

# On the **Future** *of* **Consciousness**

---

ANDREW Y. LEE

University of Toronto, Philosophy

---

## **Abstract**

I argue that the most promising direction for the future of consciousness research is to follow a path that recurs across the history of science, where an early focus on “essence questions” (“What is  $x$ ?”) is replaced by a later focus on “structural questions” (“How do we model  $x$ ?”). I discuss the prospects for investigating the structure of consciousness, and I argue that to investigate such questions is to pursue an “objective phenomenology.”

---

## A Vision for the Future

What might state-of-the-art consciousness research look like a century from now? Any answer to this question will be a matter of speculation. But let's speculate—if only for a moment.

On one optimistic vision of the future, the study of consciousness will have progressed from an emerging field into a mature science. Our scientific theories illuminate not only the conscious experiences of humans, but also those of animals and artificial intelligences. Our mathematical models map the spaces, structures, and dynamics of conscious experiences, enabling precise descriptions of their form and character. Our concepts have shifted from rough distinctions rooted in folk intuitions to refined tools that carve experiences at their joints. A new field of applied consciousness science utilizes basic consciousness research to enhance technology, medicine, art, and entertainment. While there are still deep disagreements and unresolved debates, there is increasing convergence across methods, language, and frameworks. Though the science of consciousness has not attained the status of the great natural sciences—physics, chemistry, biology—it is starting to move in that direction.

I don't know whether this scenario will ever come to fruition. Perhaps it will never be more than a speculative fiction. But science often advances in ways that are hard to foresee, and the future often turns out stranger than we imagined. This vision—I believe—is at least an open possibility. And if that's so, then we might wonder: What questions should we ask today in order to move us closer to this vision of tomorrow?

## The Purpose of this Piece

The purpose of this piece is to paint a picture of what consciousness research could become—and to articulate the direction I believe is most promising for reaching that picture. The [science of consciousness](#) is moving forward, with great momentum, at an accelerating pace. But where will we eventually land if we simply follow our current trajectory?

Occasionally, I encounter skepticism about whether there has been genuine progress in understanding consciousness. After all—some will say—consciousness has long been one of the great philosophical and scientific [mysteries](#), and we may seem no closer to solving it today than we were centuries ago. An aim of this essay is to challenge a key presupposition behind that attitude. There is indeed something deeply puzzling about consciousness. But consciousness research—I will argue—is about much more than just solving that mystery.

My view is more optimistic: I think there has been tangible progress. But even the most optimistic should agree that we are a long way from realizing that vision of the future. While the field's diversity of approaches makes it rich and exciting, it's also indicative of its fragmentation. While there are flourishing new strands of research—from animal

---

minds to altered states to artificial intelligence—our basic methods and concepts remain rudimentary. While I have my own sympathies about which approaches are most promising and which views most likely to be true, I suspect that our theories today may be merely rough outlines of a more developed science of consciousness to come.

But the field of consciousness research is still in its early stages: it has yet to find its Newton or its Darwin. The most exciting ideas are likely still to be discovered. And while we cannot yet know their exact form, we can still ask which questions we should pursue now to give ourselves the best chance of reaching those discoveries.

I'll soon offer an answer. But let me begin with a question that I think captures the current chapter in consciousness research.

### The Big Question

Consciousness is often portrayed as a singular challenge that resists standard methods of scientific explanation. The field of consciousness research is often framed as the endeavor to face up to that challenge. At the heart of the challenge is the following question:

THE BIG QUESTION

What *is* consciousness?

There are various ways to express the big question: What makes it the case that there's something it's like to be a creature? What grounds the fact that an entity has a subjective point of view? What demarcates subjects of experience from mere objects? What is the ingredient that separates conscious states from unconscious states? While there are subtle differences between these expressions, each asks a version of the big question.

When people talk about “[solving consciousness](#),” they typically mean answering the big question. When people claim that we [lack a definition](#) of consciousness, what's usually meant is that we haven't yet answered the big question. And when people talk about “the question” of consciousness—as though the subject-matter could be distilled to a single problem—they often have in mind the big question.

If you look at research in the science of consciousness, you'll find that the most common experimental paradigms likewise target the big question. Many experimental techniques are forms of [contrastive analysis](#), a method that compares two mental states that are as similar as possible, except that one state is conscious and the other isn't. The aim is to isolate the ingredient present in the conscious state but absent in the unconscious state. In other words, contrastive analysis is a method to figure out what consciousness *is*.

You might already be wondering what *else* there could be to consciousness research. Isn't it obvious that the aim is to explain what consciousness is? I agree that the big question is important; I'd even argue that it belongs at the core of the field. But I think the

conceptual weight of the big question can distort how we think about the scope and target of consciousness research.

The big question is an example of an **essence question**—a question about the underlying nature of a kind. Other examples include classic questions such as *What is water?* / *What is light?* / *What is life?* / *What is language?* / *What is computation?* These questions ask what it is in virtue of which something counts as a member of a kind or instance of a phenomenon. The answers to essence questions can often be expressed with simple identity statements, of the form ‘x is y’. Consider:

- Water is H<sub>2</sub>O.
- Light is electromagnetic radiation.
- Life is the capacity for self-maintenance, reproduction, and metabolism.
- Language is a structured system of symbols used for communication.
- Computation is rule-governed information processing.

Similarly, answers to the big question can usually be expressed as follows:

THE BIG ANSWER

Consciousness is \_\_\_\_\_.

A “big answer” tells us what it is that makes an entity conscious. Any such answer draws a line in reality, dividing conscious entities from nonconscious entities. If we know the big answer—and can determine whether an entity has the property that marks that divide—then we will be able to tell whether it is conscious. Here are a few candidates for big answers:

theory	answer
global workspace theory	global workspace
higher-order theory	higher-order representation
integrated information theory	maximally integrated information
recurrent processing theory	recurrent feedback loops
substance dualism	immaterial soul

The big question is related to the **hard problem** of consciousness, the challenge of explaining why consciousness arises from physical systems at all. But the big question isn’t simply the hard problem in question form. The big question is a question of *what*; the

---

hard problem is a problem of *why*. One could accept a big answer without solving the hard problem. For example, one could accept that to be conscious is to have higher-order thoughts while still feeling puzzled about why there's anything it's like to be such systems. Perhaps solving the hard problem is necessary to fully justify a big answer. But an answer to one doesn't necessarily entail an answer to the other.<sup>1</sup>

### What the Big Question Leaves Open

The big question is important. It's one of the hardest questions in science, and it connects to some of the deepest issues in philosophy. It deserves—in my opinion—to play a prominent role in consciousness research. But its sheer size also eclipses other important areas of research. The study of consciousness—I will argue—is not simply the search for the big answer.

Suppose we suddenly discover the big answer. Imagine that we establish—beyond reasonable doubt—that to be conscious is to have property *x*. Substitute whatever you like for '*x*': having a global workspace / higher-order awareness / a maximal amount of integrated information / recurrent feedback loops / an immaterial soul / etc. Would we then have arrived at the utopian scenario described earlier?

Not quite, and not yet. Even if we were to find the big answer, many questions about consciousness would remain open. The big answer would help us determine *whether* a given entity is conscious, but it wouldn't tell us *what it's like* to be that entity. It's one thing to know *which* things are conscious; it's another thing to understand the different *ways* in which something can be conscious. The big answer marks a dividing line between conscious entities and nonconscious entities, but it wouldn't tell us how to map the space on the conscious side of that divide.

To make this concrete, think about conscious artificial intelligence. To determine whether [artificial consciousness](#) is possible at all, we need to make progress on the big question. But even if we knew that certain artificial systems are conscious, many questions would remain unsettled. What are the experiences of artificial systems like? How do they differ from human experiences, and how do they differ from each other? Could we ever understand the experiences of artificial systems with radically different cognitive architectures?

Consciousness is often described as a matter of the "[lights being on inside](#)." Within this metaphor, the big question is about the light: Which entities have an inner light, and

---

<sup>1</sup> Is it possible to accept a solution to the hard problem without taking a stance on the big question? Consider how one might think that the hard problem can be solved by (i) appealing to the distinctive cognitive roles of phenomenal concepts, or by (ii) taking consciousness to be the intrinsic nature of matter, or by (iii) denying that there even is an explanatory gap between physical facts and phenomenal facts. Each of these ideas has been floated as a potential solution to the hard problem, but each is compatible with any of the big answers mentioned in this essay.

---

which are such that all is dark inside? But what it's like to be a creature is a matter not only of whether the lights are on, but also of what's illuminated. Even if we have a theory of the light, we still might lack a theory of what becomes illuminated. What exactly is illuminated inside the mind of a bat, or an octopus, or an AI?

Perhaps answering the big question will help with answering these other questions. But knowing the essence of a kind doesn't automatically mean understanding its structure. For example, we might know that computation is the implementation of algorithmic processes without understanding the space of complexity classes, or that water is H<sub>2</sub>O without understanding fluid dynamics and phase transitions.

As an example, consider [global workspace theory](#), the view that to be conscious is to have a "global workspace"—a central executive system whose information is consumable by other cognitive systems. This theory gives us a big answer. But it doesn't yet tell us what it's like to be a given conscious subject. It's possible, of course, to extend this theory, such as by combining it with a theory of mental representation. But doing so means going beyond the big question.

Now, a significant amount of consciousness research—arguably the majority—is concerned with more than the big question. Yet the big question exerts outsized influence on how the field is framed. Consciousness is often portrayed as a [singular unsolved puzzle](#), rather than a rich subject matter with many subdomains. But when we treat the aim of consciousness research as simply finding the big answer, we undersell both what the field is already doing and what it could become.

There's no doubt that settling the big question would be a massive step forward. It would be like discovering the nature of water, or heat, or light. But just as those discoveries still left open many questions about those phenomena, so too would the big answer leave open many questions about consciousness. Finding the big answer wouldn't be enough to bring us to the vision of the future sketched earlier. What else, then, is needed?

### **From Essence to Structure**

The big question is an essence question. I mentioned earlier some other examples of essence questions, associated with other scientific fields:

- What is water?
- What is life?
- What is language?
- What is light?
- What is computation?

---

Each of these questions is important and interesting. But none defines the scope of its corresponding science; the fields of chemistry, biology, linguistics, physics, and computer science do much more than simply answer essence questions like the ones above. Perhaps those questions animated early stages of inquiry. But they aren't focal points in contemporary science.

It's useful to contrast essence questions with what I'll call **structural questions**. These are questions about modeling, mapping, measurement, and mechanisms, rather than about underlying natures. Here are some examples, drawn from the same scientific domains as above:

- How do we model fluid dynamics?
- What are the mechanisms of reproduction?
- Is light continuous or discrete?
- What are the basic building blocks of syntactic trees?
- What are the most efficient algorithms for a given computational task?

'Structure', in the most general sense, concerns the features that can be wholly captured through mathematical descriptions. Structural questions may concern the space of possible states of a system, the dimensions along which members of a domain vary, the parts into which an object can be decomposed, or the relations that a kind bears to other kinds. Structure is a matter of form—differentiation, composition, variation, pattern, quantity—rather than intrinsic nature.

The distinction between essence questions and structural questions isn't always clear cut. There are questions that lie at the borderline, or that belong to both, or that belong to neither. But the categories are intuitively compelling and—I believe—conceptually fruitful. Understanding the contrast can help us understand our present stage in consciousness research—and perhaps where we might go in the future.

### **A Pattern across the Sciences**

Here's an intriguing observation: as a scientific field matures, its focus tends to shift from essence questions to structural questions.

Take chemistry. Essence questions in chemistry include *What is water?* and *What is gold?*. The answers to these questions—that water is H<sub>2</sub>O and that gold is the chemical element with atomic number 79—were remarkable discoveries. But the science of chemistry does much more than merely answer essence questions. In fact, one of its most significant breakthroughs—the construction of the periodic table—is a discovery about the systematic relationships between chemical elements, rather than about the nature of any particular chemical element.

---

This pattern recurs in other fields. For example, it's vividly illustrated in Hasok Chang's *Inventing Temperature*, which traces the development of thermometry and shows how much of its progress was driven by questions about modeling and measuring temperature. Questions about the essence of temperature, in turn, were reframed and advanced in light of structural investigation. This case study illustrates not only the transition from essence to structure, but also how we often discover essences only by investigating structure.

While state-of-the-art research tends to be focused on structural questions, essence questions never vanish entirely. New essence questions arise (*What is dark matter?*), and old ones return in new forms (*What is intelligence?*). Over time, however, the most foundational essence questions tend to migrate from the realm of science to the realm of philosophy (*What is life?* / *What is language?* / *What is computation?*). In later stages, the essence questions that remain active (*What is working memory?* / *What is acidity?* / *What is quantum entanglement?*) are typically accompanied by a suite of structural questions.

The pattern that emerges is a transition from essence to structure. I believe this pattern may hint at the future of consciousness research.

### **The Structure of Consciousness**

Conscious experiences are richly structured. Think about how perceptual experience can be decomposed into different sensory modalities (vision, audition, olfaction, etc.); how color experiences vary in dimensions corresponding to hue, saturation, and brightness; how pain experiences come in different magnitudes; and how spatial experience decreases in acuity from the center of one's visual field to its periphery.<sup>2</sup>

Here are some of the most foundational structural questions about consciousness:

- What are the atomic elements of conscious experiences?
- What are the dimensions of consciousness?
- Does consciousness come in degrees?
- How are the quality spaces for conscious experiences structured?
- Are conscious experiences continuous or discrete?
- How can conscious experiences be mathematically modeled?

---

<sup>2</sup> When talking about the "structure of consciousness," it's worth distinguishing two senses of 'structure'. The *mathematical* sense of 'structure' used here is distinct from the *invariance* sense of 'structure' that occurs in historical phenomenology. In the latter context, "structural features" are the invariant, essential features of conscious experiences. While invariant features of experience can be described mathematically, so too can variable features of experience (such as the particular arrangement of color qualities across your visual field at this moment).

---

These questions won't be settled simply by finding the big answer. Even if we know which things are conscious, we won't necessarily know how the experiences of one conscious system relate to the experiences of another, how to carve experiences at their natural joints, how to decompose conscious experiences into parts and dimensions, or how to model them mathematically.

Some might contend that we would have to settle the big question before studying the structure of consciousness. But the history of science tends to tell a different story: we often answer essence questions only by making progress on structural questions. For example, the nature of genes became clear only after the discovery of the double-helix structure of DNA, and the equations of thermodynamics led to the discovery that heat is molecular motion. Perhaps the story of consciousness will follow a similar arc: we may not fully understand what consciousness is until we better understand how it's structured.

Structural questions about consciousness are important. Not only are they interesting in their own right, but they also form a key part of what Anil Seth has called “[the real problem](#)” of consciousness: the search for systematic mappings between brain states and conscious states. To construct such mappings, we need a better understanding of the structures of conscious experiences themselves. More specifically, we need maps and models of **phenomenology**: what experiences feel like from a first-person point of view. Otherwise, we'll be trying to build a bridge without knowing what it's meant to reach.

As a case study, consider color experience. A simple but powerful observation is that color experiences stand in systematic structural relations to each other. For example, red is more similar to orange than to green. These similarity relations can be mapped onto a [quality space](#) model—a geometric representation of a domain of qualities. Color qualities that are more similar are modeled by points that are closer in the geometric space, and the dimensions of the space represent ways of varying with respect to color (hue, saturation, brightness). Such a model enables us to understand the structure of color experience, and even to [draw novel inferences](#) about possible color experiences.

When we study the structure of consciousness, we are mapping the state space of conscious experiences. There's a vast realm of possible experiences—human, animal, artificial—and we directly access only a tiny corner of that space. To answer the big question—to find the essence of consciousness—is to identify the universal property that unites this whole space. But the aim of consciousness research ought to be not only to figure out what's common across all conscious entities, but also to analyze the structure of this space.

The study of the structure of consciousness has rich historical precedents, spanning a variety of intellectual traditions. These include classical Buddhist philosophers (especially in the [Abhidharma](#) and [Yogācāra](#) traditions), early modern rationalists (such as [Leibniz](#)), British empiricists (such as [Hume](#)), German idealists (such as [Kant](#)), 19<sup>th</sup>-century psychophysicists (such as [Fechner](#)) and introspective psychologists (such as [Wundt](#)), and

---

20<sup>th</sup>-century continental phenomenologists (such as [Husserl](#)), and early analytic philosophers (such as [Carnap](#)). Perhaps the 21<sup>st</sup>-century science of consciousness can open a new chapter—drawing on these foundations while making use of modern methods, technology, concepts, and theories to advance questions that were previously out of reach.

In fact, a number of research programs in consciousness science today are already structural in character. These include [mathematical consciousness science](#), [neurophenomenology](#), and [computational phenomenology](#), as well as theoretical frameworks that emphasize structure, such as [integrated information theory](#) and [predictive processing theory](#). There are also pockets of work on [perception](#), [attention](#), and other aspects of the mind that focus on structural questions about consciousness.

This indicates that research on the structure of consciousness is beginning to flourish. While it has yet to take center stage, there are seeds that have begun to bloom. Yet relevant work remains highly fragmented: different approaches operate with divergent methods, theories, and guiding assumptions. You might then wonder: could these diverse frameworks be unified?

### **A Common Model of Consciousness**

It's widely agreed that conscious experiences are rich, complex, and multifaceted in structure. But there's no general framework for thinking systematically about that structure. The development of such a framework would mark a major breakthrough for consciousness research.

To crystallize these ideas, imagine a hypothetical **Common Model of Consciousness**—a widely accepted framework for systematically describing the structures of conscious experiences. Such a model would enable precise descriptions of conscious experiences across a wide range of conscious subjects, integrate experimental and theoretical insights, and provide a shared foundation for future research.

The label 'Common Model of Consciousness' might, for some, evoke the Standard Model in physics. But that comparison is, at best, aspirational: the Standard Model is vastly more developed than anything we might currently hope for in consciousness research. A more cautious comparison—given the current stage of the field—is Newtonian mechanics. Even though Newton's framework was later superseded by relativity and quantum mechanics, it played a pivotal role in using mathematical tools to model physical reality. Before that era, our grasp of the physical world was largely qualitative. Newtonian mechanics offered a systematic structural framework that unified a wide range of physical phenomena and formalized our understanding of the physical world.

What might a Common Model of Consciousness look like? At this stage, it's hard to say. But one promising starting point, in my view, is to focus on two basic features of experience: qualities and locations. Conscious experiences are marked by various kinds of

---

qualities—color, sound, smell, and so on—which stand in relations of similarity, precision, vividness, and intensity. These qualities are distributed across the visual field, the auditory field, bodily space, and other kinds of locative fields. The road to a Common Model of Consciousness might begin with a general framework for modeling quality spaces, locative fields, and how they interact. Though this may initially appear simple, it quickly becomes complex once we start thinking carefully about the [formal structure](#) such a framework would require. Still, this might serve as one component of a future Common Model of Consciousness.

At this point, you might wonder whether there's a key difference between consciousness and other targets of scientific inquiry. Models are good at describing structure. But consciousness—you might think—is more than just structure. A model can't reveal the qualitative character of experience; even the richest structural description seems bound to leave something essential out. And isn't that missing ingredient precisely what we are trying to investigate? This leads to a puzzle: if structural descriptions of conscious experiences leave out qualitative character, then how could the future of consciousness research lie in studying structure?

I'll explain, in a moment, how I think this puzzle can be resolved. But let me first turn to an interlude—about a famous passage in the literature on consciousness that presents a closely related puzzle.

### Objective Phenomenology

Perhaps the most widely read article in contemporary consciousness research is Thomas Nagel's "[What is it like to be a bat?](#)" Towards the end of the article, there is a curious and cryptic passage where Nagel speculates about the possibility of an **objective phenomenology**: a way of characterizing how conscious experiences feel that would be understandable even by those unable to have the experiences in question. Is there, for example, a way of describing bat experiences that could be understood even by humans?

"Setting aside temporarily the relation between the mind and the brain, we can pursue a more objective understanding of the mental in its own right... This should be regarded as a challenge to form new concepts and devise a new method—an objective phenomenology... Though presumably it would not capture everything, its goal would be to describe, at least in part, the subjective character of experiences in a form comprehensible to beings incapable of having those experiences."

—Thomas Nagel, "What is it like to be a bat?" (1974: 449).

You might doubt that such an endeavor is even coherent. To understand what it's like to be a bat, one must have had experiences sufficiently similar to a bat's experiences. But humans—by necessity—have a bounded range of experiential capacities. This may impose principled limits on our epistemic access to other kinds of experiences. Perhaps

---

those limits would shift if we were somehow able to acquire new experiential capacities. But given the kinds of creatures that we are—with the experiential repertoire that we have—it seems impossible for us to understand what it’s like to be a bat.

The example of bats is simply illustrative. The core problem is general. How could we ever understand what it’s like to be a creature whose experiences are radically different from our own? It’s one thing to think that we can figure out whether slugs, fetuses, and AIs are conscious in the first place. But it’s another thing to think that we can come to understand what their conscious experiences are like.

We are left—again—with a puzzle. The idea of an objective phenomenology feels intriguing and appealing. But is it genuinely intelligible?

### **Objectivity, Phenomenology, Structure**

The puzzle about modeling consciousness and the puzzle about objective phenomenology share a common solution. The [solution](#) is that the way an experience is structured is itself an aspect of its phenomenology.<sup>3</sup>

What an experience is like depends not only on which qualities are present, but also on how those qualities are arranged and how they relate to one another. A full description of your visual experience, for example, must mention the structure of your visual field, the distribution of color qualities across it, and the similarity relations between those qualities. To describe the structure of an experience is therefore to describe one aspect of its phenomenology.

This insight lays the groundwork for an objective phenomenology. A structural characterization of an experience is both objective and phenomenal: objective because it can be grasped even by those unable to have that experience, and phenomenal because it describes—at least in part—what that experience is like.

Think again about what a Common Model of Consciousness might reveal about the experiences of bats, octopuses, or AIs. It wouldn’t enable us to simulate those experiences firsthand. But it might still tell us whether echolocation experience unfolds as a series of discrete pulses or as a continuous stream; whether each of an octopus’s tentacles is associated with a partially independent stream of experience; and whether an AI’s experiences contain any counterpart to our visual field. These structural descriptions would capture—in part—what it’s like to be a bat, an octopus, or a conscious AI.

---

<sup>3</sup> The idea that structural features of experience are objective was anticipated—though not developed—by Nagel himself. Shortly after the quoted passage, he speculated that “structural features of perception might be...accessible to objective description, even though something would be left out,” See Nagel [1974: 450].

---

When the experiences in question lie far beyond the bounds of human phenomenology, we may be able to glimpse only a faint trace of their qualitative character. But descriptions of their structure would still be accessible to us, at least when cast in formal, mathematical terms. And even the most abstract structural descriptions, stripped of any qualitative content, can still illuminate the contours of an experience.

Consider someone born blind, who has never had any visual experiences. To someone in that position, the visual experiences of the sighted may be as alien as the echolocation experiences of a bat. But if they're presented with a rich description of the structure of visual experience—how color qualities form a three-dimensional similarity space, how the resolution of the visual field gradually declines outside its center, how features are bound together in object perception—then they could gradually come to understand more and more about visual phenomenology. Their grasp may be limited to the form of visual experience, rather than its quality. But—as noted above—form is itself an aspect of phenomenology.

Even an ideal model of consciousness may not fully capture what it's like to see red, feel pain, or smell cinnamon. But such a model could still reveal how red fills a portion of the visual field, the quantitative structure of pain experience, and the region of a quality space that maps to cinnamon experience. Even when we have only a hazy grasp of the intrinsic qualities of an experience, we might still be able to study its blueprint. And when those qualities are radically alien, the blueprint may be the best guide we could hope for.

## The Future of the Field

This essay has sketched a vision for the future of consciousness research by extrapolating a pattern from the history of science. The study of consciousness—I've argued—is about much more than what consciousness *is*. While answering the big question would solve a major mystery, it would still leave open many questions about the kinds, characters, and structures of conscious experiences

Structural questions about consciousness are still underexplored. But interest in these questions has [gained momentum](#) in recent years, with a rich range of [new work](#). Consciousness science may already be in transition—from its essence phase to its structural phase, from speculative beginnings to a systematic science.

To fully reach that stage, we may need the kind of structural breakthrough that has transformed other fields. We might hope for the analogue—in consciousness science—of Newtonian mechanics, Darwin's theory of natural selection, Mendeleev's periodic table of elements, or Turing's theory of computation. What form that framework might take is an open question. But we can be sure that building it will require moving beyond the big question.

A structural turn in consciousness research could have implications beyond basic science. It might reframe how we think about other minds: human, animal, artificial—

---

perhaps others still unknown. It might yield insight into which features of experience and which kinds of systems matter morally. It might provide the theoretical groundwork for a future field of applied consciousness science. And it might offer new tools for understanding ourselves, by shedding light on the shapes and textures of our subjective lives.

When I think about the space of possible experiences, I find myself better able to appreciate how strange our own experiences are. It's easy to take our own mode of experience for granted, as though it were the default pattern or the universal standard. But our experiences are just a speck in a vast sea of possibilities. To see our stream of consciousness as one possibility among countless others—many far more alien than those of bats or octopuses—is to undergo a shift in perspective, like realizing that humans are neither at the center of the universe nor the apex of life.

I don't know whether the vision I sketched at the outset of this essay will ever be realized. But I began by asking what questions we should think about today to move closer to that vision—and I've now argued for an answer. If consciousness research ends up following the arc of the most successful sciences, then its future lies in asking not only what consciousness is, but also in mapping and modeling its structure. That's a pattern that recurs across the history of science—and it may well be the most promising path for the future of consciousness.

---

## References

- Albantakis L, Barbosa L, Findlay G, Grasso M, Haun AM, Marshall W, et al. (2023) Integrated information theory (IIT) 4.0: Formulating the properties of phenomenal existence in physical terms. *PLoS Comput Biol* 19(10): e1011465. <https://doi.org/10.1371/journal.pcbi.1011465>
- Aristotle. *De Anima*. Translated by J. A. Smith. In *The Complete Works of Aristotle: The Revised Oxford Translation*, edited by Jonathan Barnes, Vol. 1, 641–692. Princeton: Princeton University Press, 1984.
- Baars, Bernard. (2005). Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Progress in Brain Research* 150: 45-53.
- Berkovich-Ohana, Aviva ; Dor-Ziderman, Yair ; Trautwein, Fynn-Mathis ; Schweitzer, Yoav ; Nave, Ohad ; Fulder, Stephen & Ataria, Yochai (2020). The Hitchhiker's Guide to Neurophenomenology – The Case of Studying Self Boundaries With Meditators. *Frontiers in Psychology* 11:1680.
- Burkeman, O. (2015, January 21). Why can't the world's greatest minds solve the mystery of consciousness? *The Guardian*. <https://www.theguardian.com/science/2015/jan/21/-sp-why-cant-worlds-greatest-minds-solve-mystery-consciousness>
- Butlin, P., Long, R., Elmoznino, E., & Bengio, Y. (2023). Consciousness in Artificial Intelligence: Insights from the Science of Consciousness. <https://doi.org/10.48550/arXiv.2308.08708>
- Chalmers, David (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies* 2 (3):200-19.
- Chang, Hasok (2004). *Inventing Temperature: Measurement and Scientific Progress*. Oxford University Press.
- Churchland, Paul M. (2005). Chimerical colors: Some phenomenological predictions from cognitive neuroscience. *Philosophical Psychology* 18 (5):527-560.
- Clark, Austen (ed.) (2000). *A Theory of Sentience*. New York: Oxford University Press.
- Fechner, *Elements of Psychophysics*, trans. H. E. Adler (New York: Holt, Rinehart and Winston, 1966).
- Fink, Sascha Benjamin ; Kob, Lukas & Lyre, Holger (2021). A structural constraint on neural correlates of consciousness. *Philosophy and the Mind Sciences* 2.
- Frank, A. (2025, July 8). Why science hasn't solved consciousness (yet). *Noema*. <https://www.noema.com/why-science-hasnt-solved-consciousness-yet>

- 
- Hanlon, Michael. "The mental block." Aeon. <https://aeon.co/essays/will-we-ever-get-our-heads-round-consciousness>
- Hoel, E. (2023, August 30). Consciousness is a great mystery. Its definition isn't. <https://www.the-intrinsicperspective.com/p/consciousness-is-a-great-mystery>
- Hohwy, J., & Seth, A. (2020). Predictive processing as a systematic basis for identifying the neural correlates of consciousness. *Philosophy and the Mind Sciences*, 1(II). <https://doi.org/10.33735/phimisci.2020.II.64>
- Husserl, Logical Investigations, trans. J. N. Findlay (Routledge, 2001).
- Jarrett, Christian (2020, April 5). Consciousness: How can we solve the greatest mystery in science? BBC Science Focus. <https://www.sciencefocus.com/the-human-body/consciousness-how-can-we-solve-the-greatest-mystery-in-science>
- Kleiner, Johannes (2024). Towards a structural turn in consciousness science. *Consciousness and Cognition* 119 (C):103653.
- Kulstad, Mark and Laurence Carlin, "Leibniz's Philosophy of Mind", *The Stanford Encyclopedia of Philosophy* (Winter 2020 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2020/entries/leibniz-mind/>>.
- LeDoux, Michel, & Lau, A little history goes a long way toward understanding why we study consciousness the way we do today, *Proc. Natl. Acad. Sci. U.S.A.* 117 (13) 6976-6984, <https://doi.org/10.1073/pnas.1921623117> (2020).
- Lee, Andrew Y. (2021). Modeling Mental Qualities. *The Philosophical Review* 130 (2):263-209.
- Lee, Andrew Y. (2024). Objective Phenomenology. *Erkenntnis* 89 (3):1197–1216.
- Lee, Andrew Y. (forthcoming). The Light & the Room. *Introducing Philosophy of Mind, Today* (eds. Devin Curry, Louise Daoust). Routledge.
- Lee, Andrew Y. & Fink, Sascha Benjamin (2025). Structuralism in the Science of Consciousness: Editorial Introduction. *Philosophy and the Mind Sciences*.
- Leitgeb, Hannes and André Carus, "Rudolf Carnap", *The Stanford Encyclopedia of Philosophy* (Fall 2024 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/fall2024/entries/carnap/>>.

- 
- Mudrik, L. (2025) Towards a science of consciousness: Promises and challenges of theory convergence. *Neuroscience & Biobehavioral Reviews*, 158, 105599. <https://doi.org/10.1016/j.neubio-rev.2024.105599>
- Nagel, Thomas (1974). What is it like to be a bat? *Philosophical Review* 83 (4):435-50.
- Pitt, David (2024). *The Quality of Thought*. Oxford, UK: Oxford University Press.
- Ronkin, Noa, "Abhidharma", *The Stanford Encyclopedia of Philosophy* (Summer 2022 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2022/entries/abhidharma/>>.
- Rosenthal, D. (2025). Objective foundations for the study of mental qualities. *Philosophy and the Mind Sciences*, 6. <https://doi.org/10.33735/phimisci.2025.11526>
- Sandved-Smith, L., Bogotá, J. D., Hohwy, J., Kiverstein, J., & Lutz, A. (2025). *Deep computational neurophenomenology: A methodological framework for investigating the how of experience*. *Neuroscience of Consciousness*, 2025(1), niaf016. <https://doi.org/10.1093/nc/niaf016>
- Seth, A. K. (2016). "The Real Problem." *Aeon*. <https://aeon.co/essays/the-hard-problem-of-consciousness-is-a-distraction-from-the-real-one>
- Seth AK (2018). Consciousness: The last 50 years (and the next). *Brain and Neuroscience Advances*. 2018;2. doi:10.1177/2398212818816019
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237(4820), 1317–1323.
- Szanyi, Szilvia, "Yogācāra", *The Stanford Encyclopedia of Philosophy* (Fall 2024 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/fall2024/entries/yogacara/>>.
- Watzl, Sebastian (2017). *Structuring Mind. The Nature of Attention and How it Shapes Consciousness*. Oxford, UK: Oxford University Press.
- Wundt, Principles of Physiological Psychology, trans. Edward Titchener (London: Swan Sonnenschein, 1904).
- Zimmer, C. (2021, March 18). How close is science to solving the problem of consciousness? *Discover Magazine*. <https://www.discovermagazine.com/mind/how-close-is-science-to-solving-the-problem-of-consciousness>