Can lists of requirements help consciousness navigate its epistemological quandaries?

Abstract

Frustration has been growing with mainstay epistemological methods of logical deduction and experimental falsification for assessing theories of consciousness. This paper explores one among several alternatives being proposed: the listed requirements epistemology. A literature search identifies five papers that explicitly list requirements for assessing consciousness theories. These five lists are analysed as a promising starting point, but as yet insufficiently comprehensive to do the method justice. The longest list has 11 items, but 19 unique items are identified across the five lists and a taxonomic analysis by category further surfaces at least 30 potential candidates. Four limitations of the method are discussed, arguing that it is best treated as one practical tool as part of a broader strategy for rigorously assessing consciousness theories and identifying avenues for future research. The conclusion describes a workplan for a sufficiently complete working taxonomy to support the field's collective endeavours.

Key words: Epistemology; Correlates of Consciousness; Taxonomy; Requirements; Thought Experiments

Introduction

Researchers are increasingly concerned about a proliferation of theories of consciousness (ToCs), with a recent debate focused on what epistemology might help choose between them.

Alternative approaches, summarised in section one, such as construct-first epistemology, exclusionary empiricism, and inference to the best explanation, have recently been championed in response to frustrations over mainstay epistemologies of logical deduction and experimental falsification. These traditionally favoured epistemologies, while having important roles to play, do not yet appear to be succeeding in driving consensus or winnowing out candidate ToCs. Indeed, the list of plausible theories is growing, with one partial list identifying 22 ToCs (Seth & Bayne, 2022) and many more alternatives and variants not included. For instance, their scope excluded non-neurobiological ToCs (e.g. dualism, idealism, panpsychism variants) and McFadden (2023) identifies nine distinct variants of electromagnetic field theories that were subsumed into a single ToC in the former list of 22.

One benefit of the plurality of theories has been growing taxonomic efforts to clarify and categorise ToCs, e.g. by the aspects they seek to explain, their mechanisms/modes of explanation, or their views of consciousness functions (e.g., Signorelli et al., 2021; Sattin et al., 2021; Niikawa et al., 2022). Unlike these related efforts, the taxonomy items discussed in this paper are not different ToCs but rather the requirements which candidate ToCs are expected to meet. All theories that eventually aspire to a complete explanation of human consciousness would be in scope for such requirements, including both our subjective first person perspective and all the different conscious experiences we can have.[1]

This paper explores one recent strand of epistemological thinking, termed here the 'listed requirements epistemology', inspired by Doerig et al. (2021a:41) and the discussion, both critical and constructive, that was provoked by their "checklist of criteria" for assessing and comparing different ToCs. Doerig et al. (2021a) identify four criteria but explicitly describe their efforts as an incomplete stepping stone (p54). Given this incompleteness, section two reports on what other similar lists have been proposed, identifying a total of five lists that explicitly have comparable goals.

Section three reviews the five lists identified, concluding that they have material omissions (even collectively) and insufficient taxonomisation. For instance, the longest list has 11 requirements, but there are 19 unique items across these five lists alone with many more credible candidates from the broader literature. This lack of coverage may partly stem from an insufficiently collaborative, interdisciplinary perspective. Only one of the five lists cites one of the others. All five lists also favour, explicitly or implicitly, a particular set of ToCs, raising concerns over cherry-picked requirements and questions over their utility for comparing all ToCs from an approximately neutral perspective.

These practical shortcomings can be addressed through future research. But would such research be worthwhile? Section four discusses four objections to the listed requirements epistemology. We argue that listed requirements has the potential to support the field through its current epistemological quandaries, acknowledging the objections that limit it to pragmatic usage as a collaborative tool used alongside other epistemological strategies. Its primary value lies in providing a flexible framework within which to log all the different thought experiments, logical paradoxes, empirical evidence, and introspective phenomena that are typically only cherry-picked by researchers to argue for one ToC over another. It is easier to gain consensus that a

[1] These efforts may shed light on consciousness in non-human systems, but that is not the immediate focus. Consciousness-generating mechanisms are not necessarily identical across system types, notwithstanding the utility and plausibility of such identity assumptions. The same uncertainty applies in principle within systems, such as between different humans. Arguments for solipsism and radical scepticism notwithstanding, this paper proceeds with the assumption that mechanisms are similar enough between people that we can usefully initiate a taxonomic exercise of gathering ToC features and phenomena to account for.

phenomenon should be accounted for (even if only by denying the evidence for it or what it appears to mean), than to gain consensus on which account is most compelling.

The conclusion presents a structured workplan for addressing the practical shortcomings in the literature to date, so as to fulfil the tool's potential. Collectively, this paper and proposed workplan respond to calls from researchers such as Wiese (2018), who asks for a more mature science of consciousness that includes a systematic listing of its features to assess better which theories might integrate them, and Del Pin et al. (2021), who call for greater collaborative efforts to specify ToC criteria.

1. The epistemological quandaries of consciousness studies

Concerns have been growing that consciousness studies needs new tools for assessing which ToCs perform better than others.

Fazekas et al. (2024:1) argue that "the discipline struggles with an abundance of alternative theories", proposing a construct-first empirical epistemology as a solution. Paßler (2023:3) develops an "exclusionary" approach, as a complement to the identification approach commonly used in neural correlates of consciousness which he describes as leading to "almost every part of the brain" being a neural correlate candidate. Del Pin et al. (2021) propose a six step empirical procedure based on the epistemological principles of strong inference. Kirkeby-Hinrup and Fazekas (2021) argue for "inference to the best explanation", as a general framework that can incorporate Bayesian estimation, formal confirmation theory, and other methods. Doerig et al. (2021a:42) develop a checklist of criteria in order to tame the "bewildering number of ToCs".

These epistemological principles have developed out of frustration that the mainstay epistemologies in the field are not helping to generate consensus or to reconcile contrasting or overlapping evidence. As Kirkeby-Hinrup and Fazekas (2021:2) state, epistemologies of "logical deduction or elimination of hypotheses through falsification" are typically preferred, but may not fully meet the needs of consciousness studies at this time.

"Logical deduction" – and the narrative reasoning and thought experiments it relates to – has long been favoured in analytical philosophy. Recent examples include taking ideas to their logical consequences to exclude alternatives, such as the arguments against materialism and dualism in Goff (2019), and the metaphysical and semantic clarity used by Lee (2023) to argue for degrees of consciousness. Despite the claimed inevitability of their individual conclusions, the methods of logical reasoning have been used back and forth for millennia on the topic of consciousness. For any given argument, there are ample opportunities to reject a definition, a premise, a logical

inference, or the implications of a conclusion. These methods help us to be clear about what each position entails, but seem to have little lasting weight in driving consensus around a single specific position.

“Falsification” is typically claimed as the favoured approach of empirical science. Empirical investigation more generally has proved powerful in refining our understanding of the brain and generating empirical evidence that ToCs must account for. However, experimental approaches are unlikely to rule out all but one ToC any time soon. Fazekas et al. (2024) argue that experimental evidence, even carefully designed adversarial collaborations, typically place pressure on the auxiliary assumptions that link a ToC to empirical observations, rather than the core assumptions that the ToC proposes to account for consciousness. Such flexibility in interpretation leads to easy cherry-picking and confirmation bias in reviewing experimental evidence (Yaron et al., 2022).

The first person aspect of consciousness may in fact be outside the direct reach of third person empirical evidence, even if indirect scientific evidence should prove indispensable to the quest elsewhere. Cohen and Dennett (2011) argue that even their perfect experiment would be unable to isolate first person experiences. Indeed, leaning solely on experimental evidence arguably rules out certain ToCs prematurely. Goff (2019) has drawn on Russell and Eddington to argue that the natural sciences can only tell us the properties, behaviours, and interactions of physical phenomena, not the actual essence they are made out of. If consciousness is a feature of the essence of matter in some form, empirical experiments cannot provide the full epistemological solution to it.

Doerig et al. (2019) have argued integrated information theory and similar causal structure theories are also unfalsifiable. These theories may not be true, but ruling them out based on the limited reach of one particular epistemology seems as strong an assumption as any used to defend them in the first place. Not all statements of potential scientific value are falsifiable nor does a lack of falsifiability mean something does not exist – multiple philosophies of science exist under which meaningful scientific progress can be made (e.g. Dardashti et al., 2019; Griffin, 2012).

These limitations of logical deducation and experimental falsification have led to growing interest in alternative or complementary epistemologies. We turn now to identifying instances of one particular suggestion: the ‘listed requirements’ epistemology.

2. Identifying lists of requirements

Four additional lists with similar goals to Doerig et al. (2021a) were identified by a structured search and citation tracing strategy. The search sought only to identify papers where such enumerated list making was a major part of the presentation and where the list was comparable in target and spirit to their ToC checklist.

A ten year academic literature search was conducted in the English language in January 2024 on the Scopus database, looking for publication dates from 2014 to 2023 (inclusive) where the title had both a topic relevant keyword and a list relevant keyword. The topic relevant keywords were consciousness, qualia, and phenomenal; the list relevant keywords were criteria, account, requirements, considerations, explananda, restrictions, desiderata, taxonomy, and list (with both singular and plural forms captured). The search produced 136 results. Of the papers related to consciousness, common reasons for exclusion were focusing on ethical criteria or moral personhood, individual empirical results, or exegetical papers, as well as discussions of disorders of consciousness, clinical measures of consciousness, or neural correlates that did not produce enumerated lists specifically framed as requirements that a theory of consciousness should account for.

The Scopus database was then used to search for all papers citing those four lists, surfacing one further list meeting the criteria for inclusion (McFadden, 2023), which was itself then citation traced (with no further lists sourced). Many borderline or adjacent papers were also identified in this search strategy, providing conceptual insights into the issue of consciousness explananda or discussing findings that might be usefully incorporated into a list of requirements. This search strategy is unlikely to have identified every possible effort to itemise lists of ToC requirements but provides a sufficient base for assessing the feasibility of the listed requirements epistemology and how the broader literature might be better engaged.

Tables 1-5 present the lists from each paper in date order, including detail from the paper for how they present their lists that makes them appropriate comparators to the checklist from Doerig et al. (2021a) and any preferred ToC they endorse in the paper.

Two other papers are borderline candidates for inclusion, but are excluded since they do not provide a single enumerated list of requirements in the same spirit of the five papers selected. Nonetheless, both papers would be useful in the future exercise called for in the conclusion.

Table 1. Requirements list from Manzotti and Chella (2014)

#	Short name	Paraphrased summary
***Context of list**: Using "a set of feasible criteria" (p403) to examine theories to see if they "satisfy the proposed criteria" (p419) and "match our empirical evidence about the phenomenology of consciousness" (p421).* ***Theory motivation**: A "causal" theory of consciousness, where two events become integrated whenever they happen to cause a joint effect (p416).*		
1	Physical grounding	Consciousness cannot be just in the model that one uses to describe a phenomenon; it must be realised into some physical aspect of the phenomenon (e.g. causal relations, energy transference, or some conservation principle).
2	Causal efficacy	If a mind and its physical components are to be integrated, it must make a causal difference that they are integrated (else we could simply set aside the mental component).
3	No circularity	A theory must be able to explain what mental-physical integration is without referring to the notion of integration.
4	Scaling	A theory must no state that the mechanism for integration applies arbitrarily just to highly complex neural networks; it must also consider what happens in simple cases.

Table 2. Requirements list from Tononi (2017)

#	Short name	Paraphrased summary
***Context of list**: "identifying the essential properties of experience (axioms of phenomenal existence) and inferring the required properties of its physical substrate" (p621) [These requirements] "must be satisfied strictly for something to exist as an intrinsic entity [...] being an intrinsic entity, properly defined, is one and the same thing as being conscious." (p623).* ***Theory motivation**: Integrated Information Theory (IIT).*		
1	Intrinsic	Consciousness has intrinsic existence (i.e. the phenomenology of existence is real) – leading to the postulate that the corresponding system that is theorised to generate it (or correspond to it) must have cause-effect power.
2	Compositional	Phenomenal experience is structured out of multiple phenomenological distinctions – postulating the theorised system must have structure.
3	Informational	Consciousness has specific information-bearing contents – postulating that, if reordered, the theorised system's cause-effect patterns must change.
4	Integrated	Consciousness is unified into an irreducible experience (e.g. seeing a blue book cannot be fully reduced to seeing a colourless book plus the colour blue) – postulating every part of the theorised system must be able to both affect and be affected by the rest of the system.
5	Exclusive	Consciousness has a definite content and spatio-temporal grain fixed for any given experience – postulating the constraint on a theory to have a maximally irreducible cause-effect structure.

Table 3. Requirements list from Gambini & Pullin (2019)

#	Short name	Paraphrased summary
Context of list: *A theory must "account for a set of properties of consciousness" (p10).* ***Theory motivation:*** *Quantum theories of consciousness.*		
1	Unity	A theory must explain how we can simultaneously experience multiple items from a single perspective, unified both in the moment and over continuity in time.
2	Robust	A theory must explain how much brain damage does not destroy all consciousness; definitive loss of consciousness only occurs under severe brain damage.
3	Thalamocortical	A theory must explain the medical association of consciousness with thalamocortical functions and why such an association does not appear in other strongly interacting neural systems like the cerebellum.
4	Divided selves	A theory must explain the mechanisms by which some types of brain damage lead to kinds of divided consciousness.
5	Varied	A theory must explain how its mechanism for consciousness translates, at least in humans, into a capability of providing very rich and varied experiences.

Table 4. Requirements list from Doerig et al. (2021a)

#	Short name	Paraphrased summary
Context of list: *"[A] checklist of criteria that, we propose, empirical ToCs need to cope with" (abstract).* ***Theory motivation:*** *Empirical theories, but otherwise claim to have developed the list with no specific theories in mind, but note Fahrenfort and van Gaal (2021) argue the list is motivated implicitly by functionalism.*		
1	Paradigm cases	Where the same process or perception can occur both consciously and unconsciously (e.g. binocular rivalry, masking studies), the theory should explain what is different in a general sense across all such cases.
2	Unfolding argument	Cope with the "unfolding argument" by which any algorithm can be represented in a feedforward network without any changes in inputs or outputs (i.e. functionally identical, potential differences in speed notwithstanding); the authors demand an experiment for resolving this issue when it applies.
3	Small/large networks	Cope with the "small network argument", i.e. either accept that very small systems might be conscious and explain why this is plausible (the authors reject its plausibility) or explain why small systems are not conscious and yet other systems are despite having similar abstract features but expressed over larger scales. Further, if multiple subsystems could be conscious in a given "large network", specify which one it is (which the authors relate to the Combination Problem).
4	Other systems	Cope with the "multiple realisation argument", making clear distinctions between what systems are conscious and others are not. (This requirement was weakened in Doerig et al. [2021b] to require that ToCs should merely be explicit about whether their theory intends to or is capable of making such distinctions).

Table 5. Requirements list from McFadden (2023)

#	Short name	Paraphrased summary
Context of list: *"Criteria through which ToCs could be checked and compared" (p1) and "eleven tests of ToCs" (p2).* ***Theory motivation:*** *Electromagnetic field theories of consciousness.*		
1-4	(...)	Explicitly the same as Doerig et al. (2021a) from table 4.
5	Unity	Explain what gives rise to the unity of consciousness problem or otherwise account for the binding problem, by which distinct features can be co-experienced from a single first person perspective.
6	Neural correlates	Address why neural correlates of consciousness do in fact correlate with consciousness, e.g. absence of consciousness in certain parts of brain (cerebellum) or certain brain states (grand mal, epileptic seizures), and timing features (e.g. the psychological refractory period for conscious awareness, attentional blink, postdictive effects on conscious perception).
7	Measurable	Address the measurement problem, i.e. identify the degree to which different organisms or structures are conscious.
8	Serial processing	Account for why we experience a conscious mind that can only do one main thing at a time even while our brain non-consciously conducts many activities in parallel.
9	Intelligence	Distinguish intelligence from consciousness and account for how we can have conscious experiences of intelligent reasoning.
10	Non-epiphenomenal	Account for the emergence of consciousness through natural selection.
11	Falsifiable	Make novel testable predictions.

The first borderline paper is Anokhin (2021) who argues for a particular theory of consciousness called the cognitome, effectively identifying key "cognitive processes in the neural hypernetwork" (p934) with consciousness. Anokhin (2021) develops a "set of criteria which must be met by a complete scientific explanation of consciousness" (p921). However, the criteria appear to be distributed over several lists, which may be separately valuable for a future exercise, but results in a lack of clarity for this review. For instance, three principles are specified (p925), along with ten properties of consciousness listed in table 1 and a separate list of eight questions (p924). The eight questions are used to assess specific ToCs but are not explicitly phrased as criteria or requirements.

The second borderline paper is Del Pin et al. (2021). In the abstract, they state "we set out the features that a theory of consciousness should account for", aligning with the search strategy for the paper. However, the main text provides a few items "for purely exemplary reasons" (p6), being a subset of items from a reference publication designed to describe properties of consciousness rather than ToC requirements in the explicit sense of the other lists presented.

3. Assessment of lists to date

The five lists all address the same broad topic, albeit with different caveats. All lists target all attempts at a complete ToCs, noting that any complete ToC would by definition account for the empirical and phenomenological observations that have been related to consciousness. However, they do so with different circumscriptions.

Manzotti and Chella (2014) address only logical requirements about physical integration. Tononi (2017) addresses only phenomenological requirements identified via introspection. Gambini and Pullin (2019) specify properties of consciousness that must be accounted for, but do not necessarily rule out other types of requirement a ToC might be called to meet. It is fair to interpret such scholars as assembling lists of ToC requirements, perhaps even comprehensive lists within particular categories of requirement, but not to assume they believe no other categories of potential requirements might be useful.

An inclusive approach to developing a comprehensive list would look to combine different categories of requirement, whether across the three categories already above or other categories. Doerig et al. (2021a) begin such an inclusive approach by including different types of criteria for assessing all ToCs that claim explanatory power over empirical phenomena. For instance, they include evidence from neuroscience and psychology (item #1), from mathematical arguments (#2), and analytical philosophy (#4). However, they also describe their checklist as incomplete. McFadden (2023)

explicitly builds on Doerig et al. (2021a) to identify a list of 11 tests and no longer explicitly describes the list as incomplete.

Possible objections to an inclusive approach are discussed in section four. Let us note here just that such inclusivity does not prevent theorists from arguing against the validity of certain requirements when comparing their preferred ToC against a consolidated list. For instance, they might deny the experimental validity of the evidence a claimed requirement is based on, deny that introspected experience means what it appears to imply, or deny the requirement is correctly defined or logically coherent. Unless such denials are univocal, it is safest to include the suspect requirements and allow theorists to decide whether to deny or account for them. More generally, developing adequately comprehensive listings of requirements appears to be a strategy followed by the researchers collating these lists to date. On that basis, it is reasonable to ask how complete the existing lists are, both individually and as a collective.

3.1. Incomplete lists

Table 6 collates and deduplicates the 29 items across the five lists. Identifying duplicates is inexact, as the phrasing is often slightly different between lists. Nonetheless, an initial rationale for deduplication is where the requirements are similar enough that a ToC account for one should also address another. Should this later turn out inaccurate, the item could be split up using more precise phrasing to draw out the relevant differences. The two most common requirements appear in three lists each, being the need to account for the unity of consciousness and the need to apply their ToC mechanisms at all applicable scales.

Table 6. Consolidated ToC requirements from tables 1-5

#	Requirement	Listing 1*	Listing 2	Listing 3
1	Unity & integrated nature of experience	T	GP	M
2	Scaling & small/large networks	MC	DSH	M
3	Causal efficacy & non-epiphenomenal	MC	M	
4	Compositional/varied contents	T	GP	
5	Thalamocortical linkage (an e.g. of neural correlates)	GP	M**	
6	Robust to some brain damage (an e.g. of neural correlates)	GP	M**	
7	Paradigm cases (of conscious/unconscious processes)	DSH	M	
8	Unfolding argument	DSH	M	
9	Multiple realisability/other systems	DSH	M	
10	Physical grounding	MC		
11	No circularity	MC		
12	Intrinsic	T		
13	Informational	T		
14	Exclusive	T		
15	Divided selves	GP		

16	Measurable	M		
17	Serial processing	M		
18	Intelligence	M		
19	Falsifiable	M		

**Listings in publication date order; MC: Manzotti and Chella (2014); T: Tononi (2017); GP: Gambini & Pullin (2019); DSH: Doerig et al. (2021a); M: McFadden (2023); ** Listed via the category of neural correlates.*

Taking the five lists so far at face value, the longest list must be incomplete, as it identifies 11 items, compared to the 19 unique items in table 6. However, the incompleteness is markedly more significant than implied by just assessing these five lists. A broader reading of the literature identifies many more candidate requirements that fall into the same category as some of the items in the consolidated list so far. None of the lists propose a rationale for why some potential requirements would be out of scope, while others are included. In the absence of such a rationale, it seems we would have to expand the list to include category cohabitants, as well as considering what other categories might be missing entirely.

This lack of coverage may partly stem from an insufficiently collaborative, interdisciplinary perspective. The five lists identified do not, with one exception, cite each other as part of a collaborative (or even combative) exercise to list ToC requirements for theories of consciousness, despite such listings being a claimed goal of each paper and despite a citation tracing strategy biased towards papers that cite each other. The only valid exception to this rule is McFadden (2023) who cites and builds collaboratively on the list from Doerig et al. (2021a).[2]

All five lists also favour, explicitly or implicitly, a particular set of ToCs, raising concerns over cherry-picked requirements and questions over their utility for comparing all ToCs from an approximately neutral perspective (see the "theory motivation" comments in tables 1-5).

3.2. Insufficient taxonomisation

We discussed above that some lists had different circumscriptions, such that they would naturally identify different categories of item. Multiple categories provide a further taxonomic method for assessing completeness.[3] Labelling the different categories that items might be assigned to can draw attention to items missing from specific categories or potential additional categories to investigate. Applying this

[2] A further technical, but invalid exception is citations of Tononi's work. All the papers cite some of Tononi's work (albeit not the 2017 one), but only to discuss IIT as a theory, rather than to build on his five axioms/postulates that drive the requirements list of interest for this paper.

[3] Some items might belong to more than one potential category along different axes, in which case we have a choice over when to append and organise by multiple labels and when to privilege one set of hierarchical categories where each item appears exactly once. Categories can also be analytical in nature, in that they aid discussion and analysis, rather than necessarily lining up to fundamental ontological differences in reality. Such taxonomisation is a pragmatic, subjective, and theory-laden exercise to support analysis, rather than a fixed algorithm with only one right answer.

taxonomic technique across the 19 items from table 6 further illustrates the incompleteness of lists to date and points towards a future workplan. We group this exercise across four overarching categories: introspective explananda, analytical explananda, empirical explananda, all being phenomena or puzzles that a ToC must provide an account for (or a denial for), as well as theory desiderata, being requirements that would typically lead to a more viable, more useful ToC.

Introspective explananda

Item 1 from table 6, unity, is one of several features of conscious experience, typically reported introspectively, i.e. what it feels like to be conscious. Other items can also be mapped to the label of "introspectively experienced features". For instance, we can experience being conscious of multiple, varied contents at the same time (#4), with specific information-bearing contents that can vary over different orderings at different times (#13). Our experience also delineates a definite boundary, e.g. there is what we experience in a given moment as well as other things we are not experiencing at the same time (#14). The ToC requirement here is to provide an account of what it is that generates subjective experience, in such a way that it can experience multiple things at the same time and have an inside/outside boundary to what is experienced.

Item 12 illustrates a helpful strategy available to ToCs accounting for introspective experience. #12 refers to accepting the intrinsic phenomenology of consciousness at face value and therefore a postulate of cause-effect power, implying reference to the felt sense of causal agency. However, while any complete ToC does owe this experience an explanation, it is not given that the explanation should validate its superficial implication. In other words, accepting item 12 as a requirement does not preclude determinists or compatibilists from explaining that our felt sense of agency is misleading in one sense or another, provided they also explain why it feels the way it does.

Serial processing (#17) and intelligence (#18) can also be introspectively observed, albeit perhaps better described as capability-related observations rather than phenomenology-related properties of experience in such items as 4, 13, and 14. Experiments might also infer these capabilities as well, albeit with additional assumptions.

There is a useful distinction between these introspectively-observed features of consciousness. In some cases, we can imagine the feature not being present and yet we would remain phenomenally conscious. Nonetheless, it is sometimes present ("optional feature"). In other cases, the feature may be necessarily always present ("necessary feature"), at least in human consciousness. When that feature is absent, we are necessarily unconscious. The distinction is salient for ToC requirements. For optional features, a ToC must explain both how the human brain is capable of achieving

that functionality and integrating the experience of doing so into its consciousness, but also what consciousness consists of, given that it persists without that feature.[4] Meanwhile, for necessary features, the ToC must explain the one-to-one correlational or causal relationship between them and consciousness. Sorting specific features into one of these two categories may be a live scientific exercise, but we can point to a few examples to illustrate the difference.

Item 18 is a plausible candidate for an optional feature. Unless intelligence is defined so broadly that it loses its meaning, people can be in states in which they are incapable of intelligent reasoning and yet they are still phenomenally conscious. Item 13 is a plausible candidate for a necessary feature, noting that information in an IIT sense is minimally defined and potentially incorporates even states of minimal awareness, although meditative claims of contentless awareness would need to be accounted for separately.

If it is possible to be conscious of two, three or more concepts/qualia at the same time (#4), it seems likely that it is also possible in principle to be conscious of only a single concept/quale in a given moment. In the case, item 4 is also an optional feature. Some have claimed that unity (#1) is a necessary feature of conscious experience (Bayne, 2010). One reasoning is that if there truly were two fully disunified experiences taking place at the same time, it would instead be the case that two conscious perspectives were present, each corresponding to a separate experience. However, if a person were conscious of only a single quale, arguably there would be nothing for that quale to be "integrated" or "unified" with, making item 1 moot as a necessary feature.

Another tricky case is item 17, claimed by McFadden (2023). What should we do with claims of experienced conscious parallel processing? Perhaps not as massively parallel as unconscious processing, but nonetheless more than the "one at a time" implied by item 17. Conscious parallel processing may also be trainable. If we accept such claims, item 17 would not be a necessary feature. However, the explanandum may be more about why some brain processes are conscious while others are not – perhaps better linked to item 7 or issues around proceduralisation in learning. Alternatively, more robust scrutiny of conscious parallel processing may reveal that in fact only ever one stream is being processed consciously, but a misleading sensation of parallel processing sometimes emerges because the single stream subsumes multiple activities or delegates other activities to unconscious implementation.

The foregoing discussion of categories reveals several possible gaps in table 6. Necessary introspective explananda may be identifiable via meditator case notes and analytic philosophy, following Metzinger (2020) who identifies six features of minimal

[4] In principle, consciousness could also take the form of "at least one (or at least X) of this set of optional features must apply for a first person perspective to be present, but it could be any one (or any X)".

phenomenal experiences, which do not all have a natural fit with table 6 items (at least as currently phrased), e.g. introspective availability and transparency. Another related feature could be interruptibility, perhaps best understood as a subcategory of multiple related requirements. Any ToC must explain why consciousness, at least as we know it today, can always be interrupted by each of several factors, such as anaesthetics, sleep, meditative cessations, or damage to the brain.

Several other phenomenology-related properties of experience have also been discussed, which would need accounting for as much as those in table 6. For instance, Van Gulick (2022) talks also of semantic transparency, sensory transparency, dynamic flow, and self-perspectival organisation. Jones (2017) describes its unsolid nature.

If capabilities such as intelligence are to be included among the optional features, why not other capabilities like creativity (Baumeister et al., 2014) or emotional awareness (Tirapu-Ustárroz et al., 2003)? Perhaps intelligence is more productively treated as a subcategory label rather than an item-level requirement, i.e. to encourage its decomposition into operations - e.g. sensing, remembering, modelling, reasoning, deciding etc. - as well as the mental objects they operate on, e.g. percepts, memories, emotions, intents, thoughts etc.

Analytical explananda

Several other items in table 6 are analytical in origin, deriving from propositional linguistic thought, as distinct from introspection on the felt experience of consciousness. For instance, item 2 simply asks that whatever mechanism be proposed to explain consciousness should not be applied only at one spatiotemporal scale (unless an explicit, non-arbitrary threshold can be invoked), else we can analytically conclude the theory is incomplete or arbitrary. Here again, there are many candidate items omitted without explanation and potential subcategories. Other puzzles related to spatiotemporal scale/persistence could include the China Nation or Leibniz Mill thought experiments (Block, 1980) and personal identity puzzles (Parfit, 1984; Davidson, 1987).

Multiple realisability (#9) is an analytical insight that consciousness might be tied to phenomena that can arise in many different physical substrates, posing a puzzle to ToCs that suggest a substrate-specific implementation. But if this requirement is included, why not other analytical arguments that can also challenge substrate-specific ToCs, like the conceivability of functionally identical p-zombies (Kirk, 2023) or the recent scattered brain thought experiment based on neural recording technologies (Gidon et al., 2022)?

The unfolding argument (#8), by contrast, is a mathematical argument that tackles substrate-neutral theories equating consciousness with particular causal structures. However, why not include other analytical arguments that tackle substrate-neutral

theories, such as computationalism or functionalism? It seems bizarre to have excluded one of the most famous thought experiments in the field, the Chinese Room Argument (Searle, 1980), even if omissions of more recent critiques are understandable, such as the slicing problem (Gómez-Emillson & Percy, 2022). Bostrom's fractional qualia argument (2006) could also be placed in this subcategory. Even if Bostrom argues for biting his fractional bullet, others may disagree or provide alternative accounts, helping readers to gain a better overall grasp on the target ToC.

Physical grounding (#10) is a different kind of analytical puzzle relating more to physical realm interactions. Once a ToC has specified what consciousness is, we should ask how it can exist in and interact with the physical universe we observe. Such a question does not preclude an idealist answer, that what appears as physical emerges from the mental, but it does prompt idealism to explain how. Other physical realm puzzles could include how to avoid the infinite regress of "just more content" problems (Bayne, 2014) in a finite brain/universe, issues around the brainstorm machine (Dennett, 1997) and the vertiginous question (List, 2023), or apparent features of the physical realm that have been used to favour one ToC over another, such as frame invariance (Gómez-Emillson & Percy, 2023), ontological uncertainty (Koons, 2019), and the quantum measurement problem (Mould, 2001).

Discussions of analytical subcategories suggests a potential subcategory not present in table 6 at all. A class of semantic arguments suggest that there are mental phenomena without a physical parallel, placing a strong restriction on the types of ToC that could account for them (unless the arguments are escaped another way). This subcategory could include seminal, much-contested arguments like Mary's Room (Jackson, 1982), Inverted Qualia (Jackson, 1982), or incomputable knowledge (Penrose, 1989), as well as related arguments about the potentially non-physical nature of abstract objects (Agassi & Sagal, 1975) and even versions of Hume's missing blue argument (Hume, 1748).

Analytical explananda furnish ToCs with a different set of accounting strategies to introspective explananda. For instance, particular premises or logical transitions can be challenged. It may also be possible to accept the argument and conclusion, but reject the claimed philosophical significance or corollaries of the conclusion. For instance, a ToC could embrace epiphenomenalism as a response to item 3, but must still explain the logical argument that concludes (falsely perhaps) that evolution selected for human consciousness.

Empirical explananda

Items 5, 6, 7, and 15 are all requirements to account for scientifically-observed phenomena, i.e. to explain how the mechanism the ToC proposes for consciousness would produce the features of these phenomena. Explanations might suggest some phenomena are identical with consciousness; other features might be necessary or

sufficient conditions to cause consciousness; some might be caused by consciousness or simply happen to co-occur with it. Since empirical evidence traces back to an assumption about when to infer consciousness (normally informed ultimately by self-report as consciousness is only ever directly observed introspectively), an accounting strategy could also deny the assumption in a particular context.

As currently phrased these four items are perhaps best seen as subcategory labels, rather than specific individual requirements. For instance, Gambini & Pullin (2019) provide two requirements that would fall underneath the neural correlates test from McFadden (2023), which includes many different neural correlate. However, each of those two requirements could in turn be broken into separate items that might individually constrain different ToCs. There are multiple possible aspects of thalamocortical activity that relate to consciousness. There are multiple parts of the brain which can be damaged without disappearing all reports of phenomenal consciousness. There are in turn many more neural correlates than these two subcategories, as discussed by Paßler (2023). Likewise, there are multiple "paradigm cases" and multiple causes of divided self phenomena.

Subcategorising empirical explananda by the scale at which they are typically observed helps to identify other possible candidates. Most neural correlates are multi-neuronal (e.g. readiness potentials, a sweetspot level of synchronisation, module ignition; see e.g. Dehaene, 2014). Others span the whole brain system, such as correlates to do with brain waves (Hunt & Schooler, 2019) or brain temperature (Bond, 2023). From these two scales, we might both look to smaller scales, such as the effect of anaesthetic on microtubules within neurons (Hameroff, 2022) or other possible non-classical signals in the brain (Kerskens & Perez, 2022), or to larger scales, such as a full person or our environmental/societal interactions, including concepts around embodiment or other psychological phenomena beyond item 15, e.g. blindsight (Earl, 2014).

Theory desiderata

What is left from table 6? Items 11, 16, and 19 do not fit neatly into the explananda categories above. The requirement that a ToC's definitions and reasoning be non-circular is neither an introspection on our experiences/capabilities, nor an analytical thought experiment, nor an empirical observation. Indeed, it is also an impossible requirement to achieve in full. Eventually, all words are defined in terms of other words and all reasoning chains lead to unfalsifiable axioms. However, this requirement still has value when seen as asking a ToC to connect its terms to a broader edifice of language related to physical phenomena and patterns, so that we have some confidence that we are talking about the same specific thing (e.g. a response to the concerns of Mandik, 2016). While it cannot be achieved in full, this requirement can still be achieved better in some ToCs than others, e.g. mysterianism concedes the battle (McGinn, 1989).

Falsifiability is another example. As discussed in section 1, the core assumptions of ToCs are rarely falsified. Instead, experiments falsify some combination of core and auxiliary assumptions. Some theories may be self-avowedly unfalsifiable, but that does not guarantee their falsehood. Science may never conduct a lab experiment directly about the origins of the big bang, limited only to testing some present-day shadows of certain big bang theory variants when combined with specific auxiliary assumptions. It does not follow that all big bang theories are false, even if we might never be certain which were true. It would certainly be convenient, even desirable if the true ToC were to produce cleanly falsifiable predictions, but it may not be the case. These requirements can be categorised as scientific desiderata. We would like the ToC to achieve them as far as possible, but appreciate perfect achievement may be out of reach. Other scientific desiderata could include parsimony, consilience across the natural sciences, or symmetry. These are useful principles that might drive theories towards truth, but not explananda that a ToC must account for.

The last item from table 6 asks that the ToC allows us to measure how much consciousness is present in a given system. We can classify this as an instrumental desideratum for a theory. It would be useful if a theory delivered this, but it may not be possible. Likewise, it would be great if we had a theory that allowed the specification of particles' position and momentum to arbitrary precision, but we have good reason to believe this is impossible. Viewed from this perspective, other instrumental desiderata might include a ToC that is easy to explain to others or inspires us to live better lives. Some discussions seem motivated also by an implied instrumental desideratum that a ToC should reaffirm an interlocutor's prior spiritual/moral beliefs or pre-theoretical assumptions about what structures are conscious.

4. Discussion

The assessment above highlighted shortcomings of the five lists of ToC requirements identified. As well as having been mostly conducted without collaboration with each other and being motivated to defend particular ToC perspectives, the most severe shortcoming is the significant number of omitted items.

Adopting a taxonomic approach of specifying potential categories that each item from table 6 might fall into, we identified many other candidate items in each category. The longest list identified had 11 items, while a conservative reading of section 3 has over 15 introspective explananda, over 20 analytical explananda, over 10 empirical explananda (and many more specific items where some of those 10 are subcategory labels), and perhaps eight theory desiderata. Further scrutiny might develop formal listing principles to exclude some items or collapse others, but we can be confident that any list

respecting the style of Doerig et al. (2021a) as developed by McFadden (2023) would end up far longer than 11 items.

These practical shortcomings can be addressed with further work. The requirements listed in section 3 are not themselves intended to be comprehensive, but should now be getting a lot closer. The conclusion describes a workplan for developing a sufficiently complete list as a working tool. The question for this section is whether such work is worthwhile. Four plausible objections to the listed requirements epistemology are now discussed: type discrepancy, separate ToC targets, list overload, and epistemological contribution.

First, as noted in section 3, these lists consist of very different types of item, from introspected properties of experience through to mathematical arguments and brain scan data. Is it even meaningful to list them together? A minimal response to this objection is to embrace a taxonomic categorisation that at least separates requirements by type. But should such different types even be jointly considered in a ToC assessment exercise, as applied by Doerig et al. (2021a) and McFadden (2023)? An argument in favour is that the items all have something relevant in common. They make an ask of a ToC and a ToC's accounts across all such items would need to be self-consistent in a successful, complete theory.

To provide two simplified examples for illustration: a global workspace theory might provide an account that gels with certain introspected sensations and certain empirical observations in the brain, but might struggle to form a self-consistent account that addresses divided self phenomena as a psychology-derived explanandum (there can only be one global workspace by usual definitions) or scaling arguments as an analytical explanandum (what phase transition prevents a minimal definition of "broadcasting to multiple modules" generating consciousness in trivially many systems). By being prompted to tackle such requirements directly, the ToC could refine its position and identify useful areas for future research. As a second example, a ToC that explains all empirical phenomena in terms of a physical brain mechanism would be inconsistent (or at least incomplete) if its explanation of Mary's Room failed to reject the existence of non-physical mental phenomena. Likewise, if readers find its Mary's Room explanation unsound, that is reason for them to downweight confidence in the ToC.

Cherry-picking requirements or types of requirement that a ToC handles well, while ignoring others, is one of the problems that inspires this epistemology in the first place. Any attempt to exclude certain categories of requirements would make it easy for theorists to claim the exclusions are motivated by a desire to favour some ToCs over others, damaging the approximate neutrality of the effort that would be needed to support the field as a whole.

A second objection is that some requirements appear targeted to specific ToCs, leaving others relatively unscathed. Perhaps there should be separate lists for different classes of ToC, rather than a single combined taxonomy. For instance, the unfolding argument primarily raises questions for causal structure theories; the Chinese Room Argument for functionalism; causal efficacy requirements for dualism etc. However, responding to such requirements for non-target ToCs should be an easy exercise – a few sentences to explain why it does not apply to their account. This request costs us little. By contrast, the benefits of a single list for all ToCs are significant. It means we can have a single comprehensive list to update and improve, rather than the confusion of several partial or proliferating lists that are specific to certain ToC categories. It also means theorists know where to look directly rather than having to first categorise their ToC, which can be a contested exercise. Finally, it may shed some unexpected light about what the non-target ToCs really think and help them to develop a consistent framework overall. For instance, the Combination Problem was originally targeted by Chalmers (2016) against panpsychism but subsequent scholars have found utility in addressing it also against other ToCs (e.g. Mendelovici, 2019).

Third, where should such a search for requirements stop? If trying to account for all conscious experiences and observations we can consciously make, do we end up requiring a ToC to integrate gravity into quantum theory, solve self-referential paradoxes in logic, or explain the minutiae of every neuronal structure and mental experience? These are familiar problems in science, with fuzzy boundaries between disciplines handled pragmatically even though reality is endlessly overlapping. The proliferation of theories suggests we should be looking for more explananda and theory constraints, not fewer. Sensible definitions of scope and a taxonomic structure of categories and subcategories can be used to make it manageable, even if we end up with very many individual experiments and observations backing up low level items. It is also possible to phrase accounts for certain phenomena in general language that captures whole subsets in a single explanation. Such general accounts may be adequate for some discussions, even while inspiring other theorists to make progress exploring the specifics.

Fourth, how would a taxonomised list of requirements actually lead to consensus on a single ToC? Different theorists can still disagree on which accounts to a given requirement are satisfactory and which requirements matter more than others. Allowing scope for this flexibility is a feature not a bug, but does point to the limits of the listed requirements epistemology. Respecting responses to the other objections and following the vision of section three, we need a broad, inclusive list. Inevitably, this means some requirements will be considered irrelevant or "obviously wrong" by some theorists. In order to maintain a chance at a collaborative, comprehensive list, such items remain in scope but opposing theorists are invited to explain why they are irrelevant, even if only briefly and irreverently.

These four objections place important limits on what we should expect from a listed requirements epistemology. They point to a need for taxonomisation (embracing debate on appropriate categorisation), to a need to be pragmatic and inclusive, and to viewing it as an epistemological tool, rather than a silver bullet strategy. As a tool, it needs to be combined with an overarching epistemological strategy. For instance, turning to some of the ideas in section one, the second step of the four step epistemology in Kirkeby-Hinrup and Fazekas (2021) requires compiling all of the evidence for and against different ToCs, an exercise that would be met by the listed requirements tool (they argue to compile just empirical evidence, but this paper suggests other evidence can also be helpful). Identifying the adequacy of different ToC accounts to an itemised requirement could, at times, be met by the contrastive tests described by Del Pin et al. (2021). Nonetheless, a collective effort to build a comprehensive taxonomy of ToC requirements, albeit with items of varying quality, is nonetheless a potential point for collaboration in an often-aggressively contested field.

Conclusion

This paper has reviewed five lists of ToC requirements produced to assess the validity of competing ToCs. Despite the lists' practical shortcomings to date, we have argued that this "listed requirements" method is a useful epistemological tool to be used as part of an overall strategy, both for choosing between ToCs and for helping ToC theorists identify new areas of research that might address weaker areas. We conclude by briefly outlining a workplan to develop a future taxonomy that addresses the shortcomings of the lists identified so far. The workplan might also be a useful Schelling point and early task for the consciousness discussion forum proposed by Klein (2021).

The first step is a desk exercise to build as comprehensive a list of individual requirements and their possible categorisations as possible. For instance, many lists have been developed in adjacent areas which could be used to expand upon the items and categories identified in section 3.2. While not formally presented as ToC requirements, many of their items or ideas could be usefully reframed as such, given the inclusive character of the items already in table 6. For instance, Seth et al. (2005) list 17 basic brain facts and properties of consciousness. Chalmers (2018) describes nine categories of problem intuition for the meta-problem of consciousness. Anokin (2021) lists 10 characteristic features of consciousness. Ideas for how to measure consciousness may also produce indirect ideas for requirements, such as the five layer hierarchy from Tirapu-Ustárroz et al. (2003), the six indicators from Pennartz et al. (2019), the correlates taxonomy from Hunt et al. (2022), or the perturbational complexity index (Casali et al., 2013). Even where individual items of evidence might aggregate into a single ToC requirement, it is helpful to have those individual items

available or referenced, such as the 31 items compiled to support an access-phenomenal distinction (Kirkeby-Hinrup & Fazekas, 2021).

Such an exercise can only ever produce a working list since new requirements might always be identified via new research or good arguments identified to group/split items in a different way. The second step is therefore a collaborative, interdisciplinary, and ultimately ongoing review to expand and refine the working list. This step could also raise awareness of the tool and develop guides for using it in the context of overarching epistemological strategies. The working list/taxonomy could be available online in an easily-navigated format, with capacity for viewers to raise questions or suggestions in specific parts of the taxonomy.

The third step is to invite or commission particular ToC theorists to assess their preferred theory against the items, i.e. to provide a written account to address each item or group of items. A fourth step could then convene an interdisciplinary panel for a Delphi-style collaborative and iterative process to score the ToC accounts and form a collective view on which ones are strongest in which areas. If the exercise has been useful, it will now be possible to direct research effort to the weakest areas and have constructive discussions about our strongest current theories.

Acknowledgments: [redacted from blinded manuscript for review]

Funding: This research received no external funding or support.

Data Availability Statement: No new data were generated for this study.

Conflicts of Interest: No conflicts of interest are declared.

References

Agassi, J., & Sagal, P. (1975). The Problem of Universals. Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition 28, no. 4, 289–94.

Anokhin, K. V. (2021). The Cognitome: Seeking the fundamental neuroscience of a theory of consciousness. Neuroscience and Behavioral Physiology, 51(7), 915-937.

Baumeister, R.F., Schmeichel, B.J., & DeWall, C.N. (2014). Creativity and Consciousness. In E. Paul & S. Kaufman (Eds). The Philosophy of Creativity: New Essays. (pp. 185-198). Oxford University Press.

Bayne, T. (2010). The Unity of Consciousness. Oxford: Oxford Univ. Press.

Bayne, T. (2014). Replies to Commentators. Analysis, Volume 74, Issue 3, Pages 520–529, https://doi.org/10.1093/analys/anu066

Block, N. (1980). Troubles with functionalism. Minnesota Studies in the Philosophy of Science, 9, 261-325.

Bond, E. (2023). The contribution of coherence field theory to a model of consciousness: electric currents, EM fields, and EM radiation in the brain. Frontiers in Human Neuroscience (16). 10.3389/fnhum.2022.1020105

Bostrom, N. (2006). Quantity of experience: brain-duplication and degrees of consciousness. Minds & Machines. 16, 185–200. https://doi.org/10.1007/s11023-006-9036-0

Casali A.G., Gosseries O., Rosanova M., Boly M., Sarasso S., Casali K.R., Casarotto S., Bruno M.-A., Laureys S., Tononi G., & Massimini M. (2013). A theoretically based index of consciousness independent of sensory processing and behavior. Sci. Transl. Med. 5(198) http://stm.sciencemag.org/content/5/198/198ra105

Chalmers, D. (2016). The Combination Problem for Panpsychism. In G. Bruntrup & L. Jaskolla (Eds.), Panpsychism (pp. 179–214). Oxford: Oxford University Press. 10.1093/acprof:oso/9780199359943.003.0008

Chalmers, D. (2018). The meta-problem of consciousness. Journal of Consciousness Studies, 25 (9–10) pp. 6-61

Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. Trends in cognitive sciences, 15(8), 358-364.

Dardashti, R., Dawid, R., & Thébault, K. (Eds.) (2019). Why trust a theory? Epistemology of fundamental physics. Cambridge University Press.

Davidson, D. (1987). Knowing One's Own Mind. Proceedings and Addresses of the American Philosophical Association. 60 (3): 441–458. doi:10.2307/3131782

Dehaene, S. (2014). Consciousness and the brain: Deciphering how the brain codes our thoughts. Penguin.

Del Pin, S. H., Skóra, Z., Sandberg, K., Overgaard, M., & Wierzchoń, M. (2021). Comparing theories of consciousness: why it matters and how to do it. Neuroscience of consciousness, 2021(2), niab019. https://doi.org/10.1093/nc/niab019

Dennett, D. (1997). Quining Qualia. In N. Block (ed.). The Nature of Consciousness. Cambridge: MIT Press. p. 623. ISBN 0-262-52210-1.

Doerig, A., Schurger, A., Hess, K., & Herzog, M. H. (2019). The unfolding argument: Why IIT and other causal structure theories cannot explain consciousness. Consciousness and cognition, 72, 49-59. https://doi.org/10.1016/j.concog.2019.04.002.

Doerig, A., Schurger, A., & Herzog, M. (2021a). Hard criteria for empirical theories of consciousness, Cognitive Neuroscience, 12:2, 41-62, DOI: 10.1080/17588928.2020.1772214

Doerig, A., Schurger, A., & Herzog, M. (2021b). Response to commentaries on 'hard criteria for empirical theories of consciousness', Cognitive Neuroscience, DOI: 10.1080/17588928.2020.1853086

Earl, B. (2014). The biological function of consciousness. Front. Psychol. 5:697. doi:10.3389/fpsyg.2014.00697

Fahrenfort, J., & van Gaal, S. (2021). Criteria for empirical theories of consciousness should focus on the explanatory power of mechanisms, not on functional equivalence. Cognitive Neuroscience, 12:2, 93-94. https://doi.org/10.1080/17588928.2020.1838470

Fazekas, P., Cleeremans, A., & Overgaard, M. (2024). A construct-first approach to consciousness science. Neuroscience & Biobehavioral Reviews, 156 105480. https://doi.org/10.1016/j.neubiorev.2023.105480

Gambini, R., & Pullin, J. (2019). Physical requirements for models of consciousness. Mind and Matter, 17(2), 129-154.

Gidon, A., Aru, J., & Larkum, M. E. (2022). Does brain activity cause consciousness? A thought experiment. PLoS Biology, 20(6), e3001651.

Goff, P. (2019). Galileo's Error: Foundations for a New Science of Consciousness. New York: Pantheon, London: Rider.

Gómez-Emilsson, A. & Percy, C. (2022). The "Slicing Problem" for Computational Theories of Consciousness. Open Philosophy, 5(1), 718-736. https://doi.org/10.1515/opphil-2022-0225

Gómez-Emilsson, A. & Percy, C. (2023). Don't forget the boundary problem! How EM field topology can address the overlooked cousin to the binding problem for consciousness. Frontiers in Human Neuroscience

Griffin, D. R. (2012). Whitehead's radically different postmodern philosophy: an argument for its contemporary relevance. State University of New York Press.

Hameroff, S. (2022). Consciousness, cognition and the neuronal cytoskeleton–A new paradigm needed in neuroscience. Frontiers in Molecular Neuroscience, 15, 869935.

Hume, D. (1748). An Enquiry Concerning Human Understanding. London: A. Millar.

Hunt, T., & Schooler, J. W. (2019). The easy part of the hard problem: a resonance theory of consciousness. Frontiers in Human Neuroscience, 378.

Hunt, T., Ericson, M., & Schooler, J. (2022). Where's My Consciousness-Ometer? How to Test for the Presence and Complexity of Consciousness. Perspectives on Psychological Science, 17(4), 1150-1165. 10.1177/17456916211029942

Jackson, F. (1982). Epiphenomenal Qualia. Philosophical Quarterly, 32: 127–136.

Jones, M. (2017). Mounting evidence that minds are neural EM fields interacting with brains. Journal of Consciousness Studies, 24(1-2), 159-183.

Kerskens, C. M., & Pérez, D. L. (2022). Experimental indications of non-classical brain functions. Journal of Physics Communications, 6(10), 105001.

Kirk, R. (2023). Zombies. In N. Zalta & U. Nodelman (Eds). The Stanford Encyclopedia of Philosophy (Summer 2023 Edition). https://plato.stanford.edu/archives/sum2023/entries/zombies

Kirkeby-Hinrup, A., & Fazekas, P. (2021). Consciousness and inference to the best explanation: Compiling empirical evidence supporting the access-phenomenal distinction and the overflow hypothesis. Consciousness and Cognition, 94, 103173.

Klein, S. (2021) Proposal to create discussion forum to consider evidence on theories of consciousness, Cognitive Neuroscience, 12:2, 91-92, DOI 10.1080/17588928.2020.1846026

Koons R.C. (2019) Thermal substances: a Neo-Aristotelian ontology of the quantum world. Synthese. 10.1007/s11229-019-02318-2

Lee, A. Y. (2023). Degrees of consciousness. Noûs, 57, 553–575. https://doi.org/10.1111/nous.12421

List, C. (2023). A quadrilemma for theories of consciousness. [preprint] https://philpapers.org/rec/LISAQF

Mandik, P. (2016). Meta-illusionism and qualia quietism. Journal of Consciousness Studies, 23(11-12), 140-148.

Manzotti, R., & Chella, A. (2014). Physical integration: A causal account for consciousness. Journal of Integrative Neuroscience, 13(02), 403-427. https://doi.org/10.1142/S0219635214400044

McFadden, J. (2023). Consciousness: Matter or EMF? Frontiers in Human Neuroscience.

McGinn, C. (1989). Can We Solve the Mind–Body Problem?, Mind, Volume XCVIII, Issue 391, Pages 349–366, https://doi.org/10.1093/mind/XCVIII.391.349

Mendelovici, A. (2019). Panpsychism's combination problem is a problem for everyone. In The Routledge Handbook of Panpsychism (pp. 303-316). Routledge

Metzinger, T. (2020). Minimal phenomenal experience: Meditation, tonic alertness, and the phenomenology of “pure” consciousness. Philosophy and the Mind Sciences, 1(I), 7. https://doi.org/10.33735/phimisci.2020.I.46

Mould, R. A. (2001). The parallel principle. arXiv preprint quant-ph/0111096.

Niikawa, T., Miyahara, K., Hamada, H., & Nishida, S. (2022). Functions of consciousness: conceptual clarification, Neuroscience of Consciousness, Volume 2022, Issue 1, 2022, niac006, https://doi.org/10.1093/nc/niac006

Parfit, D. (1984). Reasons and Persons. Oxford [Oxfordshire]: Clarendon Press.

Paßler, M. (2023). The exclusionary approach to consciousness, Neuroscience of Consciousness, Volume 2023, Issue 1, 2023, niad022, https://doi.org/10.1093/nc/niad022

Pennartz, C. M. A., Farisco, M., & Evers, K. (2019). Indicators and Criteria of Consciousness in Animals and Intelligent Machines: An Inside-Out Approach. Frontiers in systems neuroscience, 13, 25. https://doi.org/10.3389/fnsys.2019.00025

Penrose, R. (1989). The Emperor's New Mind: Concerning Computers, Minds and The Laws of Physics. Oxford: OUP.

Sattin, D., Magnani, F.G., Bartesaghi, L., Caputo, M., Fittipaldo, A.V., Cacciatore, M., Picozzi, M., & Leonardi, M. (2021). Theoretical Models of Consciousness: A Scoping Review. Brain Sci. 2021, 11, 535. https://doi.org/10.3390/brainsci11050535

Searle, J. (1980). Minds, brains, and programs. Behavioral and Brain Sciences, 3 (3), 417–457. 10.1017/S0140525X00005756

Seth, A. K., Baars, B. J., & Edelman, D. B. (2005). Criteria for consciousness in humans and other mammals. Consciousness and Cognition, 14(1), 119–139. https://doi.org/10.1016/j.concog.2004.08.006

Seth, A., & Bayne, T. (2022). Theories of consciousness. Nat Rev Neurosci 23, 439–452. https://doi.org/10.1038/s41583-022-00587-4

Signorelli, C. M., Szczotka, J., & Prentner, R. (2021). Explanatory profiles of models of consciousness-towards a systematic classification. Neuroscience of Consciousness, 2021(2), niab021.

Tirapu-Ustárroz, J., Muñoz-Céspedes, J., & Pelegrín-Valero, C. (2003). Towards a taxonomy of consciousness. Rev Neurol 2003;36 (11):1083-1093. doi: 10.33588/rn.3611.2003009

Tononi, G. (2017). Integrated information theory of consciousness: Some ontological considerations. In M.Velmans & S.Schneider (Eds.). The Blackwell Companion to Consciousness (p 621-633). 10.1002/9781119132363.ch44

Van Gulick, R. (2022). Consciousness. In E. Zalta & U. Nodelman (Eds.), The Stanford Encyclopedia of Philosophy (Winter 2022 Edition). https://plato.stanford.edu/archives/win2022/entries/consciousness/

Wiese, W. (2018). Toward a mature science of consciousness. Frontiers in Psychology, 9, 693. https://doi.org/10.3389/fpsyg.2018.00693

Yaron, I., et al., (2022). The ConTraSt database for analysing and comparing empirical studies of consciousness theories. Nat. Hum. Behav. 6, 593–604. https://doi.org/ 10.1038/s41562-021-01284-5.