Beyond IIT: (how) can we model the evolution of consciousness?

Adrian Kent

*Centre for Quantum Information and Foundations, DAMTP,
Centre for Mathematical Sciences, University of Cambridge,
Wilberforce Road, Cambridge, CB3 0WA, U.K. and*

*Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, ON N2L 2Y5, Canada.**

(Dated: September 2021)

Tononi et al.’s “integrated information theory” (IIT) postulates rules for assigning measures Φ and qualia types Q of consciousness to classical information networks. We consider whether IIT is compatible with Darwinian evolution. We argue that an IIT-like theory that assigns consciousness to physical systems by relatively simple mathematical rules poses extraordinary fine-tuning problems. For example, why, among all possible lawlike theories of consciousness, do we have one that makes us conscious of a high-level narrative of our environment and actions, so accurate that it appears to us to cause our behaviour?

We introduce IIT^+ , a class of extensions of IIT in which Φ and/or Q influence the network dynamics. We argue that IIT^+ -like theories, unlike IIT-like theories, offer at least partial explanations of how some key features of consciousness evolved. We conclude that if one takes seriously Darwinian evolution and the case for an IIT-like theory, one has to take seriously the case for an IIT^+ -like theory.

INTRODUCTION

No position on the mind-body problem is plausible. Materialism: implausible. Dualism: implausible. Idealism: implausible. Neutral monism: implausible. None of the above: implausible. But the probabilities of all these views get a boost from the fact that one of the views must be true. [1]

From the perspective of fundamental physics, primitive natural phenomena ought to have lawlike explanations. The known physical laws say nothing about consciousness; not only do they give no way of characterizing it, but nothing in them – with the possible exception of the quantum measurement postulate, for which the case is arguable but very unclear – suggests it should even be a feature of the universe. This inspires the thought that there should be new laws in the form of a relatively concise set of equations defining an elegant physical theory of consciousness (EPTC) that describes how conscious minds inhere in and relate to matter and space-time. Observing that some types of matter form magnets led eventually to Maxwell’s theory and the concept of electromagnetic field, and ultimately to identifying spin as a quantum observable and to quantum electrodynamics. Perhaps, then, empirically guided theoretical work could eventually lead us to the right concepts and equations for an EPTC?

Needless to say, putting it mildly, there are difficulties with this idea. We try here to assess the possibilities as well as the problems. We focus in particular on whether and how an EPTC could possibly be compatible with Darwinian evolution. We use definitions and concepts from Tononi et al.’s “integrated information theory” (IIT) of consciousness [2] to illustrate the arguments. Although Tononi and other leading proponents frame it differently, IIT can naturally be interpreted as a proposal for an EPTC, or at least a work in progress towards this ultimate goal. Our arguments do not depend on the details of IIT being valid (indeed in our view they probably are not). Its role in our discussion is as a concrete example that illustrates two points. First, a mathematically defined EPTC could have (at least some) explanatory power and could be (at least to some degree) empirically confirmed. Second, however, an EPTC that, like IIT, restricts its role to describing how consciousness supervenes on matter is incapable of explaining how key features of consciousness evolved.

We argue further that EPTCs can have significantly greater explanatory power, and in particular can explain significantly more about the evolution of consciousness, if they assign a dynamical role to quantities that (they postulate) define some aspect of consciousness. Some aspects of consciousness would still elude explanation. For example, it is hard to see how an EPTC can completely answer James’ question as to why pleasure and pain qualia are well correlated with evolutionary benefit and cost. Also, scientific conservatism inclines against postulating a new type of dynamical law, and the scope for such laws is quite theoretically and empirically constrained.

These would be very strong objections if a more immediately plausible theory of consciousness with greater explanatory power were available. However, we argue, none currently is. In particular, there seems no immediate prospect of a theory that explaining how the intrinsic qualities of consciousness – the sensation of redness caused by seeing red, and so on – evolved. This suggests we should perhaps restrict attention to theories with the more limited aim of explaining how consciousness could have evolved to have the type of information content it does. Here, we

argue, EPTCs involving new dynamics could conceivably give good explanations. We conclude that, relative to other contenders, they are better motivated and more promising than is generally recognised.

IIT AS AN EPTC

Tononi et al.’s work towards developing an “integrated information theory” (IIT) of consciousness [2] has attracted widespread attention. It has been criticized for a variety of reasons (e.g. [3–6]). Specific criticisms aside, it is reasonable to be sceptical about the prospects of any first principles theory of consciousness that is based on the limited high-level data presently available. Still, even critics and sceptics generally recognize IIT as a serious proposal that gives an interesting perspective on the problems of consciousness.

Tononi et al. take consciousness as primary and argue that the nature of conscious experiences, including their integration into a unified consciousness, can be characterised by a relatively small set of axioms. These motivate postulates about necessary features of physical substrates underlying conscious experiences, and in particular the causal relations between components of those substrates. As Tononi puts it [7] “... these postulates are inferences that go from phenomenology to physics, not the other way around. This is because the existence of one’s consciousness and its other essential properties is certain, whereas the existence and properties of the physical world are conjectures, though very good ones, made from within our own consciousness.”

IIT’s proponents generally appear to prefer to avoid committing to a definite ontology. However, its postulates are quite consistent with a physical realist worldview. On this view, we can take IIT as a theory that predicts which physical systems have an associated consciousness, quantifies how conscious any given conscious physical system is, and aims (albeit in a way that seems to need further elaboration) to describe what it is conscious of. In principle, given a complete physical description of the world, it is supposed to predict how many consciousnesses can be found at any given time, the systems to which they are attached, and what each is conscious of. In practice, both proponents and critics of IIT tend to speak like physical realists. Discussions of the level of consciousness IIT assigns to, for example, the cerebellum [7], a photodiode [8] or a simple logical array [4, 9] suggest little real doubt that these systems exist.

In any case, as we want to consider lawlike physical theories of consciousness, we will not entertain any such doubt here. We assume that both the physical world and consciousness are real, with the physical world described by a given physical theory with a given definite ontology, both defined by mathematical laws. We want to investigate the prospects for and potential implications of some form of EPTC, which would add a lawlike description of consciousness to our lawlike description of the natural world. We use IIT, read as an EPTC, as an example. In principle the underlying physical theory could be classical mechanics, a realist version, extension or modification of quantum mechanics or relativistic quantum field theory, general relativity, or some theory unifying quantum theory and gravity. IIT’s proponents suggest [10] that it should apply to continuous classical models, with system-dependent spatio-temporal coarse-grainings derivable in principle from its postulates. Extensions of IIT to quantum theory have also been considered [3, 11–15]. The status of these various proposed extensions seems not fully settled. Instead, discussions of IIT and what it says about the conscious states of physical systems generally consider discrete Markovian network models with transition matrices associated with a fixed discrete time step (see e.g. [7]).

Clearly, this does not fit naturally with our best fundamental physical theories, relativistic quantum field theory and general relativity [3]. Still, we don’t actually know whether quantum theory plays an essential role in a fundamental theory of consciousness (if indeed there is one). Nor do we know any relevant empirical facts about the implications of relativity for consciousness (though relativistic causality must surely at least constrain the information accessible to any given consciousness), nor any connection between consciousness and gravity. Non-relativistic classical dynamics, within its domain, is a very good theory of the material world. It is reasonable to hope that non-relativistic classical theories of consciousness might turn out to be useful, even if more general theories will ultimately be needed.

Also, IIT analyses of fixed network models are not directly applicable to creatures whose bodies and brains grow and change within dynamical environments, to reproduction, or to evolution. Discussing, for example, what IIT says about the consciousnesses to be found on and around Earth and their evolution, thus requires some leaps of faith. One might perhaps hope that a model of a gigantic dynamic discrete network, with links and nodes joining, changing and leaving, could in principle capture enough qualitative features of organisms living and evolving in an environment to give a coherent story. Alternatively, one might hope that combining dynamic network models of subsystems (such as living creatures) with a more realistic physical model of their environment might. Or one might hope that some extension of IIT applies to quantum systems in space-time, and that classical discrete network models normally (on and around Earth) give a good qualitative approximation. None of these is evidently true, though each of them could be. IIT’s proponents clearly intend that it *should* apply to brains in environments, following the causal chain back to cosmic rays impacting organisms and beyond [10]. As far as we are aware, there is no rigorous underpinning

to this ambitious vision, nor any imminent prospect of one. One just has to take a view on whether at least one of these possibilities is plausible enough to continue the discussion, accepting that any IIT-based discussion of the emergence of conscious life on Earth presently necessarily has to follow Tononi and collaborators [16, 17] in combining mathematical models with intuitions that cannot presently be carefully justified.

We tactically and charitably suspend scepticism here. Our focus is not on analysing whether IIT could ever be made coherent enough to rigorously establish intelligible accounts of the initial development of consciousness on Earth and its subsequent role in the course of terrestrial evolution. but on what the implications would be *if* it – or some qualitatively similar theory – could. We will argue that any such accounts would necessarily require almost incredible fine-tuning. In other words, we will argue that, even giving IIT-type theories the benefit of the doubt on other points, they cannot explain the evolution of consciousness.

We take a similar charitable stance on other criticisms of IIT, which we briefly now review, along with possible defences.

Motivation for axioms

IIT’s axioms may indeed, as critics argue, involve seemingly ad hoc choices, but they are at least motivated by real data, including the apparent near-unconsciousness of the cerebellum, the levels of consciousness reported in various sleeping and waking states, and the effects of drugs on consciousness. They also offer some sort of theoretical explanation for some intriguing features of consciousness – in particular that consciousnesses are separate and internally unified. One could hope that further work might eventually produce refined axioms that are more directly empirically supported.

Aaronson’s hyper-conscious expander

IIT assigns high measures of consciousness to very simple logic arrays [4], which runs contrary to many people’s intuitions. However, as Tononi notes [9], the relevant intuitions could simply be wrong. Aaronson [4] accuses Tononi of a double standard here: if we accept the intuition (supported by IIT) that the cerebellum is at most minimally conscious, we should also accept the intuition (contrary to IIT) that a simple 2D grid of logic gates is also at most minimally conscious. Tononi and colleagues counter that the theory is based on analysing and extrapolating from *human conscious experience*. Our own experience and other humans’ reports of theirs are our only sources of real data about consciousness. For this reason, they assign some weight to intuitions about the conscious states associated with parts or states of the human brain, particularly if they fit into a parsimonious theory. In contrast, they give essentially no weight to pre-existing intuitions about the consciousness of non-biological systems. This seems a reasonable rationale: of course there is no guarantee that it will turn out to be a productive way of formulating theories of consciousness, but it is one of very few sensible options.

IIT and the boundaries of consciousness science

The privileging of intuitions based on human biology is relevant to another closely related criticism of IIT: that it breaks the link between consciousness and reports of consciousness and so removes the possibility of empirical test. Unless we hope to develop and confirm an IIT-like theory entirely via our own conscious experience, we have to assume that other humans, when their brains are functioning normally, can and normally do report (many aspects of) their conscious states (fairly) accurately.[28] Proponents and critics of IIT agree [17, 18] that artificial devices can be constructed so that they respond to stimuli in the same way as humans, and in particular give the same reports about their consciousness, while according to IIT they are unconscious. So, if we accept reports as conclusive evidence of consciousness, then IIT is wrong. If we accept that the same set of reports can correspond to either to a rich human conscious state or to an unconscious computer state, then, some argue (see e.g. [18]), we abandon empirically based science. Here again IIT’s defenders (e.g. [19]) stress the theory’s basis in the connection between human experience and human biology and the crucial role of parsimony in science. The aspiration is to develop a theory whose account of human consciousness is sufficiently parsimonious and accurate to justify assigning some credence to its extension to non-biological systems.

Differing stances on the scope and limits of science and the scientific method applied to consciousness are evident here. Essentially, the question at issue is the following. Suppose we could build a neuronal-level imaging device

connected to an immensely powerful super-computer programmed using IIT 94.0 – a hypothetical future version of IIT refined by much theoretical and empirical work, whose postulates everyone finds awesomely simple and elegant. When the device scans a human subject it states whether they are conscious, and (if so) describes their consciousness in full detail. Suppose too that we test it on large numbers of subjects, of the widest possible variety of ages, genders, and cultural backgrounds – perhaps even on all living humans – all of whom report that the descriptions are accurate.[29] And now suppose we try the device on other potential subjects: various animals, plants, rocks, logic arrays, computers running AI programmes, and so on. Should we place any credence in its statements about whether they are conscious and its descriptions of their consciousnesses? For what it is worth, we agree with IIT’s proponents that we should give these statements *some* credence out of a preference for parsimony. “IIT 94.0 applies universally” is a more parsimonious theory than, say, “IIT 94.0 applies to all humans (or all biological organisms), and only humans (or only biological organisms) are conscious”, for example. Indeed, theories like this last are even more complicated and less parsimonious than they may initially seem, because slippery slope arguments (using manoeuvres such as sequentially replacing small biological components of a brain by functionally equivalent artificial components) seem to preclude a suitably succinct definition of “human” or “biological organism”.

SUMMARY SO FAR

There is good motivation for seeking an elegant physical theory of consciousness. IIT (suitably re-interpreted) is worth considering as a concrete example of an EPTC. The arguments above against it so far are serious but not completely compelling; the defences above seem just about tenable, for now. However, they raise further difficult questions. For example, are there ((infinitely?) many?) other possible sets of axioms that could also be motivated? And is there a justification for privileging intuitions about biological consciousness that survives more careful scrutiny?

These, we will argue, lead to problems that make not only IIT, but any EPTC of its type, essentially untenable. We will then consider whether some of those problems could in principle be solved by different types of EPTC, which we illustrate by toy theories that we characterise as type IIT⁺, defined by adding qualitatively new features to IIT-type theories. In particular, we will argue that, if IIT 94.0 (or any theory that is IIT-like, in a sense to be elaborated) were to work as described, then fine-tuning considerations would suggest there should be a better IIT⁺-like theory that *could* be better supported by evidence.

DEFINITIONS

We work within a simplified physical ontology, in which the biosphere and environment on and around Earth is taken to be modellable by a finite Markovian network in Galilean space with a discrete time coordinate labelled by the non-negative integers. The nodes $A_i (1 \leq i \leq n)$ have finite discrete state spaces that undergo transitions at discrete integer times with the new state probabilities determined by their inputs. We write $N = \{A_i\}$ for the set of all nodes. The nodes include all neurons in all species, and much else, so n is extremely large. The inputs to A_i at time t are the states of a set of connected input nodes I_i (which may include A_i itself) at time $t - 1$, transmitted along wires from each $A_j \in I_i$ to A_i . We take these wires to be non-intersecting. The network state at positive integer time t is given by $\mathbf{x}(t) = (x_1(t), \dots, x_n(t))$, where $x_i(t) \in \{0, \dots, n_i - 1\}$ is the value of the random variable $X_i(t)$ associated with node A_i . Writing $\mathbf{X}(t) = (X_1(t), \dots, X_n(t))$, the Markovian condition is

$$P(\mathbf{X}(t)|\mathbf{X}(t-1)) = P(\mathbf{X}(t)|\mathbf{X}(t-1), \mathbf{X}(t-2), \dots, \mathbf{X}(0)).$$

In this ontology, IIT identifies a number of independent conscious systems associated with each possible network state at each integer time $t > 0$. If IIT is correct and the model is applicable, these should include at least one conscious state for each awake human (and also for many if not all of those asleep); the expectation is that they would also include at least one conscious state for each higher-order animal that is awake, and perhaps conscious states even for the most primitive organisms, conscious states for some artificial computers, and perhaps (maybe low-level) conscious states for other natural and artificial systems. According to IIT, conscious systems are identified by defining [2, 7] a quantity Φ calculated from the network state and transition matrix, and identifying conscious systems with maximisers of Φ . Here candidate systems are subsets $S \subseteq N$ of the set of nodes. A system S is defined to be a maximiser if $\Phi_S > \Phi_{S'}$ for all $S' \subset S$ and $\Phi_S > \Phi_{S''}$ for all $S'' \supset S$. This allows the possibility that there could be many (perhaps mostly fairly minimal) consciousnesses associated with an individual human or animal brain or body. However the expectation is that, at least for humans and animals with qualitatively similar brain organization, IIT

will identify a “main consciousness” that has a much higher value of Φ than any others. For normally functioning humans with undamaged brains, the main consciousness should be the consciousness that we are familiar with and report on.

Whether any of these expected predictions of IIT actually hold true in realistic models is unknown. Even for a single brain, calculations of Φ are too complex to be feasible. The arguments that IIT should give an intelligible picture that accords with our intuitions about humans and animals are themselves based on intuition and difficult to assess. Suspending scepticism, we assume that something like the expected picture does indeed emerge.

We can then focus on subsets S in states $\mathbf{x}^S(t) = \{x_i(t)\}_{i \in S}$ that define a single animal’s “main consciousness”; the intuitions underlying IIT suggest most or all of the elements of S will be neurons in the animal’s cerebral cortex. The corresponding conscious state is postulated to be described by a *conceptual structure*, a set of the form

$$\{(S_1^c, p_1^c, S_1, S_1^f, p_1^f, \phi_1), \dots, (S_N^c, p_N^c, S_N, S_N^f, p_N^f, \phi_N)\}. \quad (1)$$

Here the S_i (*mechanisms*), S_i^c (*past purviews*) and S_i^f (*future purviews*) are subsets of S , the $\phi_i > 0$ are real numbers, p_i^c and p_i^f are probability distributions on S_i^c and S_i^f respectively. Each set $(S_i^c, p_i^c, S_i, S_i^f, p_i^f, \phi_i)$ is supposed to correspond to a *quale*, or elementary conscious sensation, with ϕ_i defining its intensity; the composition of qualia at the given intensities is supposed to define the conscious state. As Tononi puts it:[7]

“The overall “form” of the conceptual structure or quale sensu lato (Q) ... is identical to the quality of the experience, how the experience feels. ... Different forms correspond to different experiences: they feel the way they do – red feeling different from blue or from a headache –because of the distinct shape of their qualia.”

The full definitions of conceptual structure [2, 7] are complicated and may not be fully settled (a new version, IIT 4.0, is promised). The precise details are not so critical to our discussion: the key point is that the conceptual structure is completely determined by the network state and its transition probability matrix.

Qualia and information

It is hard to see how mathematical quantities defined from information flows on networks can describe our consciousness. Tononi et al. recommend [2, 7] picturing the descriptions given by the current definitions by plotting sets $(S_i^c, p_i^c, S_i, S_i^f, p_i^f, \phi_i)$ as stars in a high-dimensional graph, with brightnesses corresponding to the intensities ϕ_i . This helpfully emphasizes that there are natural distance measures on the space of conscious states, and suggests that states can be grouped by closeness and by other similarities. But, however long one stares at one of these “constellations”, it does not evoke any ordinary conscious experience; conversely, the experience of seeing a blue book does not suggest any particular constellation. Future refinements of IIT’s definitions could alter the mathematical representations of conscious states, but the knowledge argument [20, 21] suggests the gap between mathematical representation and experience is unbridgeable. A brilliant scientist, equipped with a supercomputer and fine-grained real-time brain scanner, growing up in an experientially impoverished environment, could in principle identify the IIT constellation characterising each experience he has. Still, given a constellation corresponding to a qualitatively new experience E , he could not know what it represented before actually experiencing E .

We attempt here another charitable (though unauthorised) interpretation and partial defence of IIT’s strategy. Arguably, the goal of a theory like IIT, properly understood, need not be to explain the *qualities* of experience – the redness of red, the nature of the peace arising from harmony with companions, and so on. Its more limited (though still extremely ambitious) goal is to characterise what can be said about the content of experience in mathematical (and in the broad sense information-theoretic) terms – how much one is conscious of, how conscious one is of x relative to y , how closely sensations of red and blue are related compared to sensations of red and green or red and cold, and so on.

There may be a helpful analogy here with how physics describes the material world in terms of mathematical quantities and relationships – position, relative motion, expectation value, mass density. Quantum physics predicts statistically how these quantities and relationships change – what the chances of an electron being detected at a certain position when sent through a double slit interferometer are, for example. This still seems to leave further questions unanswered, namely: What actually *is* an electron? Or indeed a double slit interferometer? These are not questions about the interpretation of quantum theory in the standard sense, or about alternatives such as Bohmian mechanics, many-worlds theories or dynamical collapse models that give alternative ontologies. In each case, those ontologies are still defined mathematically. The question here is about the entities the mathematics represents. What

are the inherent natures – the quiddities, in Russell’s language [22, 23] – of electrons, interferometers, detectors? Or of mass densities or space-time, in ontologies in which these are fundamental?

Not all questions have meaningful answers. Perhaps these don’t: certainly it is a struggle to imagine what form a satisfactory answer could take. If there is a scientific account of the quiddities of matter, it seems well beyond anything we can currently conceive. We distinguish between electrons and mathematical quantities we use to represent them, and the distinction seems fundamental. Nonetheless, physics works very well by speaking about the mathematical quantities (such as wave functions) and relating them to other mathematical quantities (such as probabilities of detection at a given position), while keeping silent about quiddities.

It is harder to accept the possibility that it might not be meaningful to discuss the qualities of our experiences. Unlike electrons, experiences present themselves to us with a Cartesian immediacy. But it’s perhaps nonetheless plausible that a very good – useful, predictive, insightful, explanatory – theory of consciousness could represent experiential states mathematically, while remaining silent on why mathematical quantities A and B translate into seeing red and blue respectively and on the intrinsic qualities of these subjective experiences. We might expect such a theory to tell us that seeing two similar shades of red, red_1 and red_2 , are described by relatively similar but distinct mathematical structures, without describing the nature of those subjective experiences – just as physics tells us that electrons and muons have similar properties, including identical charges and spins, but are distinguished (inter alia) by their masses, without describing the intrinsic nature of electrons and muons (or indeed of charge and mass). Perhaps this may even be the best we can hope for from a theory of consciousness for the foreseeable future. Any possible scientific account of the intrinsic qualities of experience may also be well beyond anything we can currently conceive.[30]

So, continuing to suspend scepticism, we accept IIT’s working hypothesis that analysing the human brain’s function at (for example) the neuronal level could eventually give a mathematical dictionary for familiar experiences that characterises their degrees of relationship in ways generally aligned with introspective exploration, so that the experiences of seeing colours have some characteristic mathematical features, with the experiences of seeing nearby shades of red closer than seeing either is to seeing blue, the experiences of hearing sounds have distinct characteristic mathematical features, with those of hearing a given note on two different musical instruments closer than either is to hearing a wolf snarl, and so on.

EVOLUTION AND THE FINE-TUNING PROBLEM

Modulo a good deal of charitable optimism, then, we have the following picture.

There is a set of mathematical laws of consciousness, given by IIT, or by some future extension, which in principle apply to quantum systems defined in continuous time and space. These might be a combination of something like the present version of IIT with some rules for prescribing discrete classical network approximations to quantum systems, or some more fundamental modification of IIT. They characterize the consciousnesses arising in the physical world. The consciousnesses characterized include, for most humans, their “main consciousness”. This is what we normally refer to simply as our consciousness; we are normally able to report its contents, and normally have a sense of being able to do so accurately, though perhaps not completely. They also very likely characterize a main consciousness for most present-day higher animals, and for animals that were on Earth in the past. Our credence in this diminishes as we consider less complex organisms – worms, plants, bacteria, ... – and as we go further back in time. The laws may also perhaps characterize other, perhaps minimal, independent consciousnesses associated with human brains and bodies, of which “we” – identified with our main consciousness – are unaware; we will not need to discuss these, nor their animal counterparts, nor the intriguing possibility [17] that octopuses and other creatures with radically different neural configurations from ours may have multiple complex consciousnesses rather than one main consciousness. The laws might also characterise further, perhaps minimal, consciousnesses associated with computers, machines, or other inanimate objects; we will not need to discuss these either.

In this picture, the IIT-like laws define an elegant physical theory of consciousness. Although their form and content is unlike known physical laws, they have a similar status. That is, as far as we understand it, it is a contingent fact that they hold in our universe: infinitely many other laws are logically possible. For example, as we currently understand it, any of the 19 parameters defining the standard model could be different; it is just a brute fact about our universe that they take the value they do, and similarly for other features of the model. A physical theory with so many brute facts is unsatisfactory: parsimony encourages the hope that some of these will be explained by a future unified theory that in some sense has fewer brute facts (for example by explaining relations between the standard model parameters). Still, a physics without brute facts seems inconceivable – absent a presently unimaginable transformation in our conceptual framework – since there are many mathematically consistent ways to build theories.

Relatively simple examples of ways in which IIT's laws could have been different include identifying conscious systems with all subsets of N or identifying them in a way that depends probabilistically on their Φ value and those of their subsets and supersets (and so, in either case, allowing overlapping consciousnesses), defining Φ to depend on the number of nodes in S and their logical functions as well as on the state and transition matrix, framing definitions in terms of some probability distribution of present states (perhaps peaked around the actual state) and altering the definitions of qualia by rules similar to any of these. The consciousness level and qualia of a network could also depend on properties of the distribution of outputs given (for example) uniformly distributed inputs, or on the properties of other (for example, according to some given cost function, simpler) networks that approximately or exactly reproduce its input-output transitions. Readers can no doubt easily imagine many further possibilities.

The computational complexity of applying IIT's laws to the human brain precludes explicit calculations of a main human consciousness and its qualia, and the same is true of alternatives that similarly require individual calculations for each partition of a set or subset. However, most such alternatives would clearly predict qualitatively different consciousnesses and qualia. If some future version of IIT does indeed accurately characterise our main consciousness, only a tiny fraction of the alternatives would predict that anything remotely resembling it is associated with a normally functioning human brain: almost all of them would predict consciousnesses far from our experience, likely in most cases with nothing resembling what we would understand as a coherent mind. We can represent this abstractly as a list of theories $T_1, \dots, T_N$, where T_1 is IIT, T_i and T_j predict qualitatively different consciousnesses for the human brain for any distinct pair i, j , and N is a very large number.

This illustrates the fine-tuning problem. As we understand it, the process of Darwinian evolution is a consequence of the physical laws governing the material world. Those laws include quantum randomness, so it was not predictable in advance that humans, or creatures with nerve cells and neurons, or any form of life, would appear on Earth. But given that primitive life did evolve, we understand that evolutionary pressure drives towards maximizing survival and reproduction chances. Our brains and bodies have evolved forms that allow evolutionary beneficial behaviour in a wide range of contexts and environments, while using resources relatively parsimoniously. Our neural network design is very probably not the mathematically optimal solution to any given input-output problem, but it is a very good solution to a wide range of problems – far better than one we could currently design, and incomparably better than a random network. Our brains did *not*, as we currently understand it, take the form they do *because of* the form of IIT or any IIT-like theory. Yet, if IIT (aka T_1) is correct, we happen to live in a world in which our brains *both* are evolutionarily highly adapted for our survival *and* have a main consciousness that is not only coherent but also gives a high-level narrative about our behaviour convincing enough that we are tempted to believe it causes our behaviour. This seems an extraordinary coincidence, since there is no *a priori* reason why these features should coincide. For, as far as we understand it, we could equally have been in a world where our brains had the same form but T_2, T_3 or $\dots T_N$ were correct. In most of these cases our consciousnesses would have been incoherent or non-existent; in most of the rest they would have had no evident connection with our behaviour. We would be evolutionarily successful, but unaware of our triumph.

The point applies equally well if any IIT-like EPTC holds: why does it just happen to belong to the astronomically tiny subset that supplies a relevant-seeming consciousness to the end-products of billions of years of Darwinian evolution?

ASPECTS OF THE FINE-TUNING PROBLEM

Looking more closely at consciousness and the fine-tuning problem, one might ask the following:

1. Why is there consciousness at all?[31]
2. Given that consciousness exists, why are *we* conscious? Why does consciousness have anything to do with human brains?
3. Given that we are conscious, why does our consciousness focus on aspects of our environment and thoughts relevant to survival, well-being and reproduction?
4. How and why can we speak about the fact of consciousness and the quality and contents of our consciousnesses?
5. Why do pleasure and pain qualia correlate with evolutionary benefit and harm?

Depending on one's stance on theories of consciousness, these questions may appear answerable, completely unanswerable, partially answerable, or potentially (partially) answerable given further progress on one's preferred theory.

In this section, we try to answer them, to the limited extent possible, from the perspective of someone seeking to develop an IIT-like EPTC.

1. The existence of consciousness is a brute fact of nature. Descartes, Tononi and others would argue it is the primary and only indubitable fact. Physical realists would say that the existence of a material world is so compelling a postulate, because of its greater explanatory power than any form of idealism, that we should accept this as another brute fact. On either view, the existence of consciousness is not something for which an explanation could be possible, within our current conception of science, any more than material physics can explain why is there something rather than nothing.

Perhaps, some might suggest, there are hints of a future conception of science in which the two problems become one; perhaps, as Wheeler and others have suggested, some future post-quantum theory will make it clear that the material universe requires observers, that existence requires information about, following the slogan “It from Bit”. It is hard to assess such poetic ideas scientifically (even, or perhaps especially, if one likes them as poetry). They are, in any case, completely separate from the idea of an IIT-like EPTC, so we will not discuss them further here.

2. We have to accept as a brute fact that consciousness depends on information flow rather than, say, mass or charge density. More precisely, we have to accept that the type of consciousness exhibited by our main consciousnesses does. We can focus on this type, since we have no evidence for other types, and they do not arise in IIT-like theories.

Now, according to IIT-like theories, it is a lawlike fact about nature that something roughly characterisable as “integrated information flow” is the physical correlate of consciousness. If IIT intuitions are correct, it may also “follow” – in the sense that we may be able to sketch a high-level explanation drawing on ecology, biology, paleontology, chemistry, classical dynamics, quantum theory and neuroscience, inter alia – from the laws of material physics that evolutionary successful complex species in complex competitive environments tended to evolve brains and nervous systems in which said integrated information flow is a prominent feature. Even if so, the fact that both the fundamental theory of consciousness and the history of terrestrial evolution select this same feature is unexplained.

3. If IIT intuitions are correct, it must be that the processes directly involving decision or learning about beneficial or detrimental events, features and actions and environmental monitoring are implemented within the subcomponent of our brain selected by an IIT-like theory as the “main complex” of our neural network. The IIT-like theory further defines a “conceptual structure” associated with the main complex, and it must be that this describes our main consciousness. That is, the qualia defined mathematically by the IIT-like theory for this conceptual structure represent a high-level narrative about evolutionarily relevant aspects of our internal and external environment. Nearly all the other processes that our brains control – including regulation of heart rate, breathing, digestion, automated learned routines, and also including much pre-processing and post-processing of our conscious experience – must not feature significantly in this conceptual structure, since we are not conscious of them. This could either because they do not take place within the main complex or because they do not play a significant part in the cause-effect relationships defining the conceptual structure (according to IIT) or whatever plays this role (in a more general IIT-like theory). One reason why they might not take place within the main complex is that they are processed without involving information integrated over many sites. Another could be that, while their processing involves some integrated information processing, it takes place within other “minor complexes” in the brain – separate systems that, according to IIT’s postulates, have separate (perhaps low-level) consciousnesses.

None of this is obviously true. Identifying the conceptual structure and associated qualia for a human brain, according to IIT’s definitions, is too complex a computational problem to be tractable. The connection between the mathematical objects that IIT, or an IIT-like theory, define as qualia, and the types of qualia we actually experience, is so theoretically elusive that it is hard to know whether this aspect of the narrative can be confirmed or falsified at all. (Even if we could compute them, how could we tell whether the mathematical objects identified by an IIT-like theory are the “right” quantities to describe our qualia?) It’s thus hard to be clear whether this narrative is presently more than wishful thinking.

4. There is no direct answer as to why we can speak about the fact of consciousness. In a network model of brain function, our speech corresponds to network outputs. According to IIT-like theories, the qualia defining our main consciousness correspond to complex mathematical structures defined by information flows within the network. These have no simple relationship to the network’s outputs. Again, calculating these structures is a problem of enormous computational complexity. As far as we understand it, no feature of our brain’s neural networks is designed to address this problem. Indeed, it seems extremely unlikely that our brains have the capacity to do so, still less to do so in the 0.5 seconds or so that elapses between experiencing stimuli and vocalizing our conscious awareness of them.

Even if we set aside these seemingly compelling objections, and look for a less direct answer, similar problems arise. For an information processing network to report on its own qualia seems to require it to comprise something like a primary network A – which actually experiences the qualia – together with a monitoring network B that reproduces many, if not all, of the calculations prescribed by IIT to identify the qualia associated with network A, and acts on the

results, for example by speaking about them. Plausibly, Network B would give these calculated data some collective name, “consciousness” to distinguish them from other things it may report. But why would evolution produce such a monitoring network? How could it be implemented in a brain, again bearing in mind the enormous computational complexity of the calculations? Would it not be redundant to evolve both network A and the apparently unconscious network B that calculates what network A is conscious of? Network B seems to need an internal model of network A, since it needs to know the logic gate at each node in order to calculate the qualia, as well as the current state. But if it has that, (why) do we need network A at all? Could it not simply accept network A’s inputs and use its model of A to calculate the ensuing states and qualia?

And how is network B supposed to learn the IIT rules that allow it to compute the consciousness of network A? According to IIT, the qualia of network A have no direct effect on its states or outputs, and so no effect on network B. Network B cannot learn anything directly from network A’s consciousness. So it can only have evolved to carry out the calculations it does because calculating and using the data that (according to IIT) define the qualia of network A also happens to be an evolutionarily optimal use of resources. If this is true, it must be true independently of whether the data do in fact correspond to consciousness. Again, we have an extraordinary coincidence: the high-level narrative that happens to be evolutionarily optimal to construct also, by chance, happens to be the narrative that we are conscious of, although the second plays no causal role in generating the first.

5. This is very difficult for any IIT-like theory to explain. One might speculate that our brain’s networks might perhaps be set up so that thoughts and activities associated with learned avoidance are associated with a class of IIT-qualia whose mathematical representations have one or more distinctive common features, while those associated with learned seeking are associated with another class with one or more different distinctive common features. Perhaps qualia with the first feature(s) are associated with experiences we class as “painful”, while those with the second are associated with experiences we class as “pleasurable”?

One obvious problem is that this *is* pure speculation: there is no evidence for it, nor any sketch of an argument as to why it might be true. It seems as though avoidance-class and seeking-class information processing can crudely be characterised by the output in a 1-bit channel (where 0 codes for avoid and 1 for seek). But it isn’t evident that the outputs of complicated information processes involving recursive and integrated information flows should generally have any close correlation with the characteristics of the qualia that IIT associates with the process.

Another is that it seems the associations of pleasure and pain could equally well have been reversed: avoidant and seeking behaviour could be associated with pleasure and pain qualia respectively, since the qualia have no effect on our behaviour. If this were the only other option, we might possibly accept that there was a 0.5 chance of the (broadly) harmonious correlation we experience, and that we were just fortunate. But this seems inadequate, since we can also ask why the mathematical features characteristic of avoidant and seeking behaviour aren’t associated respectively with sensations of seeing blue and green, or feeling hot and cold, and so on.

A possible response here could be to deny that pleasure and pain have intrinsic meaning. On this view, pleasure and pain are just words we associate with qualia evoked by avoidance and seeking, along with clouds of related words – comfort, beauty, fragrance, disharmony, agony, stench, and so on. To say pain is intrinsically something we prefer to avoid is a confusion: we have evolved to avoid certain physical states, which are associated with qualia that we give the name “painful”. We may have the illusion that our avoidance arises because pain qualia are intrinsically “bad” experiences, but in fact this is mistaken learned association: it is the states that we are avoiding, not the pain.

This is no doubt indeed arguable. After all, some argue that consciousness is pure illusion: the claim here is weaker – distinctive experiences exist, but the intrinsic qualities we assign to them do not. Even this partial deflation of experience, though, undercuts some of the motivation for IIT or any IIT-like EPTC, for which the Cartesian immediacy of experience is a central motivation. It simply feels incompatible with our direct experience to say that, while seeing blue and green are distinct sensations, they have no intrinsic qualities. Denying the reality of qualities seems like bleaching a qualia theory into quasi-behaviourism or quasi-illusionism.

A weaker and perhaps more tenable position is to accept that IIT-like EPTCs presently have to limit their ambitions to characterising information-theoretic aspects of experience. Such a theory could say whether two qualia are distinct, and perhaps give some measure of their similarity, while remaining silent about the intrinsic qualities of the associated experiences. This view allows that statements about the intrinsic qualities are meaningful. In particular, it accepts that it is puzzling that evolutionarily beneficial experiences are generally pleasurable. However, it admits that EPTCs, as presently conceived, cannot resolve the puzzle. The question then becomes whether, on this point, EPTCs suffer by comparison to other types of theory. We will return to this point below.

IIT⁺: BEYOND IIT

IIT's difficulties in explaining the evolution of consciousness motivate looking for theories with greater explanatory power. Specifically, they motivate us to consider theories in which some information about the qualia *is* available to the network. This can be achieved by postulating a modified dynamics in which the network state at time $t + 1$ is not obtained from that at time t by simply applying the Markovian transition matrix, but also depends explicitly on the mathematically defined qualia at times $\leq t$.

There are many ways in which this dependence could be implemented. For example, the states at time $t + 1$ could, with some perhaps small probability, encode aspects of the qualia by some modulation of the Markovian states at $t + 1$. Because we are interested in the potential explanatory power of *any* modified theory, we consider a stylised modification with maximal power. This allows us to examine what in principle could or could not be explained by a modified theory.

We thus define IIT⁺ to be a theory that, like IIT, applies to a discrete time network N with logic gate nodes. We take the network N to be a subset of B (which we think of as something like a human brain), and to include the nodes associated with the main consciousness. The network N functions in the standard Markovian way. However, we suppose that there are extra output channels from N that feed, along with the outputs of N , into the rest of the network comprising B . These channels output, at time $t + 1$, a complete (at least up to finite precision) description of all the qualia that IIT ascribes to N at time t . The model here is *not* of the type discussed in the previous section, in which some large subnetwork N' of B monitors N and carries out the complex calculations required to produce the qualia. We suppose rather that some modification of the known dynamical laws of nature applies in the context of networks that (according to IIT) are conscious and causes the description of the qualia to be directly output.

We emphasize again that this is a stylised model. There is no suggestion here that extra output channels of the type described actually exist in human or animal brains. We are interested, rather, in the possibility that some subtler modified dynamical law might apply, and that this hypothesis might be able to give a better explanation than IIT does of the evolution of consciousness.

Of course, postulating modified dynamics that depend on qualia states is a radical hypothesis, whether the modifications are subtle or blatant. Scientific conservatism and parsimony weigh quite strongly against it. However, we want to consider whether additional explanatory power might weigh in favour of keeping the hypothesis on the table. We consider the balance of these arguments later.

EXPLANATORY POWER OF IIT⁺-LIKE THEORIES

We now consider the five questions previously discussed from the perspective of someone seeking to develop an IIT⁺-like EPTC.

1. IIT⁺, like IIT, has no explanation for the existence of consciousness: it is a brute fact of nature.

2. For IIT⁺ to be relevant to any account of the evolution of consciousness, the physical correlates of consciousness must include the types of physical system that arise in terrestrial living organisms, and specifically in their information processing networks. If most physical systems have no physical correlates of consciousness, then it appears just a fluke that our precursors did. If most physical systems have physical correlates of consciousness (even if generally minimal), then it is less surprising that they should have.

Either way, given that physical correlates of consciousness arise in living creatures, IIT⁺ allows a sketch of an explanation as to why these correlates (and hence consciousness) would be selected for and organised into complex conscious systems. This is that because conscious networks output their qualia as well as their Markovian outputs, they have a key computational advantage. Effectively, they solve a problem of high computational complexity – calculating the qualia – in a single step. Plausibly, this computational advantage gave early organisms a small but useful evolutionary advantage, which increased over time as networks became more complex and the additional computational power was better exploited.

3. A direct description of evolutionarily critical features in the external and internal environment allows fast and well-calibrated responses. IIT⁺-type theories would suggest that this is why we have evolved information processing networks that generate consciousness of precisely these features, taking advantage of the modified dynamics that allow information about the qualia to be fed directly into the networks defining our brain.

Although highly advantageous, this presumably has some costs, which depend on the details of the IIT⁺-type theory. We expect some cost in building a network complex enough to generate qualia that describe many evolutionarily critical features in our environments, in processing the information about consciousness that enters the extended network, and in creating the channel that allows it to enter the network. Presumably the channel cost increases with its bandwidth.

These cost-benefit trade-offs presumably explain why while consciousness has – broadly speaking – evolved to include the most evolutionarily salient information, it excludes most other information.

If some of the content of our consciousnesses nonetheless has relatively little value this is presumably because there is no attainable mechanism that identifies and selects precisely the evolutionarily salient and nothing else. As spandrels illustrate, evolution is a probabilistic process and an imperfect optimiser. However, it is striking that our consciousnesses don't include much (if any) *bizarrely* irrelevant content. We can give a plausible account of how most (perhaps even all?) of the information we're conscious of is at least *related* to evolutionarily salient information. For example, although the content of our dreams often includes relatively little information directly related to our immediate external environment, topics related to evolutionary benefit – such as status, the evasion of threats, social relationships, food or sex – are not uncommon themes. Such dreams (and others) have effects on our waking behaviour that are (at least plausibly) evolutionarily beneficial – and (I would suspect) did so even before Freud and earlier oneirocritics.

On an IIT⁺-like hypothesis, the cost of increasing bandwidth explains why we get an executive summary of our situation, not every detail. For example, we know whether we feel hot or cold, and may feel impelled to act on this. However, we aren't directly aware of how dilated our capillaries are, information that is necessary for unconscious temperature regulation but adds no useful information that would affect our deliberate responses. Similarly, we know whether our muscles are tired, but not the percentage of lactic acid they contain or their pH.

4. If an IIT⁺-like theory is correct, we can speak, and also communicate non-verbally, about consciousness and qualia because they directly input information into our neural networks, through some mechanism created by modified dynamics (whose details are presently unknown).

We have developed this ability because it is evolutionarily valuable to produce a communicable narrative about our qualia, since the selection principle discussed in point 3 implies this narrative is focussed on evolutionarily critical topics. Part of the value may arise from being able to communicate this conscious narrative — through behaviour as well as verbally — to ourselves, facilitating reinforcement, memory, reflective processing and learning. There is a major further advantage in being able to communicate the narrative to others, to facilitate social behaviour, including cooperation and conflict management. This may well include inter-species communication. For example, a dog's signals of anger or contentment – which we do well to heed! – may be produced most quickly and reliably from a high-level narrative about its qualia.

5. As in the discussion of IIT above, there are two separate issues here: Can we explain why the types of qualia associated by IIT⁺ with beneficial and harmful events and actions have distinct characteristic signatures? And can we explain why the former are associated with pleasurable experiences and the latter with painful? IIT⁺, unlike IIT, can give a plausible sketch answer to the first. As discussed above, IIT⁺ suggests that evolution makes use of the qualia associated with network processes as a computational resource, provided by the modified dynamics, to help with quick decisions and actions. In order to maximise the speed of response, evolution has thus arranged the network processes so that qualia associated with processes leading to avoidance are easily characterizable, as are the qualia associated with processes leading to seeking, and the two are easily distinguishable. Once the qualia are input to the network, these features can thus very swiftly be identified, allowing the appropriate behaviour to be swiftly initiated.

On the second issue, though, the questions raised for IIT also apply to IIT⁺. Why should the intrinsic experience associated with avoidance-inducing qualia be painful rather than pleasurable? Why should avoidance- or seeking-inducing qualia not be associated with distinct neutral experiences, such as seeing blue or green?

The same responses are possible. Again, the most tenable seems to be accepting that, for now, even EPTCs with modified dynamics have to limit their ambitions to characterising information-theoretic aspects of experience. IIT⁺-type theories have greater explanatory power than IIT-type theories on the information content of consciousness, but still struggle to explain anything about the intrinsic nature of experience.

CONCLUSIONS

IIT⁺-type theories can explain far more about the evolution of consciousness than IIT-type theories. Both types of theory take the existence of consciousness as a given. Both have to assume that the physical correlates of consciousness include physical systems that naturally arose in organisms at some – most likely early – point in the evolutionary process. Given this, IIT⁺ theories, unlike IIT-type, allow a sketch explanation of why physical systems associated with complex consciousnesses were selected for, why human consciousness is focussed on an evolutionarily salient narrative about our physical states and environments, why we can speak accurately about our consciousness, and why distinctive types of conscious sensation are associated with evolutionarily beneficial and harmful processes. These are very substantial problems, not only for IIT, but for most accounts of consciousness.

The explanations come at a high price. IIT⁺ adds a further speculative hypothesis – that the known physical dynamics are modified by some direct effect of qualia, represented mathematically by information-theoretic quantities along the lines of those proposed for IIT – to IIT, which is already a speculative theory. Scientific conservatism and parsimony weigh strongly against postulating new dynamical laws, particularly since the known physical laws are very well tested in many regimes and fit into a mathematically unified structure that does not appear to admit easy and elegant tweaks. While modified dynamical rules have been seriously considered in theoretical physics – for example to resolve the quantum measurement problem [24, 25] and to unify quantum theory and gravity [26, 27] – it is generally seen as problematic that they are in apparent tension with the relativity principle and with conservation laws.

On the other hand, abandoning Darwinian evolution is an even less attractive alternative. Maintaining that Darwinian evolution applies to the material world, while accepting that the universe has laws seemingly designed from the beginning to equip humans with a narrative consciousness relevant to our survival, is also a very unsatisfactory position. For these reasons, without underestimating the huge difficulties and drawbacks, we find IIT⁺-type theories more scientifically attractive than IIT-type.

These are not the only possible options, of course. Given the difficulties with both alternatives, one should ask whether perhaps the entire project of an EPTC is misguided. Perhaps, in the light of this discussion, other theoretical approaches to consciousness seem more plausible, despite their own drawbacks? An additional reason for asking this question is that neither IIT-type nor IIT⁺-type theories seem able, even on our most charitable constructions, to address the intrinsic nature of experience.

Can other theoretical approaches to consciousness do better, though? The cases against functionalism and illusionism seem very strong. Functionalism effectively defines mind-states in terms of their role in chains of causation that start and end in the material world. Illusionism denies the existence of mind states altogether. Neither helps explain the intrinsic qualitative nature of experience. Neither requires new dynamics. In principle this would be a great advantage over IIT⁺-type theories. But if – as it seems to many – it is achieved by not taking the phenomenon of conscious experience seriously, the comparison is unfair: one can avoid all difficulties in science by ignoring inconvenient data.

Dualism offers some explanation for some of the aspects of consciousness we have discussed. If minds are independent of bodies but have causal influence over them, then we at least have a potential mechanism allowing us (our bodies) to speak about the experiences of our minds. But dualism requires a far more radical revision of known science than IIT⁺-type theories. Moreover, a dualist account of the evolution of consciousness seems utterly fantastic. We have to suppose that minds somehow attached themselves to bodies, probably at an early stage of the evolutionary process. Either minds then somehow reproduced, by some mysterious mechanism, in parallel with bodies, or else each new body somehow, at an early stage of its development, attracts a new mind. We also seem to need to suppose that minds are able to attach to bodies with a variety of types of association between mind-states and body-states, and that the nature of this attachment is itself subject to evolution, which selects for mind-body combinations in which pleasure qualia (which minds, on this view, tend naturally to favour) are associated with evolutionarily beneficial states.

Idealist accounts of the evolution of consciousness also seem both fantastic and challenged by the evidence for Darwinian evolution. On an idealist view, our experiences are consistent with a highly ordered external world containing stars, galaxies, and planets, including Earth, and with a standard account of the evolutionary history of life on Earth – but none of these external objects and processes are actually real. It is extraordinarily mysterious that our experiences should be structured in this way. The evolution of life, or consciousness, on an idealist view, is not a fact to be explained, but an erroneous theoretical construction based on – but wrongly extrapolating from – experience.

Perhaps the alternative that gains most credence from our discussion is that our current philosophical, scientific and mathematical concepts are inadequate to explain either the intrinsic qualities *or* the information content of consciousness. Of Chalmers' options, "None of the above" seems the hardest to understand – by definition, it involves some so far unimagined line of thought – but also the one whose (im)plausibility is hardest to evaluate. In 1900, no list of options for explaining unresolved problems in physics would have pointed towards quantum theory or general relativity. We may be far less well conceptually equipped now for a science of consciousness than classical physicists were then. Still, one can only work with what one has.

It seems then that IIT⁺-type theories deserve to remain on the table, despite all their problems and the very strong *a priori* weighting against them as scientific hypotheses. Every other current option for explaining the features of our consciousness seems at least comparably implausible. Alas, the IIT⁺ hypothesis is under-specified, giving no clear hint of how or where to look for evidence. Perhaps the best hope for progress thus comes from research that is independently motivated. For example, finding stronger evidence that an IIT-type theory approximately characterises aspects of experience would add to the motivation for investigating IIT⁺-type theories, given the strong case for preferring IIT⁺-type theories to IIT-type. Successful brain simulation projects for living organisms would

bolster the case against IIT⁺-type theories, particularly if one could verify that every step of the simulation follows standard network dynamics; conversely, unexplained discrepancies between simulated and organic networks would favour IIT⁺.

ACKNOWLEDGEMENTS

This work was supported by an FQXi grant and by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation.

References

-
- * Electronic address: A.P.A.Kent@damtp.cam.ac.uk
- [1] David Chalmers. Idealism and the mind-body problem 1. In *The Routledge handbook of panpsychism*, pages 353–373. Routledge, 2019.
 - [2] Masafumi Oizumi, Larissa Albantakis, and Giulio Tononi. From the phenomenology to the mechanisms of consciousness: integrated information theory 3.0. *PLoS Comput Biol*, 10(5):e1003588, 2014.
 - [3] Adam B Barrett. An integration of integrated information theory with fundamental physics. *Frontiers in psychology*, 5: 63, 2014.
 - [4] S. Aaronson. Why I am not an integrated information theorist (or, the unconscious expander). URL <http://www.scottaaronson.com/blog/?p=1799>.
 - [5] Michael A Cerullo. The problem with Phi: a critique of integrated information theory. *PLoS Comput Biol*, 11(9):e1004286, 2015.
 - [6] Adam B Barrett and Pedro AM Mediano. The Phi measure of integrated information is not well-defined for general physical systems. *Journal of Consciousness Studies*, 26(1-2):11–20, 2019.
 - [7] Giulio Tononi. Integrated information theory. *Scholarpedia*, 10(1):4164, 2015.
 - [8] Giulio Tononi. Consciousness as integrated information: a provisional manifesto. *The Biological Bulletin*, 215(3):216–242, 2008.
 - [9] Giulio Tononi. Why Scott should stare at a blank wall and reconsider (or, the conscious grid). URL [availablefromwww.integratedinformationtheory.org](http://www.integratedinformationtheory.org).
 - [10] Erik P Hoel, Larissa Albantakis, and Giulio Tononi. Quantifying causal emergence shows that macro can beat micro. *Proceedings of the National Academy of Sciences*, 110(49):19790–19795, 2013.
 - [11] David J. Chalmers and Kelvin J. McQueen. Consciousness and the collapse of the wave function. *arXiv preprint arXiv:2105.02314*, 2021.
 - [12] Elias Okon and Miguel Ángel Sebastián. A consciousness-based quantum objective collapse model. *Synthese*, pages 1–21, 2018.
 - [13] Elias Okón and Miguel Angel Sebastián. How to back up or refute quantum theories of consciousness. *Mind and Matter*, 14(1):25–49, 2016.
 - [14] Elias Okón and Miguel Angel Sebastián. The subjective-objective collapse model: Virtues and challenges. In S. Gao, editor, *Consciousness and Quantum Mechanics*. Oxford University Press, Forthcoming, expected 2021.
 - [15] Kobi Kremnizer and André Ranchin. Integrated information-induced quantum collapse. *Foundations of Physics*, 45(8): 889–899, 2015.
 - [16] Giulio Tononi. *Phi: A Voyage from the Brain to the Soul*. Pantheon, 2012.
 - [17] Marcello Massimini and Giulio Tononi. *Sizing up consciousness: towards an objective measure of the capacity for experience*. Oxford University Press, 2018.
 - [18] Adrien Doerig, Aaron Schurger, Kathryn Hess, and Michael H Herzog. The unfolding argument: Why IIT and other causal structure theories cannot explain consciousness. *Consciousness and cognition*, 72:49–59, 2019.
 - [19] Naotsugu Tsuchiya, Thomas Andrillon, and Andrew Haun. A reply to “The unfolding argument”: Beyond functionalism/behaviorism and towards a truer science of causal structural theories of consciousness. 2019.
 - [20] Martine Nida-Rümelin. Qualia: The knowledge argument. *The Stanford Encyclopedia of Philosophy (Summer 2015 Edition)*, 2015. URL <https://plato.stanford.edu/archives/sum2015/entries/qualia-knowledge/>.
 - [21] Torin Alter. The knowledge argument against physicalism. *The Internet Encyclopedia of Philosophy*, 2017. URL <http://www.iep.utm.edu/>.
 - [22] Torin Alter and Derk Pereboom. Russellian Monism. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2019 edition, 2019.

- [23] David Chalmers. Panpsychism and panprotopsyism. *Consciousness in the physical world: Perspectives on Russellian monism*, pages 246–276, 2015.
- [24] Gian Carlo Ghirardi, Alberto Rimini, and Tullio Weber. Unified dynamics for microscopic and macroscopic systems. *Physical Review D*, 34(2):470, 1986.
- [25] Gian Carlo Ghirardi, Philip Pearle, and Alberto Rimini. Markov processes in Hilbert space and continuous spontaneous localization of systems of identical particles. *Physical Review A*, 42(1):78, 1990.
- [26] Lajos Diosi. A universal master equation for the gravitational violation of quantum mechanics. *Physics Letters A*, 120(8): 377–381, 1987.
- [27] Roger Penrose. On gravity’s role in quantum state reduction. *General Relativity and Gravitation*, 28(5):581–600, 1996.
- [28] Arguably, we confront a similar issue even if we build a theory only on our own experience: we still have to rely on video or other external records of our own reports, or on our memories, which perhaps may only be internal records of internal reports.
- [29] We will not discuss in detail here the complication that this would be at least slightly unrealistic even if IIT94.0 existed and were correct. There would still be exceptional cases, where the device tells us that the subject has formed the intention of lying and we are persuaded that they are indeed lying, or where the subject’s brain is not functioning normally. To deal with this, we have to suppose that we also find the explanations of these anomalous cases sufficiently parsimonious and convincing.
- [30] Perhaps this is even true of non-scientific accounts. We are familiar with discussions of what experiences are *like* in literature, art and ordinary conversation, but these seemingly always ultimately relate one set of experiences to another. Can language ever describe what experiences *are*?
- [31] As Rabi asked about the muon, who ordered *that*?