

THE IN A HURRY SERIES

Consciousness

in a Hurry

The hardest problem, made approachable

Max Brocklesby

BOOKSINAHURRY.COM

Contents

The Whole Thing in One Page 4

Why You Should Care 6

The Core Ideas 8

How It Actually Works 22

What People Get Wrong 32

Use It 37

Terms 41

Go Deeper 47

Notes and Sources 48

Bibliography 54

The Whole Thing in One Page

A brain scan can record changes while someone sees red. It does not display red as that person experiences it. Consciousness is familiar from within and unusually difficult to explain from outside. The popular picture of an inner light coming on when enough brain is active hides the difficulty. There is no little observer inside the head to whom the brain can deliver its finished work.

The subject becomes clearer when one word does less work. Is any experience present? What does it contain? Which information is available for reasoning, voluntary action and report? Why does the relevant physical activity involve experience at all? These questions distinguish presence, content, access and explanation without assuming that they name separate biological compartments. The final question is the hard problem associated with David Chalmers. Whether it needs a new ingredient in nature or a better account of what brains do remains disputed.

The empirical progress comes from breaking apart things that normally coincide. Visual information can guide behaviour without ordinary acknowledged seeing. Attention can be occupied while an unexpected event goes unnoticed. A dream can contain a world while its owner barely responds to the room. Severe injury can remove the route for a visible answer while leaving capacities that a different test detects. Consciousness cannot be identified with any one of those outward performances without examining how the connection works.

Nor is experience only a picture of the surroundings. Bodily feeling gives it urgency; the sense of being located in a body supplies a point of view. Experiments on bodily ownership show that even the feeling of which hand belongs to you depends on the combination of signals. This makes selfhood something to explain, not evidence for an extra spectator and not proof that the person is unreal.

Brain organisation matters more than a simple count of activity. Intervention can reveal causal contributions, and measures of the response to

stimulation can help distinguish states. Theories then propose different decisive operations. Global workspace accounts emphasise broad availability; recurrent-processing accounts emphasise local feedback. Higher-order theories concern representation of one's own mental states. Integrated information theory makes a stronger identification between experience and intrinsic causal structure. Predictive approaches explain how bodily and sensory content can be constructed. Their disagreements become useful when they produce different testable expectations.

Those expectations have begun to face deliberately adversarial experiments. The results constrain particular claims rather than deliver a final winner. A test using visible images cannot automatically settle every kind of experience. Equally, a failed brain-based command task can miss someone whose ordinary behaviour already demonstrates awareness. Knowing what a measurement can miss is part of knowing what it means.

The practical consequence reaches patients, other animals and artificial systems. Human-like performance and experience must not be confused, but uncertainty does not erase the possible cost of being wrong. Evidence should determine confidence; the stakes and cost of protection also enter decisions. The private familiarity with which the subject begins ends in a public responsibility: judge another's possible experience by how well the available methods can reveal it, rather than by the convenience of the answer.

That is the book.

Why You Should Care

In 2006, researchers asked a woman diagnosed as being in a vegetative state to imagine playing tennis, then to imagine walking through her home. She was inside a functional magnetic resonance imaging scanner. The two instructions produced different patterns of activity resembling those of healthy volunteers doing the same tasks. Her body gave no reliable answer to commands that an examiner could see. The recorded activity supplied evidence that she could follow them nonetheless.

The experiment changed the meaning of an absence. It did not show that every unresponsive patient is secretly intact, or that a scanner can open a hidden mind like a book. It showed that the failure of one response route could conceal capacities detectable through another. That distinction matters to anyone assessing a person whose ability to communicate has been damaged. Silence may be a fact about an impaired channel, rather than the whole person.

The problem is not confined to rare neurological cases. Every general anaesthetic asks medicine to alter consciousness while managing the body's other functions. Every assessment of animal suffering must consider evidence from a creature that cannot provide our kind of report. Every claim about a machine's feelings asks whether familiar words justify familiar conclusions when they are produced by unfamiliar mechanisms. The science cannot make the associated decisions disappear, but it can improve what those decisions rest on.

There is also a less practical reason to care. The thing being explained is the presence of a world at all. Light entering an eye and neurons responding to it are physical events. A colour appearing in experience is part of what we are trying to understand through those events. There need be nothing exotic about the example. A patch of blue or a persistent itch contains the difficulty. Intelligence, eloquence and self-reflection are unnecessary for the basic question to arise.

Ordinary familiarity can conceal the work. A hand feels like your hand

without an argument. A face is seen as a face rather than as separate measurements of edges and contrast. Fear can give a situation an urgency that a description of the situation alone does not possess. Some of these features can be investigated by changing the relationship among sensory signals, some by changing brain activity, and some by examining what a person can do when a particular capacity is lost.

The fascination is that the parts do not always move together. Experience can be vivid and mistaken. Information can be useful without being consciously acknowledged. A memory can be missing even when the event it would have recorded occurred. These are not reasons to distrust everything the mind presents. They are openings into how its achievements are produced, and reasons to avoid treating the feeling of effortlessness as an explanation.

The disagreements are worth understanding because the obvious replies often miss their target. Telling someone that pain involves neurons does not yet explain why it hurts. Accusing an illusionist of pretending not to feel pain misses what that philosopher is disputing. A person can take neuroscience seriously and still find the hard problem pressing. Getting these arguments straight lets you disagree with an idea rather than with its least convincing impersonation.

This book cannot supply an agreed final theory. It can supply a model of the subject that makes the disagreement intelligible and the experiments worth following. There is enough established knowledge to reject a surprising number of confident mistakes, and enough unresolved work to make the remaining questions substantial. The woman in the scanner gives that work a concrete demand. Before judging the presence of another mind, we need to understand how a mind can be present without the answer we expected.

The Core Ideas

Experience Is the Thing to Explain

You can point to a red cup, measure the light reflected from it and record activity in a visual system. None of those measurements is the redness as it appears to you. The appearance belongs to the event being investigated. Explaining how a system sorts wavelengths is one achievement; explaining its experience of colour is the ambition that makes consciousness a distinctive subject.

Consciousness, in the sense used here, is having experience. Seeing, hearing, bodily feeling, emotion and conscious thought all qualify. The experience need not be articulate, accurate or impressive. A confused dream counts. So does a dull ache. Reflecting on your own existence is one possible content of consciousness, not an entrance examination that every conscious creature must pass.

Thomas Nagel sharpened the issue in 1974 by asking what it is like to be a bat. Learning how echolocation works can tell us a great deal about the animal's relation to its surroundings. Imagining ourselves flying in the dark is another matter: that exercise starts with human capacities and changes the setting. The bat's sensory world need not resemble the picture we invent. The distinction is between knowledge about an experiencing creature and occupying its point of view, not between knowledge and complete ignorance.

The felt character of experience is often called its phenomenal character, or described in terms of qualia. These words identify a target; they need not commit us to little private objects floating inside a brain. A toothache is something undergone. A description of it can help someone else recognise, predict or treat it without giving that person the toothache. There is nothing scientifically disreputable about that difference.

In 1995 David Chalmers named the hard problem. Suppose a theory explained every discrimination, remembered detail, spoken report and

adjustment of attention. Why would carrying out those functions involve experience? Chalmers contrasted this with problems about how the functions operate. He called those problems easy in a technical sense: an account of the mechanism could, in principle, explain the specified performance. Solving them may take generations. The contrast concerns what would count as an answer, not how quickly a laboratory can obtain one.

Frank Jackson's imaginary scientist Mary makes the pressure easier to feel. By stipulation, she knows every physical fact about colour vision but has never seen colour. When she first sees red, does she learn something? The apparent answer is yes. The argument asks whether her new knowledge shows that the physical account left something out. Mary is not waiting for a better textbook: the thought experiment grants her all the physical knowledge there is. That extravagant condition is what gives the argument its force.

A physicalist has several replies. Mary might acquire a new ability to recognise and imagine red, or a new way of knowing a physical fact she already understood descriptively. The experience and its neural basis might be one event encountered in different ways. On that view, expecting a description to supply the experience confuses explaining a state with inducing it. Opponents answer that even perfect knowledge of the machinery still appears to leave its felt character unexplained. Neither reaction by itself settles the argument.

The first discipline is to preserve the question without deciding the result by definition. Human experience depends closely on the working brain. That leaves room to debate what kind of explanation this dependence supports. An explanatory gap may reveal something missing from our science, something misleading in our concepts, or a deeper feature of nature. Naming the gap does not tell us which.

One Word Hides Several Questions

A dreamer, a person solving a crossword and a patient unable to move all challenge the word conscious in different ways. The dreamer can have a world without responding to this one. The solver can use a fact deliberately. The patient's capacities may be difficult to establish. Placing everyone on a single scale from less to more conscious would discard much of what we need to understand.

Four questions help keep these cases straight. Is the dreamer experiencing anything at all? That is presence. What appears in the dream? That is content. Can information be used to reason, guide voluntary action or give a report? That is access. Why should the relevant brain activity involve experience in the first place? That is explanation. The questions can be distinguished without assuming that the brain contains four separate compartments to answer them.

Presence and content are easy to separate in ordinary examples. Being awake is not the same issue as seeing a particular face. A treatment could change which colours someone discriminates without removing the capacity to experience anything. Conversely, an intervention might prevent a conscious state while leaving some sensory responses intact. A useful theory must say which change it explains.

Access concerns what information can do within a system. A consciously noticed appointment can interrupt a plan, prompt a search for keys and become something you tell another person. Much information influences behaviour without being available in that way. The philosophical distinction between phenomenal consciousness and access consciousness, associated especially with Ned Block, asks whether experience may exceed what is available for such use. The distinction is influential; the extent of any separation remains disputed.

Wakefulness introduces a further issue. Arousal systems help maintain the conditions for an alert state, while the contents of that state concern what is experienced. Open eyes are useful evidence of wakefulness, but eyes can

open without reliable behavioural signs of awareness after severe brain injury. Dreaming supplies the converse warning: an unresponsive sleeper may have experiences. Arousal and awareness are related dimensions, not interchangeable labels.

Attention selects or prioritises information. Memory stores or reconstructs it. Metacognition concerns monitoring a judgement, including how confident you are that it is correct. These processes often accompany conscious access, but identifying one does not automatically identify the others. You can be confident and wrong, attend without clearly seeing, or have an experience that you later fail to remember. Each distinction explains a different kind of error.

Self-consciousness also has several meanings. Recognising yourself in a mirror, thinking about your reputation and feeling located inside a body are different achievements. Intelligence concerns capacities such as learning and solving problems. Sentience commonly emphasises the capacity for experiences with positive or negative feeling. A creature need not be a good problem-solver for something to hurt, and impressive problem-solving is not itself a demonstration of hurt or pleasure.

The distinction becomes useful when the same apparent failure has different causes. A sleeping person may miss a spoken word while dreaming of something else. A waking reader may miss it because attention is elsewhere. An injured patient may understand it without producing an answer. The missing response is common to all three; the conscious situation need not be. A theory may eventually explain several of these capacities through one mechanism. Separating the questions now prevents an observation about one from prematurely answering all the others.

The Mind Comes Apart Under Pressure

In a published study, a man known as TN walked around obstacles in a corridor despite damage to primary visual cortex, an early cortical stage of visual processing, in both hemispheres and the loss of ordinary conscious vision. The striking fact was not that he had acquired a new sense. Visual information could still guide behaviour through surviving capacities that did not provide his usual experience of seeing. His performance made the distance between successful action and acknowledged perception difficult to ignore.

Blindsight is the wider phenomenon. Some people with damage to primary visual cortex can locate or discriminate stimuli in an affected part of the visual field better than chance, despite reporting no normal visual experience there. Cases vary: some include vague feelings rather than complete absence of awareness. That variation matters. The general lesson concerns a dissociation between visual processing, performance and ordinary seeing, not an unlimited unconscious intelligence waiting behind everyone's damaged visual system.

The familiar gorilla experiment breaks a different assumption. Daniel Simons and Christopher Chabris asked observers to count basketball passes; many failed to notice an unexpected person crossing the scene. Attention to the assigned task helped them do one thing while leaving another unreported. The result is about selection under particular demands. It does not establish that people usually see nothing outside a narrow spotlight, or that an absent later report can settle every question about what was experienced.

Brain surgery can separate routes that normally cooperate. In some people treated for severe epilepsy, cutting the corpus callosum, the large fibre bridge between the cerebral hemispheres, greatly restricts communication between them. In classic split-brain tasks, information briefly presented to one side of the visual field could guide a hand yet be unavailable to the systems supporting a spoken answer. Each eye contributes to both visual fields; the experiment does not assign one whole eye to each hemisphere.

The dissociation shows why a spoken answer is not a complete inventory of a brain's capacities. It also leaves a question. Has the surgery divided experience itself, or mainly the routes by which information can be compared and expressed? Different tasks and patients support a less tidy picture than two independent people occupying one skull. The operation creates a powerful test of integration, but it does not provide an agreed method for counting subjects.

Dreams take the separation into ordinary life. You may experience fear in a place that is not the bedroom while showing little response to the room. Later, you may remember the dream badly or not at all. Experience, connection to the environment and later recollection are three distinct things. Their separation is not a rare exception that a theory can leave for its final footnote.

A theory built only around ordinary waking conversation could make speaking, remembering and experiencing look like one package. Dreams and damaged brains undo the packaging. They reveal which operations can come apart, so a theory has to explain both their normal coordination and their independent failures. A familiar sign of consciousness remains useful; it ceases to be a definition.

Being Someone Is Part of the Construction

The hand in front of a participant was made of rubber. Their own hand was hidden. In Matthew Botvinick and Jonathan Cohen's 1998 experiment, matched stroking of the visible artificial hand and the unseen real one could make touch seem located at the rubber hand and alter bodily self-identification. Changing the timing weakened the illusion. What normally feels beyond negotiation, which hand is mine, proved responsive to the relationship among sensory signals.

A familiar conclusion goes too far: the self is therefore unreal. The experiment showed something narrower and more informative. Bodily ownership depends on how vision, touch and information about body position are combined. A system that ordinarily places sensations on the

correct body can be misled under arranged conditions. Its susceptibility explains something about how it works; it does not turn every successful result into a fiction.

This makes selfhood a proper part of consciousness research. Experience usually arrives from somewhere. The world is seen from a viewpoint, touches are located on a body, and actions often feel as though they are being performed by that body. These features can be distinguished from the autobiographical self: the history, commitments and social identity you can describe. Remembering your school and feeling the position of your left hand are both about you, but they are not the same mental operation.

Some of the information comes from inside. Interoception concerns sensing the body's internal condition. Breathlessness, fullness, nausea and the bodily changes accompanying emotion are not ornaments added to an otherwise visual mind. They are ways an organism's condition can become part of experience. A theory illustrated entirely by coloured squares must eventually account for the urgent difference between looking at a square and needing air.

Predictive approaches, including Anil Seth's account of interoceptive inference, propose that the brain estimates the causes of its signals rather than passively displaying them. In this picture, perception depends on expectations meeting incoming evidence. The importance assigned to a mismatch depends partly on how reliable the signal is expected to be. For bodily experience, the system also acts to regulate the conditions it is estimating. An organism has to keep itself alive, not merely produce an accurate internal diagram.

Imagine a pounding heartbeat while waiting to give a speech, then while running for a train. A similar bodily signal need not mean the same thing in both situations. In the predictive account, context helps shape its interpretation; the feeling is no less real for being interpreted. Predictive accounts explain how experience might be organised; prediction alone is not an agreed test of consciousness.

Feeling also has direction. An experience can be pleasant, unpleasant or comparatively neutral. Philosophers and researchers call this valence. The negative character of pain is distinct from detecting tissue damage, just as enjoying music is distinct from identifying its pitch. This matters for other animals: avoiding a stimulus and undergoing an aversive experience are related possibilities that require different evidence. It matters for machines for the same reason. A numerical penalty in a learning algorithm is not, by its name alone, suffering.

Putting a tiny spectator behind the eyes would only postpone the explanation. How does that spectator see? Giving it another internal display repeats the problem on a smaller scale. The work has to be done by the system, not delegated to another unexplained mind inside it. Nor must selfhood arrive as one complete package: a creature could have bodily experience without a verbal life story. The body and the self enlarge the target. Consciousness is a world encountered from a position, often with something at stake for the one encountering it.

The Brain Needs Organisation, Not Just Activity

A brain that stops responding to a command has not necessarily stopped processing its sound. Under some anaesthetic conditions, sensory responses remain measurable while the wider organisation of activity changes. Counting responses therefore misses the central question. What can those responses influence, for how long, and within what state of the rest of the brain?

Neurons affect other neurons through connections whose consequences depend on timing, inhibition, excitation and ongoing activity. A burst can remain local, provoke a stereotyped response or participate in changing patterns across a network. These are different physical arrangements even when all contain plenty of electrical activity. Consciousness research asks which arrangements matter for the state as a whole and which support particular contents.

The cerebral cortex, the folded outer sheet of the brain's hemispheres, is

important, but it does not operate alone. The brainstem, connecting the brain with the spinal cord, and other structures beneath the cortex help regulate arousal. The thalamus, deep within the brain, participates in extensive interactions with cortex. These systems help establish the conditions under which particular contents can occur. Explaining why a brain remains capable of experience is different from explaining why, at this moment, a face rather than a sound becomes its content. The two explanations have to fit together, but they need not point to exactly the same circuitry.

Intervention supplies a stronger test than a correlation. In a 2020 study, stimulation of a central thalamic region in anaesthetised macaques restored wake-like behaviour and changed cortical activity. This implicated a particular circuit in controlling arousal under those conditions. It did not identify a universal switch or give researchers direct access to the animals' experience. The achievement was causal and specific: changing this part changed how the larger system operated.

Another approach perturbs the cortex and examines the response. A transcranial magnetic stimulation pulse induces activity, and electrodes on the scalp record the subsequent electrical pattern. The perturbational complexity index, PCI, developed by Adenauer Casali and colleagues, measures how differentiated and distributed that response is using a compression-based calculation. A response that repeats everywhere can be described economically: the same thing happened again. A more varied, less repetitive response across places and time needs a longer description.

In the conditions studied, conscious wakefulness tended to permit a richer response than states associated with no reported experience. Later validation included cases in which people were unresponsive yet subsequently reported experiences, such as dreaming or ketamine-related states. The value of the method is that a person need not understand a command or move a hand for the brain to be probed. That removes some barriers faced by behavioural tests, although interpretation still depends on validation and the condition being examined.

PCI does not measure the richness of someone's thoughts, and it is not the formal quantity proposed by integrated information theory. A useful clinical indicator can survive disagreement over the theory that helped inspire it. Equally, a successful distinction between selected states does not guarantee accurate classification of every injured brain. The measurement gains its authority from what it predicts under testing, not from sounding close to a philosophical definition.

The emerging lesson is about organisation. Signals must do more than exist, and more activity is not automatically more experience. Which forms of interaction are essential remains contested. The brain's dependence on its physical organisation is much firmer ground than any particular claim to have found the one pattern that explains consciousness completely.

Rival Theories Propose Different Work

Imagine a laboratory observer briefly shown a face. Some visual processing occurs even on trials when the observer says they saw nothing. On other trials the face becomes clear, can be compared with a remembered person and can guide a deliberate answer. Theories disagree about the operation that makes the difference. They are not competing names for the fact that neurons fire.

Global neuronal workspace theory proposes that information initially processed within specialised systems becomes conscious when it enters a wider, recurrent network. Recurrent means that influence runs back through connections rather than travelling once in a forward chain. A sufficiently strong representation produces an amplified, sustained episode often called ignition. Its availability to otherwise partly separate capacities explains why the noticed face can affect reasoning, working memory and voluntary action together.

The workspace is not a room containing a viewer. It is a proposed organisation of information exchange. Nor does global mean that every neuron receives the message. The claim concerns broad availability across relevant systems. Proponents aim to explain conscious processing itself

through this organisation; critics question whether the commonly measured signatures identify experience or the additional work required to use and report it. Restricting the theory to a theory of talking would remove its central claim before testing it.

Recurrent-processing theory places the decisive event earlier and more locally. A forward sweep can classify features without experience. Feedback within sensory systems, on this account, establishes the organised percept. Wider access enables reasoning and report but need not create the experience in the first place. For the briefly shown face, the crucial question is whether it was experienced before, or without, becoming globally available. The challenge is to establish that distinction without treating every failed report as evidence for hidden experience.

Higher-order theories ask what makes a mental state something the system is aware of having. A first-order state represents the face. A higher-order representation concerns the system's own state, in roughly the way that a confidence estimate concerns a judgement rather than the external object alone. Different versions specify that relationship differently. The higher-order representation need not itself be conscious, verbal or a deliberate act of introspection. There is no requirement to keep forming another conscious thought about the previous conscious thought indefinitely.

This family makes self-monitoring relevant without requiring a philosophical inner monologue. It also faces a demanding question: does evidence about confidence and the monitoring of perception reveal the mechanism of experience, or a capacity operating alongside it? Someone can be unsure what they saw while still experiencing something. A theory must explain the relation, not identify conscious seeing with accurate self-assessment.

Integrated information theory begins from features its proponents take to characterise experience: it is specific, structured and unified for the subject. IIT then proposes corresponding requirements for a physical system's intrinsic causal organisation. Its central claim is stronger than the observation that connected brains are conscious. It identifies experience

with a particular irreducible cause-and-effect structure, considered from the system's own causal perspective.

The unusual word intrinsic matters. A database can tell its user a great deal, but usefulness to a user is not IIT's measure. The question is what causal differences the physical components make to one another, and what would be lost if that causal whole were divided. The theory's formal requirements select particular wholes rather than assigning one big mind to any collection of interacting objects. Applying the formalism to real brains is difficult, and both its starting assumptions and its empirical consequences are disputed.

The contrast with functional approaches becomes sharp for machines. Two devices might generate the same answers while having different internal causal structures. IIT allows that their consciousness could differ. A functionalist instead looks for the organisation of causal roles that matters to mentality, potentially across different materials. Neither position licenses judging a device from its conversational output alone; each must specify what lies behind that output.

Predictive approaches address a partly different task: how sensory evidence and bodily regulation shape what is experienced. They can be combined with other accounts rather than always competing as a fifth mutually exclusive entry. A prediction machine still needs an account of why some of its activity is conscious. In turn, a theory of global availability needs to explain why the available content has the character it does.

The theories therefore overlap, diverge and sometimes explain different things. Their names are less useful than the distinctions they force: local versus widespread interaction, representation of the world versus representation of one's own state, functional organisation versus intrinsic causal structure. A convincing theory must turn those distinctions into predictions that its rivals cannot all claim after the result arrives.

Other Minds Turn Knowledge into a Decision

A person can tell you that something hurts. An animal cannot give the same kind of report, and an artificial system can generate the sentence without settling whether it feels anything. The tempting shortcuts point in opposite directions: demand human speech from everything, or attribute experience wherever behaviour resembles ours. Both let the test's convenience determine the answer.

For other animals, the strongest case comes from converging evidence. Neural organisation, evolutionary relationships and flexible responses can reinforce one another. Withdrawal from injury is relevant but insufficient on its own, because protective processing need not be felt. Changes in learning or choices that reveal the continuing negative value of an event add a different kind of information. No single human-style performance should stand in for the entire question.

Robyn Crook's 2021 octopus research illustrates the distinction. Under the study's conditions, animals avoided a place associated with an injurious injection. Those given the injection also preferred a place associated with local pain relief, whereas uninjured animals did not show the same preference. Together with physiological evidence, this supported an ongoing aversive state rather than only an immediate withdrawal reflex. It was evidence about the tested animals and conditions, not a demonstration that every invertebrate experiences human-like pain.

An expert declaration issued in New York in 2024 judged the evidence strong for consciousness in other mammals and birds, and a realistic possibility across a wider range of animals, including all vertebrates and many invertebrates. A declaration records a judgement by its signatories; it is not a new experimental result. The welfare argument attached to it is that a credible possibility of experience should count in decisions, even where certainty is unavailable.

Artificial systems lack some of the shared biological background that makes inference between humans so strong. Their behaviour can also be selected

directly to resemble human speech. That does not establish unconsciousness. It makes the route from a convincing sentence to a conclusion about experience less secure. A system's training, architecture and operation matter alongside its answers.

Patrick Butlin and colleagues have developed a framework of possible indicators drawn from consciousness theories, including recurrent processing, global availability and higher-order monitoring. This is an organised way to investigate systems, not a validated diagnostic test. Each indicator inherits uncertainty from its parent theory, and superficial imitation of a function may differ from implementing the proposed mechanism. The useful question is what the system does internally, not how many theoretical labels can be attached to it.

Belief and action then separate. How likely is experience, given the evidence? What should we do, given that possibility and the consequences of error? These are related questions with different inputs. A modest chance of serious harm may justify an inexpensive precaution without justifying a declaration of certainty. Conversely, a weak resemblance does not require treating every responsive device as a suffering subject.

The condition that began the enquiry now has a practical consequence. Because another's experience is not available in the same way as your own, good judgement depends on the reliability of the routes by which it becomes known. Science can strengthen those routes and reveal where they fail. What it cannot responsibly do is turn a convenient absence of response into permission to stop asking.

How It Actually Works

One eye receives one image; the other receives a different one. Instead of maintaining a simple blend, the observer's experience can alternate between them. Nothing comparable need change on the displays. Binocular rivalry gives consciousness research an unusually useful arrangement: the physical stimulation can stay approximately constant while what the person sees changes.

Making experience vary

The experiment begins with a separation that ordinary viewing rarely provides. Outside the laboratory, a change in seeing usually follows a change in the world, the eyes or the observer's activity. If someone notices a passing car, increased brain activity could reflect the car's arrival rather than the fact that it became conscious. Rivalry reduces that difficulty.

Researchers can look for activity associated with a change in the experienced image while holding much of the sensory input steady.

The obvious way to mark each change is to ask the observer to press a button. That creates another problem. Pressing requires deciding what is currently visible, keeping track of the alternatives and preparing movement. Activity associated with the button may be mistaken for activity necessary to the percept. The experiment has separated experience from one confound while introducing others.

No-report methods try to reduce this burden. In some rivalry displays, eye movements or pupil changes vary with the dominant percept. Once their relationship with reported perception has been established, they can provide a less intrusive indication of changes. Stefan Frässle and colleagues used such measures in a 2014 study and found that some frontal activity associated with active reporting was reduced when report was removed. That challenged a straightforward interpretation of the original activity. It did not establish that frontal systems never contribute to consciousness.

Visual masking offers a complementary manipulation. A brief target is

followed or surrounded by another pattern that makes it difficult to see. Researchers vary the timing, strength or arrangement, then ask about the target and measure behaviour. An unseen or poorly seen target may still influence a later response. The important contrast is not between no processing and processing. It is between different consequences of processing the same kind of information.

Even a simple question needs care. A person who declines to say they saw a faint stimulus may be cautious rather than devoid of visual evidence. Forced-choice discrimination can reveal information that a yes-or-no visibility answer misses, while a graded visibility scale can separate a vague impression from clear identification. Confidence provides another measurement, but it also introduces judgement about judgement. None is automatically the privileged report of experience in its pure form.

Consider a hypothetical test of hearing a faint tone. One listener says yes whenever there might have been a sound; another answers yes only when certain. Counting their yes answers would mix sensitivity with willingness to commit. Trials containing no tone help reveal false alarms, while comparisons among responses help separate the two issues. Consciousness experiments face a related difficulty when a change in reported visibility is treated as a change in experience. The task must make it possible to distinguish better perception from a different threshold for saying that something was perceived.

A strong experiment therefore changes one relevant condition, checks what else changed and uses more than one kind of measurement where possible. This is how a private phenomenon becomes scientifically tractable. The observer does not hand over an experience. They provide behaviour whose relation to experience can be probed, challenged and compared with physiology.

From a brain image to a mechanism

Francis Crick and Christof Koch helped direct modern research towards the neural correlates of consciousness in 1990. The programme made a strategic distinction: identifying the neural mechanisms associated with conscious experience could proceed before agreement about its ultimate explanation. Finding reliable relationships would constrain the philosophical argument rather than wait for it to end.

The target is often described more precisely as the minimal neural mechanisms sufficient for a particular experience, given suitable background conditions. In practice, researchers begin with candidates that covary with what people report. A candidate might belong to the experience itself, to a prerequisite such as arousal, or to a consequence such as remembering and describing it. The term correlate names a research target, not a licence to collapse these possibilities.

To follow that fleeting face percept, researchers need to know both where activity changes and when. Electroencephalography records electrical activity from the scalp with fine timing but difficult source localisation. Magnetoencephalography records magnetic fields associated with neural activity. Functional magnetic resonance imaging helps locate patterns across the brain, but uses a slower, indirect signal: changes in blood oxygenation. Intracranial electrodes record close to neural tissue, usually in patients implanted for clinical reasons. They offer proximity at the cost of coverage, since they sample only the places where electrodes have been put.

No instrument supplies every advantage at once. A rapid electrical change might be well timed but hard to locate; a spatial pattern might be clear while its sequence is uncertain. Agreement across methods is useful because their limitations differ. Disagreement can also be informative when it exposes the exact scale or operation on which a claim depends.

Recording still leaves causation unresolved. Stimulation, damage and controlled changes in brain state allow stronger questions. Does disrupting a

process remove a capacity? Does inducing a pattern change an experience? Does the effect persist after controlling for a damaged sensory route or a lost ability to respond? Necessary and sufficient are demanding words. A mechanism can be required in one route to a percept without being the only route any brain could use.

Putting theories at risk

In the Cogitate collaboration, researchers and proponents of integrated information theory and global neuronal workspace theory specified contrasting predictions before testing them. The study, published in 2025, involved 256 participants across functional imaging, magnetoencephalography and intracranial recordings. Participants viewed visible images presented for different durations. The experiment asked where and how the contents and timing of experience were reflected in brain activity.

To find where image information is represented, researchers can train an analysis system on recordings made while known images are being viewed, then ask it to identify images from other recordings. If it succeeds, the measured signal contains information the system can use. This is decoding, not a demonstration that the sampled region produces experience on its own. Conversely, failure may reflect the recording or analysis rather than an absence of information. A theory about where a representation must be available needs to survive that distinction.

The results supported elements of both accounts while challenging agreed predictions. Among the difficulties were the absence of predicted sustained synchronisation between posterior areas for IIT and limited evidence for predicted prefrontal representation of some stimulus features and for an ignition response when an image disappeared. The published conclusion was not that one theory had defeated the other. Particular commitments had encountered data they did not explain as expected.

A response published later in 2025 by Lionel Naccache and other workspace proponents disputed how much this threatened their framework.

They argued, among other points, that a response to an image disappearing depends on whether its disappearance is attended and noticed. They also stressed that presenting only visible stimuli does not directly test the contrast between conscious and unconscious processing. The disagreement concerned the connection between a broad theory, a specific prediction and the experimental task.

The next step follows from the objection. If noticing the disappearance is required for offset ignition, an experiment can vary attention to that disappearance and ask whether the predicted response follows. Such a test would turn a retrospective explanation into a prospective risk. Agreeing on predictions in advance remains valuable even when their theoretical importance is later contested: it preserves a record of what the participants expected before seeing the answer.

The study's images also mark its boundary. Successful decoding of faces or objects concerns those contents in the tested human tasks. It does not yet explain nausea, bodily ownership or the experience of an unfamiliar animal. The value of the collaboration lies in making several large theoretical claims answerable to particular observations, with disagreements precise enough to guide the next comparison.

What would count as an explanation?

Suppose the neural work went perfectly. You could predict what someone would see, alter the experience by changing the brain, and explain why they reported it. Would that finish the job? A physicalist holds that conscious life belongs wholly within physical reality. Some expect neuroscience to identify experiences with physical processes; others emphasise the functional organisation those processes realise. Both can accept that today's explanation is incomplete without concluding that a non-physical ingredient must be added. The dispute concerns what a completed explanation would look like.

The question of function crosses this divide. Broad access to information could help an animal compare competing goals and use a remembered

situation to choose an unfamiliar action. Those are potential benefits of mechanisms that some theories identify with consciousness. Whether feeling supplies an additional evolutionary advantage depends on whether it makes sense to separate feeling from those mechanisms in the first place. A functionalist may regard them as one phenomenon, not two traits needing independent explanations.

Dualist approaches deny that the physical description exhausts the mental. Substance dualism treats mind and matter as different kinds of thing. Other versions propose irreducible mental properties rather than a separable soul. Their attraction is that they take the apparent difference between experience and objective description seriously. Their difficulty is explaining how those properties fit the reliable causal relations between brains, behaviour and the physical world.

Panpsychism changes the starting point. Instead of beginning with matter that has no experiential aspect and asking how experience appears, it treats some basic experiential aspect as fundamental. This need not mean that a rock thinks or that an electron has human emotions. It creates a different problem: how could those basic features compose, or relate to, the unified experience of a person? Moving experience into the foundations changes what must be explained; it does not complete the explanation.

Illusionism, defended by Keith Frankish, challenges the way the target is commonly described. It proposes that introspection misrepresents our mental activity as possessing phenomenal properties at all, rather than denying only that those properties are mysterious or non-physical. The work is then to explain the discriminations, reactions and self-reports, including why we are convinced that something further remains. Calling this an illusion does not mean that pain behaviour is pretence or that suffering may be ignored. It is a claim about what our access to our own states reveals.

A further thought experiment exposes the disagreement. A philosophical zombie is stipulated to duplicate a human's physical organisation and behaviour while having no experience. It would make the same reports, including insisting that pain hurts. Chalmers uses the apparent

conceivability of such a duplicate in an argument against physicalism. A physicalist can reject the move from an apparently coherent description to a real possibility: perhaps specifying the complete physical duplicate already specifies the experience, whether we recognise that or not. No laboratory has manufactured the stipulated comparison. Its role is to test an argument about what follows from a complete description.

These positions disagree about explanation, not just anatomy. Finding a reliable neural signature will not by itself choose among all of them. Yet they are not insulated from scientific progress: accounts of self-monitoring, integration and report can make one interpretation more plausible or expose its costs. The useful division is between what an experiment establishes and what a theory adds to explain it.

Losing contact without losing a world

A sleep laboratory can turn a later recollection into a more closely timed observation. Wake sleepers repeatedly and ask what, if anything, they were experiencing just before the awakening. Record brain activity beforehand. The report remains dependent on memory, but the delay is shortened and the preceding state is known more precisely than it is in a story told at breakfast.

Francesca Siclari and colleagues used this approach to study reports of experience in both rapid-eye-movement and non-rapid-eye-movement sleep. Their 2017 findings linked reported dreaming with local changes in posterior cortical activity. The crucial lesson is not a universal dream location. It is that the broad label sleep, and even the distinction between its major stages, does not settle whether experience was occurring on a particular occasion.

General anaesthesia changes the conditions pharmacologically. Anaesthetic agents affect different molecular targets and circuits; they are not interchangeable ways of pressing the same button. Depending on drug, dose and circumstances, responsiveness, memory, connection to the environment and experience can separate. Unresponsiveness may

accompany absence of experience, or experiences disconnected from the room. Later amnesia adds another possibility: an event occurred but left no recoverable record.

Movement is particularly misleading when considered alone. A neuromuscular blocking drug prevents muscles from responding; it does not by itself remove experience or provide pain relief. In clinical anaesthesia, such drugs are part of a wider plan. Their scientific lesson is stark: eliminating a route for expression cannot be treated as eliminating the state that might have been expressed.

Conversely, detecting a brain response to sound during an unresponsive state does not demonstrate that the person heard it consciously. Some sensory processing can persist without a reportable percept. Researchers therefore distinguish detecting stimulation, being connected to the surroundings, having any experience and later remembering it. Anaesthesia is valuable to the science because these conditions can change in controlled and partly reversible ways, not because unconsciousness is one uniform state created by every drug.

Asking for an answer without movement

At the bedside after severe brain injury, the first challenge is careful observation. Coma involves an absence of both wakefulness and behavioural evidence of awareness. Unresponsive wakefulness syndrome, also called the vegetative state, includes wakeful features without reliable behavioural signs of awareness. A minimally conscious state includes limited but definite evidence of awareness. These are clinical categories based on the evidence available, not direct readings of a private scene.

The Coma Recovery Scale-Revised structures the examination across auditory, visual and motor responses, oromotor and verbal function, communication and arousal. Repeated assessment matters because performance can fluctuate. A command may fail because of impaired hearing, language, attention, movement or the effects of medication. Locked-in syndrome makes the distinction especially clear: severe motor

restriction can coexist with preserved consciousness, although other impairments may also be present.

The mental-imagery approach used in the 2006 case tries to bypass one obstacle. Following an instruction to perform a particular mental task can produce a sustained, task-specific pattern rather than merely a brief response to hearing words. In a study published in 2010, one patient used different imagery tasks to indicate answers to questions. This demonstrated a possible communication route for selected individuals, not a general ability to read thoughts from every injured brain.

The larger six-centre study published by Yelena Bodien and colleagues in 2024 put numbers around the distinction. Among 241 participants without observable command-following, 60 produced responses to tasks detected by functional imaging, electroencephalography or both. That is the often-quoted 25 per cent. It was a convenience sample from specialist centres, assembled from data collected between 2006 and 2023, not a random survey of everyone described as unresponsive.

The second denominator is just as instructive. Of 112 participants who could follow commands observably, only 43 produced a detected response on the brain-based tasks. Those tests therefore missed many people whose bedside behaviour already supplied evidence of awareness. A negative result could not safely carry the same weight as a well-controlled positive one. The difficulty was not philosophical doubt about whether other people have minds; it was demonstrated failure of a demanding test to register known capacity.

A patient-level meta-analysis published in July 2026 brought together evidence from 56 studies identified in a search ending on 1 May of that year. Detection varied with clinical state and cause of injury. The combined work strengthens the case for covert command-following as a real clinical phenomenon, while the differences among patients and methods prevent its prevalence from becoming one universal percentage. A newer synthesis is not a newer census of the same population.

These findings do not supply an individual prognosis or prove that someone has intact memory, reasoning or quality of life. They show why assessment should look for more than one route of response. Recovering a channel of communication and establishing the full life available through that channel remain different tasks.

How we know

The most persuasive evidence combines changes in reported experience with behaviour, neural recording and interventions. Each approach constrains a different alternative: whether a stimulus was processed, whether a report was possible, whether a brain event was a cause or a consequence. Convergence matters more than the visual appeal of a scan.

Much experimental work concerns adult humans, visual stimuli and tasks designed for a laboratory. Extending its conclusions to emotion, infants, unfamiliar animals or artificial systems requires more than retaining the same technical vocabulary. Clinical studies add important cases but face damaged sensory and motor routes, fluctuating states and incomplete assessment. Their participants are seldom representative of every patient to whom a headline refers.

Reports are evidence, not perfect transcripts; absence of report is evidence whose force depends on the route's reliability. A neural measure can improve that evidence without becoming a direct sample of experience. Theories then organise the findings and propose explanations. Their empirical predictions can be tested even where their philosophical interpretation remains disputed. The result is substantial knowledge with identifiable limits, rather than a choice between certainty and knowing nothing.

What People Get Wrong

“Awake means conscious, asleep means unconscious”

The everyday distinction works well enough to ask whether someone is ready for breakfast. It is a poor guide to whether they are experiencing anything. A sleeper may be in the middle of a dream; a person with severe brain injury may have open eyes without reliable signs of awareness. Wakefulness and experience usually travel together, which explains the shortcut. The exceptions show why the shortcut cannot serve as a definition.

Dreaming is not confined to one sleep stage. Conversely, the wakeful features of unresponsive wakefulness syndrome do not establish awareness. Neither finding licenses a simple reversal in which every sleeper is experiencing something and every unresponsive waking patient is conscious. The correction is to keep arousal, connection to the environment and experience separate long enough to ask which one has been measured.

This changes the questions asked of an experiment. A signal that distinguishes sleep from wakefulness might identify the physiology of those states without detecting every episode of experience. A finding about whether somebody responds to the room may miss a dream. The broad label is useful, but it is not the answer to every more precise question beneath it.

“No answer means nobody is there”

An answer depends on more than having something to say. It requires a usable route from the stimulus through understanding and response. Ordinary conversation hides this dependence because its components usually work. Severe injury can make the dependence decisive. Hearing, attention, language comprehension or movement may fail before a clinician has learned what experience is present.

The reverse mistake is to treat any brain response as a hidden answer. A response to a sound could reflect processing that does not involve understanding a command. This is why sustained, instruction-specific modulation carries more evidential weight than a general increase in activity after speech. A test must distinguish deliberate participation from the alternatives its design makes possible.

The correction changes how silence is interpreted. A negative result is stronger when the test reliably detects the relevant capacity in comparable people who are known to possess it. When that reliability is poor, silence leaves more possibilities open. It warrants neither a declaration of complete absence nor an assurance of intact consciousness. The limitation belongs to the inference, not to the value of the person being examined.

“Consciousness requires human-level intelligence”

Language makes consciousness unusually easy to discuss in humans. That convenience encourages a mistaken definition: the creatures able to describe themselves must be the creatures able to experience anything. Under that rule, absence of an elaborate report would defeat the case before evidence about feeling had been considered.

Intelligence and experience answer different questions. Solving a puzzle establishes a capacity to solve that kind of puzzle. An aversive experience establishes that something is bad for the subject undergoing it. Sophisticated learning can contribute evidence about experience, especially when choices reveal flexible responses to benefit and harm, but there is no agreed intelligence score above which feeling begins.

The same caution applies to development. Infants cannot supply adult verbal reports; research must use other evidence rather than define them out of the enquiry. Research identifies several potential markers of early experience well before language, although their interpretation and the timing of its onset remain disputed. They cannot be read directly from the first use of a word or a successful adult-style test.

For animals, the appropriate comparisons concern their sensory capacities, bodily organisation and ways of responding to their own environment. Requiring a creature to perform like a human can confuse failure to meet our task with absence of a mind. The correction is not that all living things feel. It is that human-like cleverness cannot be the sole admission ticket.

“Scientists have found the consciousness centre”

The headline is attractive because a place sounds like an explanation. Yet finding a region involved in seeing a face is different from finding what keeps a person capable of any experience. The first concerns content; the second concerns state. Combining them into one search for a consciousness centre hides that difference before the experiment starts.

Necessity asks whether a capacity is lost when a region's contribution is removed. Sufficiency asks whether the relevant activity can produce it under specified background conditions. The difference matters in an experiment. Damage might block incoming visual information, disrupt the activity supporting its experience, or prevent the person giving an answer. All can spoil performance, but only one was the proposed mechanism of experience. A damaged connection can be indispensable without being the place where everything happens.

A strong account therefore needs both location and operation. It should explain why a particular circuit contributes to a particular state or content, and how that contribution depends on the rest of the brain. This allows local specialisation without a single anatomical spectator. The useful question becomes what the region does in the system, rather than whether somebody has finally found the spot where you reside.

“A split brain contains two complete people”

The claim persists because classic experiments are striking: one route can guide a response that another cannot explain. It is tempting to picture a second person behind the inaccessible information. The experiment, however, establishes a division between capacities before it establishes a division between subjects.

Information integration, control of action, verbal report and the sense of self need not divide along an identical line. Operations differ in completeness, people differ in retained connections and patients can adapt. The findings therefore do not support treating every person who has undergone callosal surgery as the same clean experiment repeated. Contemporary research continues to dispute how unified their experience and agency remain.

There is no need to weaken the original discovery to avoid exaggerating it. Disconnecting major communication routes reveals that the apparently seamless mind depends on coordination that can be altered. The open question is how those changes map onto experience. Two isolated answers to a laboratory task are evidence about access; counting whole experiencing selves requires more. A memorable interpretation should not be allowed to outrun the experiment that made it memorable.

“The hard problem proves that the mind cannot be physical”

The argument begins with a genuine pressure: an account of what a system does can seem to leave its experience unexplained. It gains force from the immediacy of pain or colour compared with descriptions of electrical signals. The mistake is to turn that contrast into an established proof of a non-physical mind.

An explanatory gap and a gap in reality are different claims. The first concerns what we can currently explain or derive from a description. The second concerns what exists. Moving from one to the other requires an argument. Physicalists can accept that current neuroscience has not explained consciousness while denying that the failure demonstrates a separate mental ingredient.

There is an equally careless dismissal on the other side: declaring that consciousness is brain activity and treating the explanatory work as completed. A label does not specify which activity, why it matters or how the different characters of experience are accounted for. The dispute is substantial precisely because both easy escapes leave work undone.

The productive stance is to ask what a proposed explanation commits itself to. Does it identify experience with an operation, explain why we think something further is missing, or add a new basic property? Those are rival answers to examine. None becomes correct merely because another answer has difficulties.

“An AI that describes feelings has demonstrated them”

In a human conversation, a report of fear arrives with a background of shared biology, development and experience. A generated sentence can resemble that report while having a different history. Training may reward language that fits a context, including language about feelings. The words alone do not reveal whether the system has the state they describe.

This does not prove that artificial consciousness is impossible. It means behavioural resemblance needs interpretation. A useful investigation would ask what produced the report, how stable it is across contexts, what internal processes are involved and whether those processes satisfy a defensible theory. A prompted denial of consciousness is no more a final metaphysical authority than a prompted declaration.

The opposite shortcut, declaring that computation can never feel, also needs an argument about which physical or biological properties matter. Theory-based indicators are a way to organise that research, not a scorecard whose total establishes consciousness. Until a method is validated, precise percentages can give uncertainty an undeserved appearance of measurement. The appropriate conclusion may be unresolved, but it should be unresolved for stated reasons rather than decided by the most convincing sentence.

Use It

Suppose a headline announces that a scanner has detected consciousness in someone who cannot respond. The important question is what the person was asked to do. A brief response to sound, a sustained response to an imagery instruction and a complex response to magnetic stimulation establish different things. The distinctions in this book become useful when they change how you read that headline, rather than leaving you impressed or suspicious of every scan.

Name the claim before judging it

Replace the broad word conscious with the specific claim being made. Is there evidence of experience at all, of a particular content, of flexible access to information, of self-monitoring, or of a proposed explanation? An account can succeed on one question without answering the rest. It can also claim that two questions have the same answer, but that identity needs to be stated and defended.

Imagine an announcement that a system can identify which picture a volunteer is viewing. Everyone tested might have been awake throughout. The achievement would concern information about visual content in a brain signal, without yet testing how to distinguish a conscious state from an unconscious one. A content decoder and a detector of consciousness face different tests.

The habit protects useful science from inflated promises. Instead of dismissing the result because it has not solved everything, you can recognise what it explains. Instead of accepting the headline, you can locate the step that turns a narrow result into a much larger claim. The question becomes more informative precisely because it becomes smaller.

Examine the route by which an answer could arrive

When a response is absent, work backwards through what the task required. Was stimulation delivered in a usable form? Could it be detected and understood? Did success require sustained attention, a particular kind of imagery or a movement that may have been impaired? The missing answer could originate at several points, only some of which concern experience itself.

This does not require imagining limitless hidden capacities. It requires asking whether the test was capable of registering the capacity under the circumstances. The 2024 clinical study supplied an important check: many participants who followed commands visibly did not produce a detected response in its brain-based tasks. That observation changes the weight of a negative result more than a general assurance that the technology is sophisticated.

The same logic applies to unfamiliar animals. A task can be difficult because its demands are poorly matched to the species, rather than because the relevant experience is absent. For artificial systems, the difficulty may run the other way: producing the expected answer can be easy for reasons unrelated to the proposed inner state. Look at the whole route, not only the final response.

Follow the intervention, not the brightness

When a claim is based on a brain image, ask what changed and what was measured. A region can become more active because a stimulus arrived, because the person noticed it, because they decided what to say, or because they prepared the answer. A coloured patch in a figure does not identify which interpretation is correct. The comparison between conditions does the explanatory work.

An intervention can go further by changing the candidate mechanism. Even then, the relevant distinction is between disrupting experience and disrupting the means of demonstrating it. A task that requires speech will fail

if speech production fails. The failure does not alone show that the person no longer experiences the stimulus. A useful control asks whether another response route remains available.

This lens also helps with claims that an activity or treatment increases consciousness. More electrical activity, greater complexity or a larger response is a measurement, not yet an improvement in experience. Ask what the change predicts in the people and states tested. Keeping the measurable claim separate from the attractive interpretation makes it possible to evaluate either without accepting both.

Make the explanation risk being wrong

Ask two questions together: what does the theory predict, and what would its strongest rival predict instead? Finding feedback in a conscious brain does not by itself favour recurrent-processing theory over a workspace theory that also requires feedback. The decisive difference concerns which feedback is sufficient, where it occurs and whether wider availability is required. Shared predictions cannot decide between the accounts that share them.

For the competing accounts to part company, researchers would need a case in which local recurrence and wider access come apart. They would then have to establish what, if anything, was experienced. Asking for a report can recruit the processes whose necessity is disputed. That is why another bright brain image may settle less than a carefully designed task: the useful result has to distinguish the rivals, not merely resemble both.

The same discipline applies after an unexpected result. A revised theory may be better than the version that failed, but it must make another commitment. In the Cogitate dispute, an attention-dependent account of offset ignition calls for tests that vary attention to disappearance. An explanation earns confidence by making that next comparison possible, not by finding words that accommodate the previous result.

Separate your estimate from your precaution

Evidence about experience and a decision about protection are not the same calculation. The evidence determines how plausible a claim is. Action also depends on the possible harm, the burden of protection and the alternatives available. There is no need to pretend that uncertainty has vanished before deciding that a modest precaution is justified.

Take a hypothetical policy for an unfamiliar animal. Several independent findings suggest that injury may create an aversive experience, but researchers disagree about the interpretation. Avoiding an unnecessary harmful procedure could be reasonable without announcing that its consciousness has been proved. A much more burdensome measure would demand a different assessment of consequences. The principle is proportionality, not an instruction to maximise protection regardless of cost.

Artificial systems require the same separation, while also taking human consequences seriously. Caution about possible experience does not require trusting a system's emotional claims, granting its requests or ignoring the people affected by its deployment. Neither casual dismissal nor automatic personification is a substitute for evidence. The aim is to make the reason for a decision visible, including the uncertainty it is designed to manage.

The limits

These lenses do not diagnose a patient or settle a family decision about treatment. Clinical assessment combines history, repeated examination and specialist interpretation; a general book cannot supply those particulars. Nor does it offer a consciousness test for a pet or a chatbot. Its contribution is to improve the questions asked of people claiming to possess one.

They also cannot turn philosophical disagreement into a numerical answer. A percentage has meaning only when there is a defensible basis for assigning it. Disagreement among theories is not automatically a probability distribution, and several indicators derived from the same assumption are not several independent confirmations. Counting them may increase the

appearance of confidence without increasing its justification.

Finally, understanding that experience is constructed does not make everyday perception worthless. The point of explaining a mechanism is to understand its successes as well as its failures. A mind that can be misled is still capable of discovering how the misleading occurred.

The one thing to keep

Return to the woman asked to imagine playing tennis. The experiment did not reveal a miniature person on a screen. It changed the route by which an answer could become observable. The absence of one kind of answer had concealed a capacity that another method could detect.

Silence needs interpreting. Its force depends on the channels that remain open, the task being asked and the reliability of the instrument listening. The same scrutiny applies to an eloquent answer: fluency also has a mechanism, and the mechanism matters to what the answer establishes.

Consciousness is something you encounter whenever a colour appears, a thought forms or a feeling matters. You need not solve its ultimate explanation to recognise the achievement involved. What should change is how you move from that familiarity to claims about another subject. An absence of evidence can deserve further investigation; convincing behaviour can deserve closer analysis. Between them lies a more demanding respect: take experience seriously enough to ask how you know it is there, and how you could have missed it.

Terms

Consciousness

The presence of subjective experience: there being something it is like for a subject at a time. This includes sensation, emotion, imagery and thought, without requiring reflection on oneself.

Phenomenal consciousness

Experience considered by how it feels or appears to the subject. It is distinguished conceptually from information available for flexible use; whether the two can come apart is disputed.

Access consciousness

Information being available for reasoning, report and flexible control. Access concerns what a representation can do within a system. Workspace theories explain conscious processing through this form of availability.

Qualia

The qualitative features of experience, such as the appearance of red or the sting of pain. The term identifies an explanatory target without settling what those features ultimately are.

Hard problem

Chalmers's name for explaining why physical or functional processes involve experience, rather than only explaining their operations. Accepting the question does not require accepting his proposed answer to it.

Explanatory gap

The apparent distance between a physical account of a process and understanding why it has its experienced character. The existence of an explanatory gap does not itself establish dualism.

Physicalism

The view that conscious states belong wholly within physical reality. Physicalists disagree about identity, realisation, function and reduction. The view leaves substantial scientific and philosophical work to be done.

Dualism

The view that mental and physical reality differ fundamentally in properties or substances. Versions differ over whether this requires a separate mind and how mental causation relates to physical processes.

Functionalism

The view that mental states are characterised by their causal roles. Different materials might realise the relevant organisation, making functionalism important to arguments about whether artificial systems could have experiences.

Panpsychism

The view that experience, or basic experiential properties, is fundamental and widespread in nature. It faces questions about how such features relate to unified minds, rather than automatically explaining them.

Illusionism

The proposal that introspection misrepresents mental activity as possessing phenomenal properties at all. It aims to explain the resulting judgements and reports, not dismiss the ordinary importance of pain or distress.

Interoception

Sensing the body's internal condition. It contributes to bodily feeling and emotion, reminding consciousness research that its subject includes an organism's lived condition as well as information about external objects.

Arousal

Physiological readiness maintained by interacting brain systems. It supports wakefulness and is relevant to sleep, anaesthesia and coma. Arousal and awareness are related, but they are not interchangeable measurements.

Responsiveness

Observable reactions to stimulation or commands. A response depends on sensory, cognitive and motor capacities. Its absence must therefore be interpreted in light of the route needed to produce it.

Metacognition

Monitoring or evaluating one's cognitive activity, often studied through confidence judgements. It is relevant to higher-order theories but is not synonymous with consciousness, accurate introspection or deliberate self-reflection.

Sentience

Capacity for experience, often used especially for positively or negatively felt states such as pleasure and pain. It is important in welfare debates and is not defined by intelligence.

Neural correlate of consciousness

The minimal neural mechanisms jointly sufficient for a specified experience under suitable background conditions. Identifying a candidate correlation is only a step towards separating those mechanisms from prerequisites and consequences.

Visual masking

A method that reduces a target's visibility using another stimulus, often close in time or space. It helps compare processing with different degrees of reported awareness while controlling presentation.

Binocular rivalry

Alternation in perception when different images are presented to the two eyes. It allows researchers to study changes in experienced content while much of the physical stimulation remains unchanged.

Blindsight

Preserved visual performance in an affected field despite loss of ordinary acknowledged vision, usually after primary visual cortex damage. Cases vary, including in whether any residual feeling is reported.

Inattentional blindness

Failure to notice an unexpected visible event while attention is engaged elsewhere. It reveals limits of reportable perception under particular task demands, without implying that no information was processed.

Global neuronal workspace

A theory in which recurrent activity makes selected information broadly available for flexible use. Its proposed ignition connects specialised processing with memory, reasoning and voluntary control, without an inner observer.

Recurrent processing

Feedback within or between neural systems, rather than a purely forward sequence. Recurrent-processing theory proposes that local sensory recurrence can support experience without the wider access needed for report.

Higher-order theory

An account in which a mental state becomes conscious through a suitable representation of that state. The representation can operate outside awareness; neither inner speech nor deliberate reflection is required.

Integrated information theory

A formal theory identifying experience with a system's intrinsic, irreducible cause-and-effect structure. Its information concerns causal organisation, not the amount of useful data an external observer can extract.

Predictive processing

A framework in which perception involves estimating causes of signals and responding to prediction errors. It offers accounts of experienced content but requires further commitments to explain conscious presence.

Perturbational complexity index

A measure of the differentiated, distributed electrical response to magnetic stimulation of cortex. It supports state assessment without requiring command-following and should not be confused with IIT's formal quantities.

Disorders of consciousness

Clinical conditions involving impaired wakefulness or behavioural evidence of awareness, including coma, unresponsive wakefulness syndrome and minimally conscious state. Categories depend on careful assessment and do not directly reveal experience.

Cognitive motor dissociation

Task-related evidence of command-following in brain activity despite absent observable command-following. It establishes a dissociation between tested capacities, not intact cognition or a particular prognosis for every affected patient.

Locked-in syndrome

Severe motor restriction with preserved consciousness, sometimes leaving eye movement as a communication route. Additional impairments can coexist. It must be distinguished from conditions in which awareness itself is impaired.

Go Deeper

Anil Seth, *Being You: A New Science of Consciousness* (Faber & Faber, 2021). Start here to find out why a world feels present and a body feels like yours. Seth makes predictive perception and bodily regulation concrete through personal and clinical examples, keeping the science attached to lived experience. His account has a clear position: it is an argument for a way of understanding consciousness, not a neutral catalogue of every theory. That gives the book its shape. The chapters on embodiment and the constructed self are a natural next step from the rubber hand and the body's internal signals.

Stanislas Dehaene, *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts* (Viking, 2014). Choose this to see how an experiment turns the uncertain moment of noticing into something testable. Masking, thresholds and the global neuronal workspace supply its main machinery. Dehaene's confidence is part of the appeal and part of what to examine: read his case alongside accounts that place more weight on posterior processing or remove the demand for a report. Later disagreements do not make the experimental chapters obsolete. They make it useful to understand exactly what those experiments showed and where a theoretical interpretation entered.

David J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory* (Oxford University Press, 1996). This is the demanding philosophical route. Chalmers explains why he thinks functional explanation leaves experience open, then develops a form of property dualism. The book is rigorous, slow and controversial. Read the opening arguments first; later chapters matter most once the distinction between logical possibility and empirical evidence is clear. The book shaped the modern debate so thoroughly that many current positions are best understood as replies to it. Expect argument rather than an introductory survey, and allow time for the thought experiments.

Christof Koch, *The Feeling of Life Itself: Why Consciousness Is*

Widespread but Can't Be Computed (MIT Press, 2019). Koch presents integrated information theory and follows its implications into animals, machines and simple systems. It is valuable because it makes a theory risk counterintuitive conclusions instead of hiding them. It is also advocacy from a central participant. Read it beside the 2025 adversarial results, the subsequent debate and the formal IIT literature. The contrast with Dehaene is useful: both take consciousness seriously as science while locating the decisive explanatory burden in strikingly different places. Koch is accessible, but the formal theory beneath the book requires separate technical reading.

Notes and Sources

Important empirical, clinical and contested claims were rechecked on 5 September 2026. Dates below distinguish original publication from the periods covered by empirical work. Bibliographic titles retain their published spelling. The book uses experience as its central sense of consciousness; its distinctions among presence, content, access and explanation are an organising framework, not a claim that four separate biological faculties have been established.

The Whole Thing in One Page and Why You Should Care

The woman in the scanner. Owen and colleagues (2006) reported task-specific functional imaging responses to motor and spatial imagery in a patient meeting behavioural criteria for the vegetative state. The account retains the published tasks and finding without inventing a bedside scene, dialogue or private reaction. It is evidence of command-following, not proof of intact cognition in every respect. Monti and colleagues (2010) subsequently reported imagery-based yes-or-no communication in one participant; the book does not generalise that communication result to all patients with detectable task responses.

The explanatory target. Nagel (1974), Jackson (1982) and Chalmers

(1995, 1996) provide the primary philosophical arguments. A difference between undergoing an experience and reading its description is not treated here as a demonstrated difference between physical and non-physical substances. Physicalist replies are discussed in Chalmers's comparative account (2003) and in the contrasting approach developed by Frankish (2016).

The Core Ideas

Experience and access. Block (1995) supplies the influential distinction between phenomenal and access consciousness. The book distinguishes the questions without assuming that the corresponding capacities must always come apart. Mashour and colleagues (2020) state the global neuronal workspace position; that position is allowed to explain conscious processing through functional organisation rather than being restricted by definition to verbal reporting.

Dissociations. The TN account follows de Gelder and colleagues (2008), a report about one patient with bilateral striate-cortex damage. Its obstacle-navigation result is not a universal model of blindsight. Simons and Chabris (1999) support the inattentional-blindness example; no rate is extrapolated beyond their varied conditions. De Haan and colleagues (2020) review split-brain findings and disagreements. Lateralised visual-field presentation must not be confused with assigning a complete eye to each hemisphere. A division of task performance does not by itself settle the number of experiencing subjects.

Bodily ownership and internal feeling. Botvinick and Cohen (1998) support the rubber-hand example. The claim concerns multisensory bodily self-identification, not a demonstration that personal identity is illusory. Seth (2013, 2021) develops predictive and interoceptive accounts of bodily feeling. Predictive inference is attributed as a theoretical approach, not presented as a complete or universally accepted theory of consciousness. Valence names the pleasant or unpleasant character of experience; an algorithmic reward or penalty is not equated with felt value. The speech-and-train heartbeat comparison is an imagined illustration of

contextual interpretation, not a reported experiment.

Organisation and intervention. Koch and colleagues (2016) distinguish neural correlates from prerequisites and consequences, and review cortical and subcortical contributions. Redinbaugh and colleagues (2020) support the specific thalamic-stimulation finding in anaesthetised macaques. Wake-like behaviour in that preparation does not establish a universal human consciousness switch or directly disclose animal phenomenology.

Perturbational complexity. Casali and colleagues (2013) introduced PCI using transcranial magnetic stimulation and the complexity of the evoked electrical response. Casarotto and colleagues (2016) provide validation across an expanded benchmark including report-supported conscious and unconscious conditions and unresponsive patients. The retained claim concerns discrimination within studied conditions. PCI is neither a direct reading of the richness of thought nor a calculation of IIT's formal intrinsic cause-and-effect structure.

Competing mechanisms. The theory descriptions draw on Lamme (2006), Brown, Lau and LeDoux (2019), Mashour and colleagues (2020), and Albantakis and colleagues (2023), with Seth and Bayne (2022) for comparison. Higher-order representations need not themselves be conscious or linguistic. IIT 4.0 contains formal requirements beyond generic connectivity, complexity or information useful to an observer. Predictive-processing approaches can overlap with other theories. The proposed evolutionary benefits of broad access are explanatory possibilities, not a demonstrated evolutionary history of feeling.

Animal experience. Crook (2021) supports the octopus place-avoidance and pain-relief-preference example, including the comparison with uninjured controls. The study supplies converging evidence for an aversive state in the tested conditions, not direct access to the animals' experience or a conclusion about every invertebrate. The New York Declaration on Animal Consciousness (19 April 2024) is an expert signatory statement. Its distinction between strong support for mammals and birds and a realistic possibility more widely is preserved; it is not described as a new experiment

or an institutional proof.

Artificial indicators. Butlin, Long, Bayne and colleagues' article appeared online on 10 November 2025 and in the June 2026 issue of *Trends in Cognitive Sciences*. The publication date is not an observation period for a new experiment. The framework derives candidate indicators from theories and does not provide a validated consciousness classifier. This book makes no definitive claim about the consciousness of a named current commercial system.

Experiments, explanations and clinical applications

Rivalry, masking and reporting. Frässle and colleagues (2014) support the use of optokinetic and pupil measures during rivalry and the finding of reduced frontal activity when active reporting was removed. The interpretation remains specific to those tasks. Dehaene (2014) and Mashour and colleagues (2020) explain masking, conscious access and associated experimental contrasts. The faint-tone listener comparison is an explicitly hypothetical illustration of sensitivity and response criterion, not a newly reported study.

Neural methods. Crick and Koch (1990) set out an influential neurobiological programme; Koch and colleagues (2016) state the neural-correlate research target and its complications. The descriptions of recording methods distinguish electrical or magnetic signals from the blood-oxygenation response used in functional imaging. Decodability of a measured signal is not treated as evidence that the sampled region is independently sufficient for an experience.

The 2025 adversarial experiment. The Cogitate Consortium paper was published online on 30 April 2025 and in *Nature* 642. Its 256 participants were studied across different recording modalities, not all through an identical three-method protocol. The study used visible images and tested specific predictions concerning content, duration and neural interaction. The text retains challenges to sustained posterior synchronisation and

workspace-related predictions without calling the result a definitive refutation of either whole theory.

The subsequent disagreement. Naccache and colleagues' response, published on 9 October 2025, contests the implications for global neuronal workspace theory. Important points concern the attention-dependence of offset ignition, interpretation of decoding, and the absence of a conscious-versus-unconscious stimulus contrast. The response is an argument by theory proponents, not a neutral final verdict. It is included alongside the original study because the precise theory-to-task entailment is material to interpreting the result.

Philosophical alternatives. Chalmers (1996, 2003) develops the hard-problem position and compares physicalist and non-physicalist approaches. Frankish (2016) states the illusionist proposal in its own terms. Mary and the philosophical zombie are thought experiments. Neither is presented as an actual experimental case, and conceivability is not treated as an uncontested proof of physical possibility. Panpsychism is described as a family of positions rather than a claim that ordinary objects think like humans.

Sleep and anaesthesia. Siclari and colleagues (2017) studied reports following awakenings in REM and non-REM sleep. The report-based posterior association is not a universal dream detector. Mashour (2024) and Brown, Lydic and Schiff (2010) support the distinctions among experience, environmental connection, responsiveness, memory and pharmacological mechanisms. Neuromuscular blockade is distinguished from hypnosis and analgesia. The discussion supplies no drug protocol or individual clinical recommendation.

Clinical assessment. Giacino, Kalmar and Whyte (2004) describe the Coma Recovery Scale-Revised. The 2018 guideline supports standardised serial assessment and attention to confounding factors. Diagnostic categories are explained as clinical assessments, not direct observations of subjective experience. Locked-in syndrome is distinguished from disorders of consciousness without assuming that every affected person has

otherwise entirely intact cognition.

The two clinical denominators. Bodien and colleagues (2024) included 353 adults in a six-centre convenience sample, with data collected during 2006-2023. Among 241 participants without observable command-following, 60 had a detected task response. Among 112 with observable command-following, 43 had one. These are different subgroups. The first proportion is not a population prevalence; the second demonstrates limited detection in participants whose behaviour already showed the relevant capacity. Modalities and task availability varied. The manuscript does not interpret non-detection as a definitive negative diagnosis.

The newer clinical synthesis. Laigaard and colleagues' patient-level meta-analysis was published on 31 July 2026 after a literature search ending on 1 May 2026. It included 56 studies and 1,248 participants overall, with heterogeneous clinical categories and methods. The manuscript uses its bounded qualitative conclusion and associations, not a pooled rate presented as a newer estimate for the Bodien cohort's population. Associations with clinical state or cause of injury are not treated as experimental causal effects.

What People Get Wrong, Use It and Terms

Development. Frohlich and Bayne (2025) examine clusters of potential consciousness markers in infancy, favouring early experience without claiming an agreed exact onset. Their article appeared online in October 2024. The manuscript rejects verbal competence as the sole criterion while leaving the timing and character of early experience open to empirical investigation.

Inference and precaution. The practical lenses are the book's reasoned applications of the distinctions and evidence above. The picture-decoding announcement and animal-policy case are explicitly hypothetical. Proportionate precaution is a normative proposal, not an empirical finding that a particular system is conscious. A decision can consider possible harm without pretending that a disputed attribution has become certain. No

numerical probability of consciousness is assigned.

Terminology and reading. Terms preserve the same meanings used in the narrative. The four Go Deeper recommendations are original specialist books with different purposes and identifiable viewpoints. The years identify the editions cited below; later paperback dates are not substituted for original hardback publication. Subsequent theoretical and empirical developments are treated in the research notes rather than silently attributed to older books.

Bibliography

Original experiments, arguments and theoretical proposals

Albantakis, L., et al. 2023. “Integrated information theory (IIT) 4.0: Formulating the properties of phenomenal existence in physical terms.” *PLOS Computational Biology* 19: e1011465.
[doi:10.1371/journal.pcbi.1011465](https://doi.org/10.1371/journal.pcbi.1011465).

Block, N. 1995. “On a confusion about a function of consciousness.” *Behavioral and Brain Sciences* 18: 227-247.
[doi:10.1017/S0140525X00038188](https://doi.org/10.1017/S0140525X00038188).

Bodien, Y. G., et al. 2024. “Cognitive Motor Dissociation in Disorders of Consciousness.” *New England Journal of Medicine* 391: 598-608.
[doi:10.1056/NEJMoa2400645](https://doi.org/10.1056/NEJMoa2400645).

Botvinick, M., and J. Cohen. 1998. “Rubber hands ‘feel’ touch that eyes see.” *Nature* 391: 756. [doi:10.1038/35784](https://doi.org/10.1038/35784).

Brown, R., H. Lau and J. E. LeDoux. 2019. “Understanding the Higher-Order Approach to Consciousness.” *Trends in Cognitive Sciences* 23: 754-768.
[doi:10.1016/j.tics.2019.06.009](https://doi.org/10.1016/j.tics.2019.06.009).

Butlin, P., R. Long, T. Bayne, et al. 2026. “Identifying indicators of consciousness in AI systems.” *Trends in Cognitive Sciences* 30: 488-501.

Published online 10 November 2025. doi:[10.1016/j.tics.2025.10.011](https://doi.org/10.1016/j.tics.2025.10.011).

Casali, A. G., et al. 2013. "A Theoretically Based Index of Consciousness Independent of Sensory Processing and Behavior." *Science Translational Medicine* 5: 198ra105. doi:[10.1126/scitranslmed.3006294](https://doi.org/10.1126/scitranslmed.3006294).

Casarotto, S., et al. 2016. "Stratification of unresponsive patients by an independently validated index of brain complexity." *Annals of Neurology* 80: 718-729. doi:[10.1002/ana.24779](https://doi.org/10.1002/ana.24779).

Chalmers, D. J. 1995. "Facing Up to the Problem of Consciousness." *Journal of Consciousness Studies* 2: 200-219. [Author's text](#).

Chalmers, D. J. 2003. "Consciousness and its Place in Nature." In *The Blackwell Guide to Philosophy of Mind*, edited by S. P. Stich and T. A. Warfield. Blackwell. [Author's text](#).

Cogitate Consortium. 2025. "Adversarial testing of global neuronal workspace and integrated information theories of consciousness." *Nature* 642: 133-142. doi:[10.1038/s41586-025-08888-1](https://doi.org/10.1038/s41586-025-08888-1).

Crick, F., and C. Koch. 1990. "Towards a neurobiological theory of consciousness." *Seminars in the Neurosciences* 2: 263-275. [Caltech author archive](#).

Crook, R. J. 2021. "Behavioral and neurophysiological evidence suggests affective pain experience in octopus." *iScience* 24: 102229. doi:[10.1016/j.isci.2021.102229](https://doi.org/10.1016/j.isci.2021.102229).

de Gelder, B., et al. 2008. "Intact navigation skills after bilateral loss of striate cortex." *Current Biology* 18: R1128-R1129. doi:[10.1016/j.cub.2008.11.002](https://doi.org/10.1016/j.cub.2008.11.002).

Frankish, K. 2016. "Illusionism as a Theory of Consciousness." *Journal of Consciousness Studies* 23 (11-12): 11-39. [Author's text](#).

Frässle, S., J. Sommer, A. Jansen, M. Naber and W. Einhäuser. 2014. "Binocular Rivalry: Frontal Activity Relates to Introspection and Action But Not to Perception." *Journal of Neuroscience* 34: 1738-1747.

[doi:10.1523/JNEUROSCI.4403-13.2014](https://doi.org/10.1523/JNEUROSCI.4403-13.2014).

Giacino, J. T., K. Kalmar and J. Whyte. 2004. "The JFK Coma Recovery Scale-Revised: measurement characteristics and diagnostic utility."

Archives of Physical Medicine and Rehabilitation 85: 2020-2029.

[doi:10.1016/j.apmr.2004.02.033](https://doi.org/10.1016/j.apmr.2004.02.033).

Jackson, F. 1982. "Epiphenomenal Qualia." *The Philosophical Quarterly* 32: 127-136. [doi:10.2307/2960077](https://doi.org/10.2307/2960077).

Lamme, V. A. F. 2006. "Towards a true neural stance on consciousness." *Trends in Cognitive Sciences* 10: 494-501. [doi:10.1016/j.tics.2006.09.001](https://doi.org/10.1016/j.tics.2006.09.001).

Mashour, G. A., P. Roelfsema, J.-P. Changeux and S. Dehaene. 2020. "Conscious Processing and the Global Neuronal Workspace Hypothesis." *Neuron* 105: 776-798. [doi:10.1016/j.neuron.2020.01.026](https://doi.org/10.1016/j.neuron.2020.01.026).

Monti, M. M., et al. 2010. "Willful Modulation of Brain Activity in Disorders of Consciousness." *New England Journal of Medicine* 362: 579-589.

[doi:10.1056/NEJMoa0905370](https://doi.org/10.1056/NEJMoa0905370).

Naccache, L., C. Sergent, S. Dehaene, X.-J. Wang, M. Farisco and J.-P. Changeux. 2025. "GNW theoretical framework and the 'adversarial testing of global neuronal workspace and integrated information theories of consciousness'." *Neuroscience of Consciousness* 2025 (1): niaf037.

[doi:10.1093/nc/niaf037](https://doi.org/10.1093/nc/niaf037).

Nagel, T. 1974. "What Is It Like to Be a Bat?" *The Philosophical Review* 83: 435-450. [doi:10.2307/2183914](https://doi.org/10.2307/2183914).

Owen, A. M., et al. 2006. "Detecting awareness in the vegetative state." *Science* 313: 1402. [doi:10.1126/science.1130197](https://doi.org/10.1126/science.1130197).

Redinbaugh, M. J., et al. 2020. "Thalamus Modulates Consciousness via Layer-Specific Control of Cortex." *Neuron* 106: 66-75.e12.

[doi:10.1016/j.neuron.2020.01.005](https://doi.org/10.1016/j.neuron.2020.01.005).

Seth, A. K. 2013. "Interoceptive inference, emotion, and the embodied self." *Trends in Cognitive Sciences* 17: 565-573. [doi:10.1016/j.tics.2013.09.007](https://doi.org/10.1016/j.tics.2013.09.007).

Siclari, F., et al. 2017. "The neural correlates of dreaming." *Nature Neuroscience* 20: 872-878. doi:[10.1038/nn.4545](https://doi.org/10.1038/nn.4545).

Simons, D. J., and C. F. Chabris. 1999. "Gorillas in our midst: sustained inattentional blindness for dynamic events." *Perception* 28: 1059-1074. doi:[10.1068/p281059](https://doi.org/10.1068/p281059).

Research syntheses, guidance and expert statements

Brown, E. N., R. Lydic and N. D. Schiff. 2010. "General Anesthesia, Sleep, and Coma." *New England Journal of Medicine* 363: 2638-2650. doi:[10.1056/NEJMr0808281](https://doi.org/10.1056/NEJMr0808281).

de Haan, E. H. F., et al. 2020. "Split-Brain: What We Know Now and Why This is Important for Understanding Consciousness." *Neuropsychology Review* 30: 224-233. doi:[10.1007/s11065-020-09439-3](https://doi.org/10.1007/s11065-020-09439-3).

Frohlich, J., and T. Bayne. 2025. "Markers of consciousness in infants: Towards a 'cluster-based' approach." *Acta Paediatrica* 114: 285-291. Published online 14 October 2024. doi:[10.1111/apa.17449](https://doi.org/10.1111/apa.17449).

Giacino, J. T., et al. 2018. "Practice guideline update recommendations summary: Disorders of consciousness." *Neurology* 91: 450-460. doi:[10.1212/WNL.0000000000005926](https://doi.org/10.1212/WNL.0000000000005926).

Koch, C., M. Massimini, M. Boly and G. Tononi. 2016. "Neural correlates of consciousness: progress and problems." *Nature Reviews Neuroscience* 17: 307-321. doi:[10.1038/nrn.2016.22](https://doi.org/10.1038/nrn.2016.22).

Laigaard, P. P., F. W. Abia, M. Hassani, A. K. Eigenbrodt and D. Kondziella. 2026. "Cognitive Motor Dissociation in Disorders of Consciousness: An Individual Participant Data Meta-Analysis." *European Journal of Neurology* 33: e70713. doi:[10.1111/ene.70713](https://doi.org/10.1111/ene.70713).

Mashour, G. A. 2024. "Anesthesia and the neurobiology of consciousness." *Neuron* 112: 1553-1567. doi:[10.1016/j.neuron.2024.03.002](https://doi.org/10.1016/j.neuron.2024.03.002).

New York Declaration on Animal Consciousness. 19 April 2024.

Declaration.

Seth, A. K., and T. Bayne. 2022. "Theories of consciousness." *Nature Reviews Neuroscience* 23: 439-452. doi:[10.1038/s41583-022-00587-4](https://doi.org/10.1038/s41583-022-00587-4).

Books

Chalmers, D. J. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.

Dehaene, S. 2014. *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. Viking.

Koch, C. 2019. *The Feeling of Life Itself: Why Consciousness Is Widespread but Can't Be Computed*. MIT Press.

Seth, A. 2021. *Being You: A New Science of Consciousness*. Faber & Faber.