK

<sup>2</sup> Imperial College London, London, UK

<sup>3</sup> University of Oxford, Oxford, UK

<sup>4</sup> PIBBSS, Prague, Czech Republic

<sup>5</sup> University of Cambridge, Cambridge, UK

<sup>6</sup> University College London, London, UK

<sup>7</sup> Aarhus University, Aarhus, Denmark

<sup>8</sup> University of Maryland, College Park, USA

<sup>9</sup> King's College London, London, UK

\*Equal contribution

Correspondence: [f.rosas@sussex.ac.uk](mailto:f.rosas@sussex.ac.uk) (FER); [andrea.luppi@psych.ox.ac.uk](mailto:andrea.luppi@psych.ox.ac.uk) (AIL)

## Abstract

What should neuroscientific progress look like? As scientists, we want solid answers, but we also want to answer questions that matter. However, the inherent complexity of the brain forces a trade-off between these desiderata, bringing about two distinct research cultures in neuroscience that can be described as “top-down” and “bottom-up”. In this article we argue that both approaches are equally valid — but not openly recognised as such, leading to misunderstandings and unnecessary tension. Drawing on diverse fields, here we aim to dispel the latent tension between these research cultures and promote constructive interaction to advance brain science.

Passion for neuroscience often sparks from an early exposure to big-picture questions — e.g., how brains give rise to minds, or what consciousness is. Approaches to tackle such grand challenges tend to coalesce into two distinct categories, corresponding to different neuroscientific “research cultures” which diverge not only on research methods but even on *what it means to make progress*.

## Bottom-up neuroscience

Bottom-up research culture cares primarily about building knowledge over firm foundations, and drives progress through systematic accumulation of solid evidence and increasing detail. Accordingly, bottom-up approaches emphasise careful experimental design for studying neuroscientific scenarios by keeping everything as controlled as possible. When this cannot be achieved, they are willing to stretch the definition of their object of study until it becomes amenable to full experimental control. Scientific progress, then, derives from the systematic accumulation of solid evidence and increasing detail.

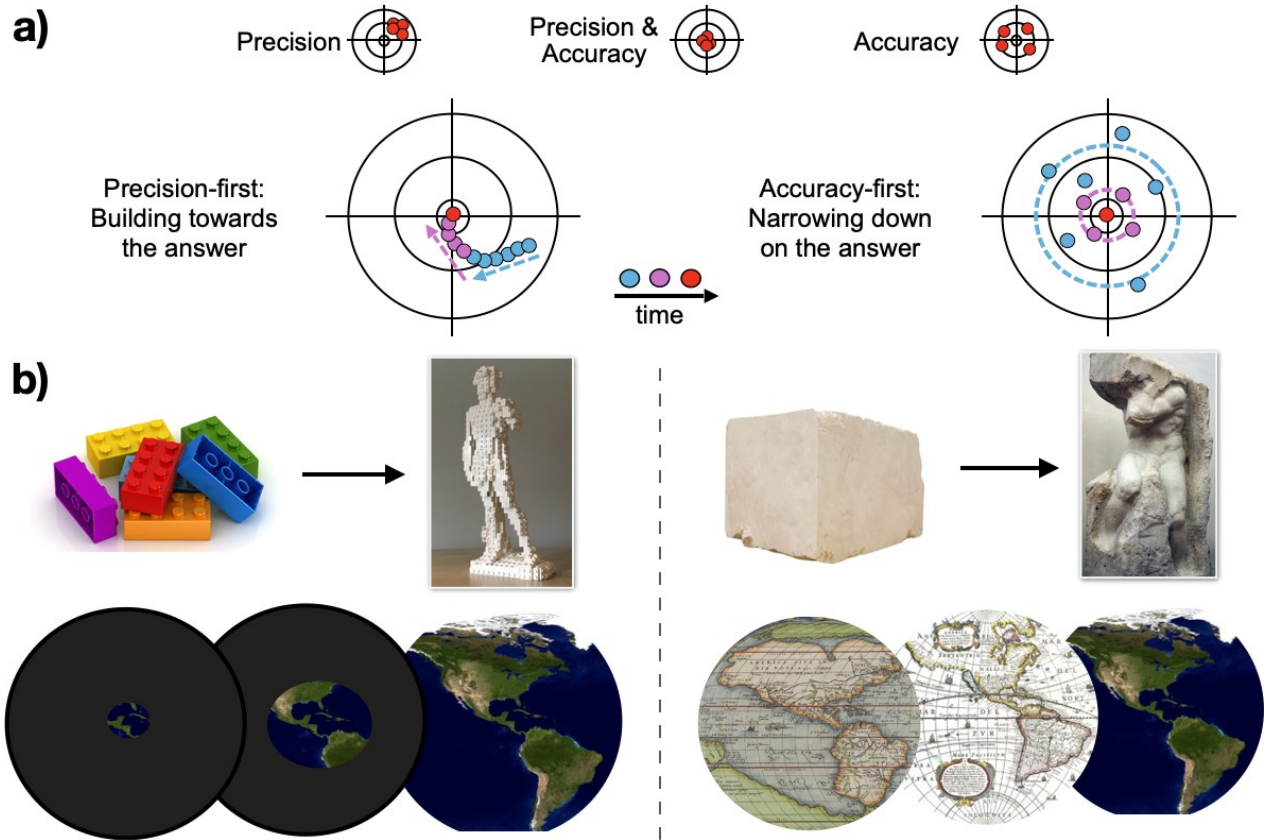
Examples include the tightly-controlled experiments of psychophysicists, or the Nobel-winning work of Eric Kandel that operationalised memory as habituation to repeated prodding in a sea slug, *Aplysia californica*. By accepting this — arguably drastic — simplification, Kandel’s team were able to leverage the exquisite experimental accessibility of *Aplysia*, characterising synaptic long-term potentiation with unprecedented detail.

## Top-down neuroscience

In contrast, the top-down research culture is driven by the desire to work on “big” questions directly, less willing 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 analytic 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 his ground-breaking results were correct, his often-incomplete arguments sometimes appeared as ‘logical jumps’<sup>1</sup>. Network neuroscience also exemplifies successful work based on top-down principles<sup>2</sup>: by

emphasising fundamental properties that the brain shares with other networked systems (e.g. cities, social networks), network neuroscience has become a flourishing neuroscientific sub-field, bridging scales and promoting translational discovery.



**Figure 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 focused around the target itself (higher accuracy). (b) Analogies from art and cartography. In art, 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.

## How the divergence can hinder progress in neuroscience

Because of the brain’s inherent complexity, from microscale to macroscale<sup>3</sup>, 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 prioritise *precision* and make progress by reducing bias from target, one step at a time; top-down approaches prioritise *accuracy* and make progress by “narrowing down” on the target (Fig.1).

It is not difficult to see how each approach may be criticised. 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 fMRI 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 about 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<sup>4</sup>. Analytic philosophers emphasised 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, on the other hand, devoted themselves to arguably “existential” questions, at the risk of failing to tame the complexity of such big-picture questions. Though 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 by misunderstanding each other’s priorities and motivations. Proponents of each culture mistakenly believe the other to be operating under their own same conception of progress, but improperly so. Such a situation is analogous to judging a statistical model purely based on 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 (e.g., 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 if 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 (though 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<sup>5,6</sup>. Conversely, model-development forces the bottom-up neuroscientist to confront the bigger picture by drawing a line between (putatively) necessary and unnecessary detail, 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<sup>6</sup>: for example, modelling the effects of pharmacological interventions by enriching one’s macroscopic model with biological detail about receptors<sup>7</sup>. Altogether, 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 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 where different dimensions of progress in neuroscience are explicitly acknowledged and their value is recognised — to temper their respective limitations and build on each other’s strengths.

## Acknowledgements

The authors thank Sebastian Raza and Helene 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 of 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. FET is funded by the National Institute for Health and Care Research (NIHR) Maudsley Biomedical Research Centre (BRC).

## References

1. Uffink, J. (2004). Boltzmann's work in statistical physics. *Stanford Encyclopedia of Philosophy*.
2. Bassett, D. S., & Sporns, O. (2017). Network neuroscience. *Nature neuroscience*, 20(3), 353-364.
3. Pessoa, L. (2023). The entangled brain. *Journal of cognitive neuroscience*, 35(3), 349-360.
4. Levy, N. (2003). Analytic and continental philosophy: Explaining the differences. *Metaphilosophy*, 34(3), 284-304.
5. D'Angelo, E. and Jirsa, V., (2022). The quest for multiscale brain modeling. *Trends in neurosciences*, 45(10), pp.777-790.
6. Shine, J. M., Müller, E. J., Munn, B., Cabral, J., Moran, R. J., & Breakspear, M. (2021). Computational models link cellular mechanisms of neuromodulation to large-scale neural dynamics. *Nature neuroscience*, 24(6), 765-776.
7. Deco, G., Cruzat, J., Cabral, J., Knudsen, G.M., Carhart-Harris, R.L., Whybrow, P.C., Logothetis, N.K. and Kringelbach, M.L., (2018). Whole-brain multimodal neuroimaging model using serotonin receptor maps explains non-linear functional effects of LSD. *Current biology*, 28(19), pp.3065-3074.