v2, 18 July 2024

When is a brain organoid a sentience candidate?

Jonathan Birch

Centre for Philosophy of Natural and Social Science,
London School of Economics and Political Science,
Houghton Street, London, WC2A 2AE, UK.
j.birch2@lse.ac.uk
http://personal.lse.ac.uk/birchj1

Abstract

It would be unwise to dismiss the possibility of human brain organoids developing sentience. However, scepticism about this idea is appropriate when considering current organoids. It is a point of consensus that a brain-dead human is not sentient, and current organoids lack a functioning brainstem. There are nonetheless troubling early warning signs, suggesting organoid research may create forms of sentience in the near future. To err on the side of caution, researchers with very different views about the neural basis of sentience should unite behind the "brainstem rule": if a neural organoid develops or innervates a functioning brainstem that regulates arousal and leads to sleep-wake cycles, then it is a sentience candidate. If organoid research leads to the creation of sentience candidates, a moratorium or indefinite ban on the creation of the relevant type of organoid may be appropriate. A different way forward, more consistent with existing approaches to animal research, would be to require ethical review and harm-benefit analysis for all research on sentience candidates.

Key words: organoids, sentience, consciousness, brainstem, regulation, precaution

1. The promise of organoid research

Biomedical research urgently needs new and better alternatives to animal models. The trend in recent decades has been towards increasing reliance on a small number of model species, especially rats, mice, zebrafish and fruit flies, and towards a troubling level of dependence on assumptions about the relevance of these model systems to human medical conditions (Farris, 2020). Many researchers and funding agencies have invested heavily in the idea that understanding the brain mechanisms of animal models will help us understand complex

conditions such as depression, anxiety, autism or schizophrenia in their human forms. But animal models are far from perfect models, leading to widespread reflection on how things could be done differently. (Shemesh & Chen, 2023; Taschereau-Dumouchel et al., 2022).

The maxim to "replace, reduce and refine" (the "3Rs") was coined more than fifty years ago and is now embedded in frameworks for the regulation of animal research around the world (Hubrecht & Carter, 2019). It crystallizes a point of wide agreement. We should aim to replace animal models with other types of model where possible, reduce the numbers of animals being used, and refine experimental techniques to minimize suffering. Yet this maxim has turned out to be compatible with a drastic increase over those same fifty years in the total numbers of animals used (Taylor & Alvarez, 2019).

So, we have two disquieting trends: growing concern about the ability of biomedical research on animal models to deliver tangible benefit, particularly in relation to neurological/mental conditions, and a growing realization that, despite widespread endorsement of the 3Rs, invasive animal research is on the rise, not on the way out. These trends raise the question: what is the alternative? To study a complex condition like depression or autism, the argument goes, you cannot simply study tissue in culture, but you also cannot study human subjects at the level of mechanistic detail required to understand how, for example, particular alleles and patterns of gene expression may influence these conditions. So, you must use animals, where the ethical limits on what can be done are more permissive and a broader range of interventions is available.

This is where brain organoids have tremendous promise. The organoid is a relatively new kind of model system with great potential for replacing invasive animal research. Organoids are models of organs constructed from pluripotent stem cells. Human stem cells can be used, leading to miniature models of human organs constructed from human tissue. Suppose, for example, you want to understand human kidney function. One option is to study the renal system of a rat or mouse, relying on the idea that this will resemble human kidney function in the ways that matter. But organoid technology gives you a new option. You take pluripotent human stem cells and induce them to differentiate into kidney cells. The kidney organoid you construct will still differ from a normal kidney in many ways, but you have a degree of control over those ways, and you can be confident that the genes being expressed are the same as those in human kidney cells.

When we are talking about kidney organoids, gastrointestinal organoids, cardiac organoids, and other types of non-neural organoid, these developments should be celebrated. We should not try to put the brakes on a programme that could turn out to deliver the alternative to animal research that has been so sorely needed for so long.

But when it is the *brain* being modelled, the work becomes more controversial, and rightly so. A brain organoid is a model constructed from pluripotent stem cells induced to form organized neural tissue. Here too, it is the use of human stem cells to create human neural organoids that is generating major scientific excitement. I will use the term "brain organoid" here, but I note

that the term “neural organoid” is also used, and the terms “cerebral organoid” and “cortical organoid” are also often used in cases where the organoid is intended to model the human neocortex. Sometimes these models have been implanted into mice, leading to functional connections with the mouse’s brain (Wilson et al., 2022), my main focus here will be on extracorporeal organoids, sustained in a controlled environment outside of a living body.

There are ethical reasons in favour of doing this research, if it allows us to model neurological conditions for which scientists currently lack good models, and if it can substitute for invasive animal research. And yet the research invokes the image, if not currently the reality, of a sentient brain in vitro, and this image fills many onlookers with a sense of horror, regardless of whether the brain is human or non-human (although I will be focusing, in this article, on human brain organoids). Even when one looks at the research as it is now, it is hard not to feel a certain unease at the idea of a miniature model of a brain constructed from brain tissue. Sometimes unease is a bias we should try to overcome. But sometimes it is pointing us in the direction of genuine moral reasons to pause the research.

We need to think seriously about which of these possibilities is the case here. In recent years, a number of bioethicists have been doing just that (Ankeny & Wolvetang, 2021; Hyun et al., 2020; Kreitmair, 2023; Lavazza, 2020; Lavazza & Massimini, 2018; Niikawa et al., 2022; Sawai et al., 2019; Sharma et al., 2021; Żuradzki, 2021). I have weighed into the debate already, advocating for a precautionary approach to these issues (Birch & Browning, 2021). This is an extension of the approach I advocate towards non-human animals in which sentience is disputed, such as shrimps and insects (Birch, 2017). But, as Steel (2015) has argued, even though formulations of the precautionary principle often fail to mention consistency, requiring consistency in our thinking about different risks is crucial if we are to avoid inconsistent packages of recommendations. In this case, we need to make sure our approach to organoids is fully consistent with our approach to animal research. In particular, we must be careful not to be overcautious regarding organoids in a way that undermines their promise as replacements for animals.

My goal in this article is to find the right balance. Sections 2 and 3 set out the ingredients of a precautionary approach to questions of sentience. The discussion here will be quite general and will not specifically concern organoids. The crucial concept introduced is that of a *sentience candidate*: a system that is not certain to be sentient, but which is sentient according to at least one reasonable, evidence-based theoretical position. Sections 4-8 ask: when is a brain organoid a sentience candidate? I argue that the presence of a functioning brainstem that regulates arousal and leads to sleep-wake cycles is enough. Section 9 considers the precautionary steps we should consider when a brain organoid is a sentience candidate.

2. A scientific meta-consensus

The term “sentience” in English comes from the Latin “sentire”, literally “to feel”. It is used in different ways in different contexts, with the idea of “feeling” providing a loose common

thread. Sometimes, people in brain organoid research use the term to mean nothing more than "responsiveness to sensory stimuli due to adaptive internal processes" (Kagan et al., 2022). When the term is used in this way, some preparations of human brain tissue are already sentient. However, I strongly recommend against using the term in this way, because it creates a large gap with how the term has come to be used in bioethics, animal ethics, animal law and the science of animal welfare.

In those fields, to say that a system is "sentient" is to say that it is capable of *valenced conscious experiences* such as pain or pleasure. That is: in at least in some conditions, there is something it is like to be that system, and the experience is either pleasant (positively valenced) or unpleasant (negatively valenced). The reason for using the term in this way is that this capacity is widely taken to be morally significant. Put simply, it is a good thing when animals have conscious pleasant experiences, whereas unpleasant experiences such as pain are a source of ethical concern. That is the sense in which I will be using the term. Not everyone would agree with that definition, and we could spend a whole article unpacking it, but this is not the place for that. Our question is whether there is good reason to think that brain organoids could already be—or have the potential to become—sentient in this sense.

There is no scientific consensus about the neural basis of sentience or phenomenal consciousness in humans, other mammals, or any other animals. Contemporary consciousness science contains a wide range of positions (Seth & Bayne, 2022). It is equally important, though, to see that an absence of consensus on a specific theory does not lead to a chaotic "anything goes" situation in which all speculation is equally valid. Evidence still constrains theorizing. Some options are serious and evidence-based, while others are not.

The concept of "meta-consensus" can be helpful for thinking about these situations. The concept is borrowed from political science (Dryzek, 2010). In very broad terms, the motivation for the concept is that people may agree about a lot, even when they disagree about the best policy. Crucially, they may still agree about the range of reasonable options, and they may agree about how these options relate to each other along important dimensions (such as more moderate to more radical). "Meta-consensus" is a term for consensus on these "meta" questions concerning the option space. Seeing a meta-consensus can be an incredibly important step towards negotiating a way forward.

To my knowledge, the concept has not yet received explicit discussion in relation to scientific disagreement. But it should. Just as finding a meta-consensus can help lawmakers move forward when they disagree, so finding a scientific meta-consensus can help scientists move forward, as well as helping outside audiences to better understand what is going on in the science. It is all too easy for a non-expert, looking in, to think "since they disagree so much, there is no reason for me to listen to a word they have to say. I'll just go with my gut feeling". That is a poor inference, and a very dangerous one too, but it can be a tempting one when scientists cannot articulate clearly what they *do* agree about.

Does meta-consensus exist in the science of sentience? I think it does. I will first present where I think the meta-consensus lies, and then explain *why* I think this:

Proposed meta-consensus:

Given our current evidence, all of the following theoretical positions about the neural system requirements for sentience (defined as the capacity for valenced experience) are realistic possibilities. None should be held dogmatically, but all should be taken seriously in practical contexts:

R1. Sentience requires distinctively primate neural mechanisms (e.g. in granular prefrontal cortex) and is absent in non-primates.

R2. Sentience requires mechanisms distinctive to the mammalian neocortex and is absent in non-mammals.

R3. Sentience requires the neocortex in mammals but can also be achieved by other brain mechanisms performing relevantly analogous functions (such as the avian pallium).

R4. Sentience does not require the neocortex even in mammals and can be achieved in at least a minimal form by integrative subcortical mechanisms crucially involving the midbrain. However, it is absent in non-vertebrates.

R5. Sentience does not require the neocortex even in mammals and can be achieved in at least a minimal form by integrative subcortical mechanisms crucially involving the midbrain. Moreover, it can also be achieved by other brain mechanisms performing relevantly analogous functions (such as the central complex in insects).

These five positions are ordered from less inclusive to more inclusive. R5 is the most inclusive, in the sense that the distribution of sentience in the animal kingdom is likely to be the widest if this position is correct, since midbrain mechanisms are far more widely shared than neocortical mechanisms.

By contrast, it is not reasonable, given current evidence, to give serious attention in practical contexts to views less inclusive than R1 (such as a view on which sentience requires a developed capacity for natural language) or more inclusive than R5 (such as a view on which the spinal cord is said to support sentience by itself in the absence of a brainstem). The evidence does not support taking these views seriously in practical contexts.

There is no consensus about which of R1–R5 is correct, and each option can be fleshed out in many different ways. At the most inclusive end of the reasonable range, Merker (2007),

Panksepp (1998), and Solms (2021) have defended midbrain-centric theories that are neutral between R4 and R5, while Barron and Klein (2016), Ginsburg and Jablonka (2019), Feinberg and Mallatt (2016), and Tye (2016) have defended versions of R5. Damasio can be placed approximately between R3 and R4, since he has often emphasized the importance of both the midbrain and some parts of the cortex (especially the insular and somatosensory cortex) (Damasio et al. 2000, 2013).

Meanwhile, many cortex-centric, computational functionalist theories, such as the global workspace theory, the perceptual reality monitoring theory, and the recurrent processing theory are most naturally interpreted as versions of R3. Both Dehaene (2014) and Lau (2022) posit important roles for distinctively primate mechanisms in the human implementation of the mechanisms they take to be responsible for conscious experience: Dehaene proposes a key role for dorsolateral prefrontal cortex in implementing the global neuronal workspace, while Lau proposes a key role for dorsolateral and frontopolar prefrontal cortex in implementing perceptual reality monitoring. However, both allow that these mechanisms may have alternative implementations in other animals. Lamme (2022), in developing the recurrent processing theory, focuses on mammalian visual areas (such as visual cortex), but recurrent processing could, clearly, be implemented by other animals in their own sensory areas. Humphrey's (2022) also falls in the R3 zone: he recognizes that the feedback loop he takes to be constitutive of conscious experience may be implemented differently in birds.

At the less inclusive end, R1 includes more demanding computational functionalist theories, on which sentience is linked to complex computations that may only be achievable by brain mechanisms distinctive to the primate lineage. The relevant mechanisms are located in granular prefrontal cortex (granular PFC), a part of the frontal lobe greatly expanded and elaborated in primates, incorporating the frontopolar, dorsolateral, and ventrolateral prefrontal cortex and characterized by a notably thick layer of granular (layer IV) cortical neurons (Preuss and Wise 2022). These brain regions are strongly linked to executive control functions. Rolls's (2004, 2014) 'higher-order syntactic thought' theory gives a crucial role to these mechanisms. LeDoux has at times appeared sympathetic to R1 and has emphasized the special processing properties of granular PFC (LeDoux 2023, pp. 758-9). However, his most recent work clarifies that granular PFC is required only for the most cognitively demanding kinds of consciousness: 'autonoetic' and 'noetic' consciousness (LeDoux et al. 2023). He allows that 'anoetic' consciousness, which I see as much closer to the idea of sentience, may be achievable in a much wider range of animals.

The R2 category includes theorists who have, for various reasons, proposed that neocortical neurons, and perhaps especially the large pyramidal neurons in layer V, may have special processing properties that allow them to support consciousness (e.g. Aru et al. 2020; Beck and Eccles 1992; Key 2015). On this view, granular PFC is not necessary, potentially allowing all mammals to meet the requirements, but there is something very special about the neocortex more generally. For Beck and Eccles (1992), for example, pyramidal layer V neurons were the most likely entry point for mental causation in the workings of the brain. On this (admittedly highly speculative) theory, the nucleated structure found in birds might not be enough.

The consensus lies not at the level of specific positions (clearly!) but rather at the meta-level, in the idea that everyone should be able to recognize any of the positions in the range R1–R5 as realistic possibilities that must be taken seriously in practical contexts. All positions in this range have some evidence behind them, conferring a degree of plausibility. Moreover, everyone should be able to agree on the ordering of these views from less inclusive to more inclusive (**Fig. 1**). Finally, everyone should be able to agree on the severe challenges facing any view that sees both the neocortex and the midbrain as unimportant to sentience, or any view that regards a functional primate brain as insufficient.

This may sound like it does not exclude very much, but it does. Consider, for example, the cerebellum. This is part of the hindbrain, at the very back and base of the brain, and it contains more neurons than any other brain region, even the cortex. There are 69 billion neurons in your cerebellum, compared with a mere 16 billion in the cortex (Herculano-Houzel, 2009). If one were trying to guess the 'seat of consciousness' in the brain using nothing but neuron counts, one would probably guess the cerebellum—and be completely wrong. There is no evidence for a role for the cerebellum in generating conscious experience and strong evidence against. The cerebellum has important roles in motor control and sensorimotor integration, and appears to be crucially involved in modelling the expected sensory consequences of our actions and registering prediction errors (Arikan et al., 2019; Johnson et al., 2019). These computations could have turned out to be essential to sustaining a conscious state, but they turn out not to be, as a matter of empirical fact. Being born without a cerebellum (complete primary cerebellar agenesis) leads to motor control problems but turns out to be compatible with otherwise normal cognitive development (Yu et al., 2014).

So, the evidence does not warrant attaching significant probability to a hindbrain-centric theory of sentience, or a theory that blithely predicts that sentience will be tied to the brain region with the most neurons with no consideration of what the neurons are doing. One cannot pluck theories out of thin air, without supporting evidence, and expect them to be taken seriously when practical questions are at stake. There are too many possible-but-very-low-probability theories, and their practical implications are so diverse that they are apt to derail discussion if we admit them to the table. In practical contexts, we need to maintain a focus on credible theories that have amassed enough evidence in their favour to merit serious discussion of their practical implications.

For another example, this time from the other end of the axis, consider a theory that ties sentience to natural language. There are serious theories, such as Rolls's (2004, 2014) 'higher-order syntactic thought' theory, that tie conscious experience to quite sophisticated kinds of thought, suggesting a narrow distribution of sentience in the animal kingdom. Yet even Rolls stops short of proposing that natural language is required for the relevant type of thought, allowing that a 'language of thought' might also be sufficient. This is a wise move, because we have clear evidence that linguistic abilities are not needed to have conscious experiences. Brain injuries to regions associated with language can lead to temporary aphasia (loss of linguistic

ability) of various kinds, but subjects, when they recover, can often vividly recount their conscious experiences during the time they were affected (Koch, 2019).

One theory that is challenging to locate in the R1-R5 range is the integrated information theory (IIT) of Tononi and collaborators (summarized in Tononi et al., 2023). In consciousness science, researchers outside the IIT camp often distinguish *fundamental IIT* (or strong IIT), the full version of the theory including a highly speculative metaphysical background picture, from *empirical IIT* (or weak IIT), which simply claims that, in the human brain, the neural correlate of conscious experience is a "posterior cortical hot zone" at the back of the neocortex and that the high causal integration of this region is what makes it apt for this role (Michel and Lau, 2020; Mediano et al. 2019, 2022). Empirical IIT is too thin a claim to locate in the R1-R5 range, but it has affinities with the R2 group, in that it ascribes a special status to the neocortex without giving any special emphasis to the *prefrontal* cortex (which is, in fact, de-emphasized).

Meanwhile, fundamental IIT appears to go dramatically further than empirical IIT, claiming that *any* causally integrated system will realize some form of conscious experience, albeit not necessarily a valenced form (and the IIT group has not yet offered a theory of valence). On this view, not only neural organoids but also kidney organoids, gastrointestinal organoids, cardiac organoids, and other types of organoid (or their constituent parts, if causal integration is greater in the parts than in the whole) would have a form of consciousness. I regard fundamental IIT as a highly speculative position, fair to discuss in the seminar room but unsupported by empirical evidence. There is therefore no reason to expand the range of realistic possibilities to make room for this view.

I am not suggesting that views outside the range R1–R5 can be decisively ruled out with absolute, 100% certainty. What I have in mind is closer to the old idea of 'moral certainty': enough confidence to justify setting aside these views when grave practical questions are at stake. Possibilities in the range R1–R5 have amassed enough evidence to deserve serious consideration when important practical questions are at stake, whereas views outside this range have not.

Many people may hope the current meta-consensus is something we can *move beyond* as new evidence comes to light. This could take the form of a narrowing of the range of realistic possibilities, a widening of that range, or a restructuring of the way we think about the range. But holding such a hope is compatible with accepting that the meta-consensus succeeds in capturing the positions we need to take seriously *now*, given the evidence we have.

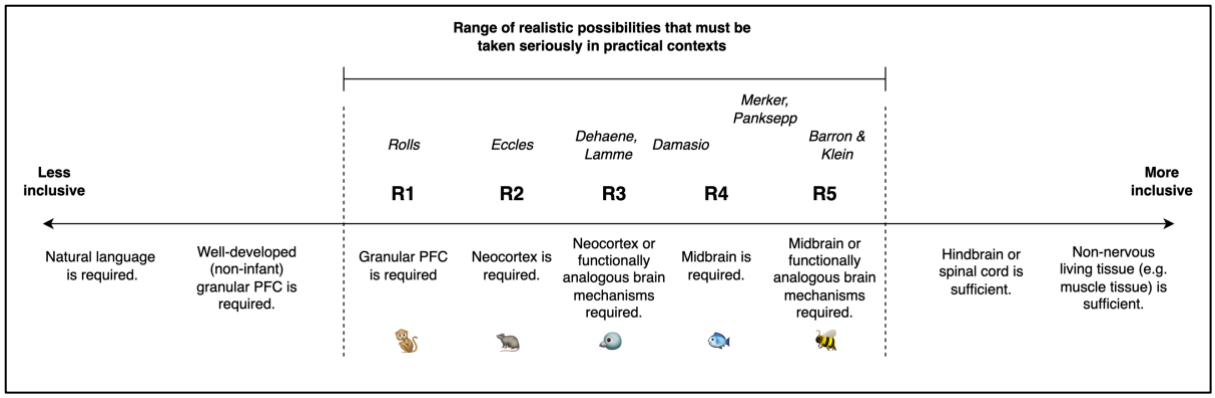

Figure 1: A proposed scientific meta-consensus on the neural basis of sentience. There is no consensus about which position within the reasonable range is correct. However, this is compatible with a meta-consensus forming around the idea that positions R1-R5 are reasonable, given current evidence, provided they are held open-mindedly. Moreover, there can be a meta-consensus on the ordering of these views from less inclusive to more inclusive, and on the challenges facing any view that falls outside this range.

3. The concept of a sentience candidate

From the idea of a scientific meta-consensus, we can construct the concept of a *sentience candidate*:

> A system *S* is a **sentience candidate** if there is an evidence base that:
>
> (a) implies a realistic possibility of sentience in *S* that it would be irresponsible to ignore when making policy decisions that will affect *S*, and
>
> (b) is rich enough to allow the identification of welfare risks and the design and assessment of precautions.

The concept of a sentience candidate is defined in terms of possibilities it would be irresponsible to ignore, given current evidence. There is, inevitably, a value-judgement involved in declaring that evidence has amassed to a point at which it is now irresponsible to ignore it in practical contexts. Judging something to be a sentience candidate is not, therefore, a completely value-neutral exercise.

Yet it is also a judgement that must be informed by the scientific meta-consensus just described. We can appeal to the meta-consensus to explain why disconnected spinal cords, zygotes, neural and non-neural tissue samples, organs other than the brain, non-neural organoids (such as kidney organoids) and unicellular organisms are not sentience candidates. One can speculate, in the seminar room, about sentience in these systems, but responsible precautionary actions cannot be based on these speculations.

To judge a system to be a sentience candidate, then, involves scientific and evaluative components: like many other judgements that have to be made at the science-policy interface,

it is a 'mixed' judgement (Alexandrova 2018; Plutynski 2017). The concept captures a delicate threshold in our evidential and practical situation. When the threshold is crossed, a substantial enough evidence base exists to allow responsible, informed discussion of possible precautionary actions.

How could a system fail to be a sentience candidate? A medically important example of the first type of case is a patient who definitively meets the clinical criteria for brain death. In practice, these criteria test for the irreversible cessation of functional brainstem activity, not literally for the death of the cells in question. It is not straightforward to establish an irreversible loss of brainstem function, leading to continuing debate about the correct criteria (Greer et al., 2020; Walter et al., 2018). But let us focus on a case where the irreversible cessation of functional brainstem activity has been conclusively established. In this case, what remains is not sufficient for sentience on any view in the R1–R5 space. This is why doctors are legally permitted to remove organs and tissues from registered organ donors who are brain dead. This is perhaps the most significant illustration of the idea that a hidden meta-consensus can exist regarding the parameters of reasonable debate when grave issues are at stake. Because we agree that, *if* all brainstem function has irreversibly ceased, the patient is no longer sentient, serious disputes can focus on the question of whether brain death has been accurately determined.

The concept of a sentience candidate is a bridging concept that helps us move from disagreement in the realm of theory to agreement on a course of action. When a being is a sentience candidate, there will be at least one reasonable, scientifically credible basis for taking steps to protect its welfare. That should trigger us to at least start talking about what the reasons *against* might be, and what an all-things-considered *proportionate* response that does justice to the reasons on both sides might look like. By contrast, if a system is not even a sentience candidate, the bar for triggering this process is not cleared.

We can capture this thought in the form of a "Sentience precautionary principle" (intended to be more general than the "Animal Sentience Precautionary Principle" defended in Birch 2017):

> ***Sentience precautionary principle***. If *S* is a sentience candidate, then it is reckless/negligent to make decisions that create risks of suffering for *S* without considering the question of what precautions are proportionate to those risks. Reasonable disagreement about proportionality is to be expected, but we ought to reach a policy decision rather than leaving the matter unresolved indefinitely.

4. Brain organoids: no risk of sentience?

With this precautionary framework in place, let us turn back to brain organoids. I want to start by considering possible reasons to think current neural organoids (at the time of writing) are not sentience candidates. A simple reason often given is their size. This is not a persuasive reason. Bees have around 1 million neurons, and they are sentience candidates. There are

existing brain organoids of a similar size, in terms of neuron count, and researchers aim to create organoids with around 10 million neurons (Smirnova et al., 2023).

A second simple reason, in my view more on-target than the first, is that organoids are not living organisms. They are pieces of tissue, and a default attitude of scepticism towards the idea of sentient tissue, outside of any living animal, is appropriate. Neuroscientists have experimented with small samples of cortical tissue for many years without anyone suggesting a risk of sentience. We must ask: given that cortical tissue samples are not normally sentience candidates, what is different about this type of cortical tissue sample that should cause us to worry? This creates a legitimate default bias against sentience if there is no evidence to the contrary.

Moreover, we should take account of what is missing from present-day organoids. Current neural organoids are typically clusters of cortical neurons, without connections to a functioning brainstem. On Merker's theory, mechanisms at the top of the brainstem, in the midbrain, are constitutively involved in conscious experience (Merker, 2007). Advocates of these theories should be sceptical of the idea of sentience in a neural organoid composed only of cortical tissue. The situation is different when an organoid is implanted into the brain of a host animal (typically a mouse or rat) to create a chimera. These chimeras are clearly sentient, but that is because the host animal is sentient, and the hard question becomes one of how the new tissue alters its cognitive capacities and welfare needs, and there are huge evidence gaps in this area (for commentary on the ethical implications of these evidence gaps, see Birch & Browning, 2021). But in the case of a cortical organoid that is not implanted into a host, midbrain-centric theories give no grounds for attributing sentience.

Here there is an interesting inversion of debates about non-mammalian animals. In the animal case, there is a certain familiar pattern: those who suspect subcortical mechanisms are constitutively involved in consciousness take the possibility of sentience very seriously in a wide range of cases, whereas those who think only neocortical mechanisms are constitutively involved are inclined to play down the risk. Current cortical organoids present us with the opposite situation. They generally lack the subcortical mechanisms taken to be so important by Merker, Panksepp, Solms, Feinberg and Mallatt, Ginsburg and Jablonka, and others. Yet they do have cortical tissue that resembles the neocortical tissue of a developing human brain. So now it is a different family of theories—neocortex-centric theories—that recommend taking the risk of sentience more seriously.

Even defenders of neocortex-centric theories, however, will normally grant a crucial role to the brainstem in supporting conscious experience in humans. The idea is typically that brainstem mechanisms, and in particular the reticular activating system, are akin to a "power cable" for conscious experience, switching it on without being part of its constitutive basis, just as your computer's power cable makes it possible to run a software programme without itself running that programme. Current organoids lack this "power cable" and accordingly display no sleep-wake cycles, to my knowledge.

Indeed, as I understand it, a major limitation of current organoids (when not implanted into host animals) is that they are not fully vascularized: they lack active blood flow. As I write, labs around the world are trying hard to overcome this limitation by joining up neural organoids to vascular organoids, with varying degrees of success (Matsui et al., 2021; Shirure et al., 2021; Sun et al., 2022). We cannot rule out the possibility that fully vascularized organoids will be developed very soon, or even by the time this article is published. But as things stand at this moment, it seems a basic pre-requisite for any cognitive function or conscious experience in a human brain is absent in brain organoids.

5. Early warning signs

For all this, there are concerning signs about the potential for organoid research to accelerate rapidly towards the edge of sentience. In the case of disorders of consciousness, the search for electrophysiological markers of conscious experience has been underway for decades. Synchronized, rhythmic oscillations of local field potentials—informally, brain waves—have long been seen as one of the most important sources of potential markers. Despite a continuing lack of consensus about exactly which oscillations matter, there is widespread consensus about the idea that they are a promising place to look.

Trujillo et al. (2019) allowed cortical organoids to develop for an unusually long period of time, 10 months, and recorded their electrophysiological activity through weekly recordings. They charted the emergence of complex oscillatory waves. They found that organoids quickly settled into a pattern of switching "between long periods of quiescence and short bursts of spontaneous network-synchronized spiking" (Trujillo et al. 2019, p. 562). These synchronized "network events" became stronger and more frequent over time, while the intervals between events became more variable.

This broad pattern of increasingly strong and frequent bursts of activity, with less predictable intervals, is also seen in the EEGs of preterm infants. In an eye-catching result, Trujillo et al. showed that a regression model predicting a neonate's developmental age from key features of its EEG recording, and trained only on data from preterm infants, could also judge the developmental age of organoids older than 25 weeks with above-chance accuracy, with moderate correlation between the predicted and actual ages.

The result must be carefully interpreted. This does not show that the organoids were in any sense equivalent to the brains of preterm infants. It is important to note, first of all, that these cortical organoids were not *brains* at all. We should take care to avoid terms such as "mini-brain" for systems like these. The organoids were formed of a single type of tissue—cortical tissue—representative of one particularly important brain region, the neocortex. The organoids were vastly smaller than an infant brain, and still lacked a brainstem and vascularization. Nor does it show that the electrophysiological activity was the same or indistinguishable in the two cases. The regression model aimed to exploit the similarities that existed, not quantify the

degree of similarity. The model identified enough similarities to inform above-chance predictions of developmental age, but this is compatible with substantial differences.

Nonetheless, the result was, to me, a wake-up call: a jolt out of complacency about the potential ethical implications of this research. Brain organoids develop, they are sometimes allowed to develop for a long time, and they develop in ways that show broad electrophysiological similarities to the developing human brain.

6. Assessing sentience candidature in brain organoids

We cannot rule out the possibility that sufficiently sophisticated organoids will soon be sentient, and we can expect the science to continue to develop extremely rapidly. So, we need to have a discussion now about what sort of warning signs might suffice to regard an organoid as a sentience candidate.

Here we run into a serious problem. In people with prolonged disorders of consciousness (another difficult case), some behaviour remains, despite the tendency to describe patients as "unresponsive", and that behaviour informs diagnosis and the design of precautions (Johnson, 2022). Clinicians (in the UK) are already advised to respond to outward signs of pain, distress, anxiety and depression on the precautionary assumption that they really do indicate those states. The behaviour may be involuntary much of the time, but it is behaviour nonetheless. Sleep-wake cycles are also present, marking a clear distinction with coma. Meanwhile, in the case of non-human animals, the most compelling and widely accepted markers of sentience again tend to be behavioural. Animal welfare experts have formulated lists of such markers, generally focusing on pain (Birch et al., 2021; Sneddon et al., 2014). Organoids present a very different kind of challenge. None of these behavioural markers of sentience are likely to be present in a typical brain organoid, because organoids are typically cut off from the sources of sensory input and motor output that are available to a complete and developed organism, and I assume this is likely to remain typical in the near-term future.

This could turn out to be an incorrect assumption. Some future organoids, even in the near term, may well have sources of sensory input and motor output. For example, a recent study showed that under the right conditions a cortical organoid can spontaneously develop optic vesicles—the developmental precursors to eyes—and it is not yet known how far this process could go, as the technology develops (Gabriel et al., 2021). Another study allowed organoids to develop in culture for a year, placed near to a spinal cord and muscle tissue taken from a mouse. The organoids "were able to innervate mouse spinal cord" and "evoke contractions of adjacent muscle" (2019), p. 669).

On this evidence, a time when organoid preparations can be joined up to both muscle outputs and sensory inputs is not far off. At that point, public concern about the research may grow. At the same time, using behavioural criteria to assess the likelihood of sentience may also become more feasible, providing a new way in which public concern could be exacerbated or at least

slightly eased, depending on the results. Negative results would still require very cautious interpretation, because a failure to display sentience-related behaviours could easily reflect a failure of coordinated muscle control and a very limited behavioural repertoire rather than a lack of sentience.

Sentience, then, may be both more likely and easier to attribute when a neural organoid is joined up to other tissues, be they themselves organoids or taken from animals. But let us focus for now on the case of a "pure" brain organoid, disconnected from any other tissues and any sources of sensory input or motor output. This is the type of case that presents the deepest puzzle. If the system is sentient, then it is what Bayne, Seth and Massimini (2020) have called an "island of awareness", unable to manifest its sentience in any of the usual ways. In this case, there is no behaviour, so we need to assess sentience candidature using only non-behavioural markers. Where do we even begin?

7. The brainstem rule

There is one important piece of common ground in this area. All reasonable views compatible with the scientific meta-consensus can agree that, in a human brain, there can be no sentience in the absence of a functioning brainstem. Agreement *that* this is the case is much wider than agreement about *why* it is the case. For the midbrain-centric family of theories, mechanisms at the top of the brainstem are sufficient for sentience without a cortex. For the cortex-centric family, midbrain mechanisms are causally but not constitutively involved. They help regulate the global state of consciousness without being part of its neural basis.

All parties can agree, however, that sentience in humans depends on brainstem activity. Without a living brainstem, a human cannot maintain coordinated patterns of global cortical activity, integrative subcortical activity or sleep-wake cycles. Theorists from right across the zone of reasonable disagreement are able to agree that irreversible loss of brainstem function implies the irreversible loss of consciousness. The main challenge is determining exactly when "irreversible" loss has occurred.

We should add a caveat in the interests of future-proofing. Strictly speaking, what is required is a functioning brainstem *or a functionally equivalent system* that regulates arousal and supports sleep-wake cycles. It could be that, in the distant future, artificial brainstems will be created to allow people to recover from currently irrecoverable brain injuries. Such a person would clearly be a sentience candidate, despite lacking a biological brainstem. This is a long way off, but what may be much closer is the possibility of a small-scale functional equivalent that is able to regulate the activity of an organoid in the same way a brainstem would. Even in the absence of a biological brainstem, we should be wary of the risks posed by attempts to use artificial brainstem-like systems to regulate and coordinate cortical activity in organoids.

This common ground is at the root of the widespread view that current organoids are *not* sentience candidates. But it also gives us one threshold for the point at which organoids *will*

become sentience candidates. If an organoid (or assembloid) is developed that has a functioning brainstem or artificial substitute that regulates arousal and leads to sleep-wake cycles then, no matter how small it is, it should be regarded as a sentience candidate. There would be at least one view within the zone of reasonable disagreement (namely a midbrain-centric view along the lines of Panksepp, Merker and Solms) on which such a system would be likely to be sentient. The outward signs of regulated arousal and sleep-wake cycles would, in this context, be indicators that the conditions Panksepp/Merker/Solms regard as sufficient for sentience are plausibly in place.

We can call this proposal the "brainstem rule":

> ***Brainstem rule*:** If a neural organoid develops or innervates a functioning brainstem (including the midbrain) that regulates arousal and leads to sleep-wake cycles, then it is a sentience candidate. A functional equivalent of a brainstem (even if artificial) would also suffice.

This is proposed as a *sufficient condition for sentience candidature*. To be clear, it is not proposed as a sufficient condition for sentience (the Panksepp/Merker/Solms view is a realistic possibility, not a certainty), nor is it proposed as a necessary condition for sentience candidature. The idea is that, when the condition is satisfied, we are in a situation in which we can no longer have confidence that sentience is absent and so should start considering precautions. The proposal leaves open the possibility that there may be *other* scenarios in which we should consider precautions. I am describing here a route to sentience candidature that runs via taking midbrain-centric theories of consciousness seriously, but there may well be other routes, running via different theories.

The proposal says "develops or innervates", highlighting two different ways in which an organoid could acquire a functioning brainstem. One is spontaneous development, along the lines of the optic vesicles spontaneously developed by an organoid in the Gabriel et al. (2021) study. The other is through innervating animal tissue, along the lines of the innervation of a spinal cord by an organoid in the Giandomenico et al. (2019) study.

We may well find that future model systems in neuroscience increasingly blur the boundary between organoids and chimeras, as more and more living brain tissue from a host animal is used in mixed human-animal "preparations". One can imagine a future variation on Giandomenico et al. (2019) that takes the whole living brainstem from a mouse, not just the spinal cord, and connects it to an organoid. Such a system may realistically possess the midbrain mechanisms that lead us to regard humans with conditions such as hydranencephaly as sentience candidates. So, the pressure of consistency should push us towards regarding this system as a sentience candidate too.

A controversial aspect of the proposal is that it implies a system that is clearly *not* a complete, embodied living organism *can* nonetheless be a sentience candidate. That is indeed what I am proposing. The brainstem rule says, in effect, that what I earlier called a "legitimate default

bias" against the idea of a sentient non-organism should not be allowed to become a hard assumption that leads us to neglect risks. Although we can be confident that destroying the brainstem "pulls the plug" (so to speak) on sentience in a living organism, we have no right to be similarly confident that taking away the rest of the body while leaving the brainstem fully functional would have the same effect. Disembodied brain organoids with functioning brainstems would be intrinsically similar systems, reached by building up rather than by stripping away.

8. Possible regulatory frameworks

The proposed "brainstem rule" leaves open what would be a proportionate response to an organoid's sentience candidature. It may be tempting to think: even if an organoid is sentient, it is at no serious risk of harm, because harm requires nociception and a capacity for bodily sensation. This, however, would be too hasty. Think, for example, of phantom pain: we know that, in adult humans, the brain mechanisms associated with pain can be triggered in the absence of a physical stimulus (Culp & Abdi, 2022). We should take seriously the risk of an organoid developing versions of the pain pathways of a normal human brain (and the pathways linked to other negatively valenced states—thirst, hunger, cold, etc.) without their usual bodily inputs, leading a risk of these mechanisms being activated unpredictably by the environment.

Among the possible responses to the risk of harm are a moratorium (time-limited ban) or an indefinite (non-time-limited) ban on the creation of these particular organoids. I say "indefinite" rather than "permanent" because governments are not able to bind their successors, so there can never be a guarantee that a ban will be permanent. I take these seriously as options that may be proportionate, and I resist the idea that they would amount to drastic or radical restrictions on biomedical research. They should be options that are on the table when we debate these issues.

There is, after all, a huge amount of valuable research that can be done on organoids without getting anywhere near the edge of sentience. Researchers could invest their time in simpler neural organoids or in non-neural organoids, such as kidney organoids and gastrointestinal organoids. A similar line of reasoning is often considered plausible in relation to embryos past the legal age limit (14 days in the UK). Yes, we could learn much from research on older embryos, but it is not in keeping with our values to run even a small risk of creating sentient human embryos solely for the purpose of experimentation, and there are many other valuable kinds of research we can prioritize instead, so we should be willing to forego the benefits. The key would be to ensure that the ban is targeted, so that lower-risk forms of organoid research are allowed to continue. An indiscriminate ban on all organoid research would be excessive and disproportionate. It would give no weight to the great promise of organoid research as a potential substitute for research on whole animals.

A less stringent response would be to allow research on sentience candidates, but subject this research to a licensing regime modelled on that of animal research. After all, most animals used

in research are sentience candidates (like insects) or sentient as a matter of consensus (like rats and mice). As a society, we permit this research even though it implies some level of suffering to sentient beings. Where research on a potentially sentient organoid might replace research on a clearly sentient animal, like a mouse or a rat, and might even be preferable on scientific grounds, consistency suggests we should at least try to weigh up the harms and benefits of the two possible projects, rather than always favouring animal research. An indiscriminate bias in favour of research on whole sentient animals rather than merely potentially sentient organoids is unwarranted.

The "weighing" here is, admittedly, very challenging, because our uncertainty about the nature of sentience is so severe that we cannot put precise, agreed probabilities on the chance of sentience in an organoid, and opinions vary widely. We cannot expect ethical review bodies to weigh the risks precisely—but I think we can expect them to weigh risk in broad, qualitative terms, and to debate whether imposing clear, known harms on clearly sentient animals is any easier to justify than imposing somewhat speculative potential harms on organoids that are merely sentience candidates.

This line of thought led me to suggest, in a piece with Heather Browning, that we should look to include potentially sentient organoids within the scope of animal experimentation legislation, such as the UK's Animals (Scientific Procedures) Act 1986, commonly known as "ASPA" (Birch & Browning, 2021). This would certainly be more appropriate than treating potentially sentient organoids as mere tissue, and also more appropriate than treating them as if they were whole embryos, when they are not.

Under ASPA, scientists proposing research projects with the potential to cause suffering to animals have to obtain a licence for the work. To be licensed, they need approval from an institutional ethical review board. The board needs to see that the scientists have carefully weighed harms and benefits and duly considered the imperative to reduce, refine, and replace. In this context, "replace" might mean the replacement of work on potentially sentient organoids with work on organoids that lack any brainstem structures and are less likely to be sentient. Researchers should be expected to make a case that they need to create a sentience candidate, and not just a simpler organoid system, to achieve the biomedical goals of the work. The ethical review board should consider whether those goals genuinely make the proposed research justifiable, and whether proportionate steps have been taken to mitigate the risks of causing suffering.

Plainly, it would be controversial to bring a form of human tissue under regulations designed for animal research, for two reasons: we are talking about tissue and not about whole animals, and we are talking about human tissue, not the tissue of other animals. In both ways, the proposal involves extending a general regulatory approach outside the context for which it was originally devised. However, I see the problems here as problems of framing and wording, not deep problems. If ASPA were to be amended to include organoids, it would be wise to rename it. Politically, it may be wiser to regulate organoid research using new legislation modelled on ASPA rather than through amending ASPA itself.

I see both of the above options—an indefinite ban or moratorium targeted at specific types of organoid, and a regulatory framework modelled on ASPA and centred on the idea of harm-benefit analysis—as options worthy of serious discussion. Which option we take depends on broader evaluative questions about the value we see, as a society, in this research, relative to the disvalue of the risks. We may also want to use both options in relation to different types of brain organoid, regulating research on some, banning research on others. Moreover, in cases where organoid research can replace kinds of animal research that are *more* harmful, targeted bans on the relevant kinds of animal research should *also* be on the table. I doubt there will be a one-size-fits-all solution, and for now I want to put both proposals on the table as options that should be debated further.

> ***Response 1 (targeted bans):*** If organoid research leads to the creation of organoids that are sentience candidates, a moratorium (time-limited ban) or indefinite ban on the creation of this particular type of organoid may be an appropriate response. Bans should avoid indiscriminate targeting of all organoid research.
>
> ***Response 2 (ethical review):*** When a neural organoid is a sentience candidate, research on it, if permitted at all, should be subject to ethical review and harm-benefit analysis, modelled on existing frameworks for regulating research on sentient animals.

To be clear, the proposals in this paper are independent of each other. So, one may still agree that my proposed responses are on the right lines even if one thinks the "brainstem rule" sets the bar in the wrong place, and *vice versa*.

9. Conclusion

To summarise the overall argument: human brain organoids are showing great promise as models of the human brain, models that could potentially replace a substantial amount of animal research. It would be hasty to dismiss the possibility they could develop sentience, (defined as the capacity for conscious experiences with a positive or negative quality). However, scepticism about this idea is appropriate when considering current organoids (at the time of writing). This is not because of their size, but because of their organization. It is a point of consensus across reasonable views that a brainstem-dead human is not sentient, and current organoids lack a functioning brainstem or anything equivalent to one. There are nonetheless some troubling early warning signs, suggesting that organoid research may create forms of sentient being in the future.

Researchers with very different views about the neural basis of sentience can unite behind the "brainstem rule": if a neural organoid develops or innervates a functioning brainstem that regulates arousal and leads to sleep-wake cycles, then it is a sentience candidate. This is proposed as a *sufficient condition for sentience candidature*. When a system is a sentience

candidate, we should take the possibility of its sentience seriously and discuss proportionate steps to protect its welfare, despite continuing uncertainty and doubt.

What steps might be proportionate? If organoid research leads to the creation of organoids that are sentience candidates, a moratorium (time-limited ban) or indefinite ban on the creation of this particular type of organoid may be appropriate, but bans should avoid indiscriminate targeting of all organoid research. An alternative approach, consistent with existing approaches to animal research, is to require ethical review and harm-benefit analysis whenever a brain organoid is a sentience candidate.

Acknowledgements

I thank Tim Bayne, Simon Brown, Heather Browning, Andrew Crump, Katha Dornenzweig, Sarah Diner, Charlotte Gauvry, Syd Johnson, Takuya Niikawa, Lorenzo Sartori, Tomasz Żuradzki, and the members of my Edge of Sentience Reading Group for their comments and advice. This research is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, Grant Number 851145.

References

Alexandrova, A. (2018). Can the science of well-being be objective? *British Journal for the Philosophy of Science, 69*(2), 421–445. https://doi.org/10.1093/bjps/axw027

Ankeny, R. A., & Wolvetang, E. (2021, 2021/01/02). Testing the Correlates of Consciousness in Brain Organoids: How Do We Know and What Do We Do? *The American Journal of Bioethics, 21*(1), 51-53. https://doi.org/10.1080/15265161.2020.1845869

Arikan, B. E., van Kemenade, B. M., Podranski, K., Steinsträter, O., Straube, B., and Kircher, T. (2019). Perceiving your hand moving: BOLD suppression in sensory cortices and the role of the cerebellum in the detection of feedback delays. *Journal of Vision, 19*(14), 4. https://doi.org/10.1167/19.14.4

Aru, J., Suzuki, M., and Larkum, M. E. (2020). Cellular mechanisms of conscious processing. *Trends in Cognitive Sciences*, 24(10), 814–825. https://doi.org/10.1016/j.tics.2020.07.006

Barron, A. B., & Klein, C. (2016). What insects can tell us about the origins of consciousness. *Proceedings of the National Academy of Sciences, 113*(18), 4900-4908. https://doi.org/doi:10.1073/pnas.1520084113

Bayne, T., Seth, A. K., & Massimini, M. (2020). Are There Islands of Awareness? *Trends in Neurosciences, 43*(1), 6-16. https://doi.org/10.1016/j.tins.2019.11.003

Beck, F., & Eccles, J. C. (1992). Quantum aspects of brain activity and the role of consciousness. *Proceedings of the National Academy of Sciences, 89*(23), 11357-11361. https://doi.org/doi:10.1073/pnas.89.23.11357

Birch, J. (2017). Animal sentience and the precautionary principle. *Animal sentience, 2*(16), 1.

Birch, J., & Browning, H. (2021, 2021). Neural Organoids and the Precautionary Principle. *The American Journal of Bioethics, 21*(1), 56-58. https://doi.org/10.1080/15265161.2020.1845858

Birch, J., Burn, C., Schnell, A., Browning, H., & Crump, A. (2021). *Review of the evidence of sentience in cephalopod molluscs and decapod crustaceans*. LSE Consulting.

Culp, C. J., & Abdi, S. (2022). Current understanding of phantom pain and its treatment. *Pain Physician*, *25*(7), E941–E957.

Damasio, A., Damasio, H., and Tranel, D. (2013). Persistence of feelings and sentience after bilateral damage of the insula. *Cerebral Cortex*, *23*(4), 833–846. https://doi.org/10.1093/ cercor/bhs077

Damasio, A. R., Grabowski, T. J., Bechara, A., Damasio, H., Ponto, L. L., Parvizi, J., and Hichwa, R. D. (2000). Subcortical and cortical brain activity during the feeling of self-generated emotions. *Nature Neuroscience*, *3*(10), 1049–1056.

Dryzek, J. S. (2010). *Foundations and Frontiers of Deliberative Governance*. Oxford University Press.

Farris, S. M. (2020). The rise to dominance of genetic model organisms and the decline of curiosity-driven organismal research. *PLoS One, 15*(12), e0243088. https://doi.org/10.1371/journal.pone.0243088

Feinberg, T. E., & Mallatt, J. M. (2016). *The Ancient Origins of Consciousness: How the Brain Created Experience*. MIT Press.

Gabriel, E., Albanna, W., Pasquini, G., Ramani, A., Josipovic, N., Mariappan, A., Schinzel, F., Karch, C. M., Bao, G., Gottardo, M., Suren, A. A., Hescheler, J., Nagel-Wolfrum, K., Persico, V., Rizzoli, S. O., Altmüller, J., Riparbelli, M. G., Callaini, G., Goureau, O., Papantonis, A., Busskamp, V., Schneider, T., & Gopalakrishnan, J. (2021, Oct 7). Human brain organoids assemble functionally integrated bilateral optic vesicles. *Cell Stem Cell, 28*(10), 1740-1757.e1748. https://doi.org/10.1016/j.stem.2021.07.010

Giandomenico, S. L., Mierau, S. B., Gibbons, G. M., Wenger, L. M. D., Masullo, L., Sit, T., Sutcliffe, M., Boulanger, J., Tripodi, M., Derivery, E., Paulsen, O., Lakatos, A., & Lancaster, M. A. (2019, 2019/04/01). Cerebral organoids at the air–liquid interface generate diverse nerve tracts with functional output. *Nature Neuroscience, 22*(4), 669-679. https://doi.org/10.1038/s41593-019-0350-2

Ginsburg, S., & Jablonka, E. (2019). *The Evolution of the Sensitive Soul: Learning and the Origins of Consciousness*. MIT Press.

Greer, D. M., Shemie, S. D., Lewis, A., Torrance, S., Varelas, P., Goldenberg, F. D., Bernat, J. L, Souter, M., Topcuoglu, M. A., Alexandrov, A. W., Baldisseri, M., Bleck, T., Citerio, G., Dawson, R., Hoppe, A., Jacobe, S., Manara, A., Nakagawa, T. A., Pope, T. M. . . . and Sung, G. (2020). Determination of brain death/death by neurologic criteria: the World Brain Death Project. *Journal of the American Medical Association*, *324*(11), 1078–1097. https://doi.org/10.1001/jama.2020.11586

Herculano-Houzel, S. (2009, 2009-November-09). The human brain in numbers: a linearly scaled-up primate brain [Review]. *Frontiers in Human Neuroscience*, *3*. https://doi.org/10.3389/neuro.09.031.2009

Hubrecht, R. C., & Carter, E. (2019). The 3Rs and humane experimental technique: Implementing change. *Animals*, *9*(10), 754. https://doi.org/10.3390/ani9100754

Humphrey, N. (2022). *Sentience: The Invention of Consciousness*. Oxford University Press.

Hyun, I., Scharf-Deering, J. C., & Lunshof, J. E. (2020, Apr 1). Ethical issues related to brain organoid research. *Brain Res*, *1732*, 146653. https://doi.org/10.1016/j.brainres.2020.146653

Johnson, J. F., Belyk, M., Schwartze, M., Pinheiro, A. P., & Kotz, S. A. (2019). The role of the cerebellum in adaptation: ALE meta-analyses on sensory feedback error. *Human Brain Mapping*, *40*(13), 3966–3981. https://doi.org/10.1002/hbm.24681

Johnson, L. S. M. (2022). *The Ethics of Uncertainty: Entangled Ethical and Epistemic Risks in Disorders of Consciousness*. Oxford University Press.

Kagan, B. J., Kitchen, A. C., Tran, N. T., Habibollahi, F., Khajehnejad, M., Parker, B. J., Bhat, A., Rollo, B., Razi, A., & Friston, K. J. (2022, 2022/12/07/). In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, *110*(23), 3952-3969.e3958. https://doi.org/https://doi.org/10.1016/j.neuron.2022.09.001

Key, B. (2015). Fish do not feel pain and its implications for understanding phenomenal consciousness. *Biology & Philosophy*, *30*(2), 149–165. https://doi.org/10.1007/s10539-014-9469-4

Koch, C. (2019, 24 September, 2019.). Consciousness Doesn't Depend on Language. *Nautilus*.

Kreitmair, K. (2023). Consciousness and the Ethics of Human Brain Organoid Research. *Cambridge Quarterly of Healthcare Ethics*, 1-11. https://doi.org/10.1017/S0963180123000063

Lamme, V. A. F. (2022). Behavioural and neural evidence for conscious sensation in animals: an inescapable avenue towards biopsychism? *Journal of Consciousness Studies*, *29*(3–4), 78–103. https://doi.org/10.53765/20512201.29.3.078

Lavazza, A. (2020, 2020). Human cerebral organoids and consciousness: a double-edged sword. *Monash Bioethics Review*, *38*(2), 105-128. https://doi.org/10.1007/s40592-020-00116-y

Lavazza, A., & Massimini, M. (2018, Sep). Cerebral organoids: ethical issues and consciousness assessment. *Journal of Medical Ethics*, *44*(9), 606-610. https://doi.org/10.1136/medethics-2017-104555

LeDoux, J. E. (2023). Deep history and beyond: a reply to commentators. *Philosophical Psychology*, 36(4), 756–766. https://doi.org/10.1080/09515089.2022.2160312

LeDoux J. E., Birch, J., Andrews, K., Clayton, N. S., Daw, N. D., Frith, C., Lau, H., Peters, M. A. K., Schneider, S., Seth, A., Suddendorf, T., and Vandekerckhove, M. M. P. (2023). Consciousness beyond the human case. *Current Biology*, *33*(16), R832–R840. https://doi. org/10.1016/j.cub.2023.06.067

Matsui, T. K., Tsuru, Y., Hasegawa, K., & Kuwako, K.-i. (2021). Vascularization of human brain organoids. *STEM CELLS*, *39*(8), 1017-1024. https://doi.org/https://doi.org/10.1002/stem.3368

Mediano, P. A. M., Rosas, F. E., Bor, D., Seth, A. K., and Barrett, A. B. (2022). The strength of weak integrated information theory. *Trends in Cognitive Sciences*, *26*(8), 646–655. https://doi.org/10.1016/j.tics.2022.04.008

Mediano, P. A. M., Seth, A. K., and Barrett, A. B. (2019). Measuring integrated information: Comparison of candidate measures in theory and simulation. *Entropy*, *21*(1), 1–30. https://doi.org/10.3390/e21010017

Merker, B. (2007). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. *Behavioral and Brain Sciences*, *30*(1), 63-81. https://doi.org/10.1017/S0140525X07000891

Michel, M., and Lau, H. (2020). On the dangers of conflating strong and weak versions of a theory of consciousness. *Philosophy and the Mind Sciences*, *1*(2), 8. https://doi.org/10.33735/phimisci.2020.II.54

Niikawa, T., Hayashi, Y., Shepherd, J., & Sawai, T. (2022, 2022). Human Brain Organoids and Consciousness. *Neuroethics*, *15*(1), 5. https://doi.org/10.1007/s12152-022-09483-1

Panksepp, J. (1998). *Affective Neuroscience: The Foundations of Human and Animal Emotions*. Oxford University Press.

Plutynski, A. (2017). Safe or sorry? Cancer screening and inductive risk. In K. C. Elliott and T. Richards (eds.), *Exploring Inductive Risk: Case Studies of Values in Science* (pp. 149-170). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780190467715.003.0008

Preuss, T. M., and Wise, S. P. (2022). Evolution of prefrontal cortex. *Neuropsychopharmacology*, *47*(1), 3–19. https://doi.org/10.1038/s41386-021-01076-5

Rolls, E. T. (2004). A higher order syntactic thought (HOST) theory of consciousness. In R. J. Gennaro (ed.), *Higher Order Theories of Consciousness* (pp. 137–172). John Benjamins.

Rolls, E. T. (2014). *Emotion and Decision-Making Explained*. Oxford University Press.

Sawai, T., Sakaguchi, H., Thomas, E., Takahashi, J., & Fujita, M. (2019, Sep 10). The Ethics of Cerebral Organoid Research: Being Conscious of Consciousness. *Stem Cell Reports, 13*(3), 440-447. https://doi.org/10.1016/j.stemcr.2019.08.003

Seth, A. K., & Bayne, T. (2022, 2022/07/01). Theories of consciousness. *Nature Reviews Neuroscience, 23*(7), 439-452. https://doi.org/10.1038/s41583-022-00587-4

Sharma, A., Zuk, P., & Scott, C. T. (2021, 2021/01/02). Scientific and Ethical Uncertainties in Brain Organoid Research. *The American Journal of Bioethics, 21*(1), 48-51. https://doi.org/10.1080/15265161.2020.1845866

Shemesh, Y., & Chen, A. (2023, 2023/03/01). A paradigm shift in translational psychiatry through rodent neuroethology. *Molecular Psychiatry, 28*(3), 993-1003. https://doi.org/10.1038/s41380-022-01913-z

Shirure, V. S., Hughes, C. C. W., & George, S. C. (2021). Engineering Vascularized Organoid-on-a-Chip Models. *Annual Review of Biomedical Engineering, 23*(1), 141-167. https://doi.org/10.1146/annurev-bioeng-090120-094330

Smirnova, L., Caffo, B. S., Gracias, D. H., Huang, Q., Morales Pantoja, I. E., Tang, B., Zack, D. J., Berlinicke, C. A., Boyd, J. L., Harris, T. D., Johnson, E. C., Kagan, B. J., Kahn, J., Muotri, A. R., Paulhamus, B. L., Schwamborn, J. C., Plotkin, J., Szalay, A. S., Vogelstein, J. T., Worley, P. F., & Hartung, T. (2023, 2023-February-28). Organoid intelligence (OI): the new frontier in biocomputing and intelligence-in-a-dish [Frontiers in Science Lead Article]. *Frontiers in Science, 1*. https://doi.org/10.3389/fsci.2023.1017235

Sneddon, L. U., Elwood, R. W., Adamo, S. A., & Leach, M. C. (2014, 2014/11/01/). Defining and assessing animal pain. *Animal Behaviour, 97*, 201-212. https://doi.org/https://doi.org/10.1016/j.anbehav.2014.09.007

Solms, M. (2021). *The Hidden Spring: A Journey to the Source of Consciousness*. Profile Books.

Steel, D. (2015). *Philosophy and the Precautionary Principle*. Cambridge University Press.

Sun, X.-Y., Ju, X.-C., Li, Y., Zeng, P.-M., Wu, J., Zhou, Y.-Y., Shen, L.-B., Dong, J., Chen, Y.-J., & Luo, Z.-G. (2022, 2022/05/04). Generation of vascularized brain organoids to study neurovascular interactions. *eLife, 11*, e76707. https://doi.org/10.7554/eLife.76707

Taschereau-Dumouchel, V., Michel, M., Lau, H., Hofmann, S. G., & LeDoux, J. E. (2022, 2022/03/01). Putting the "mental" back in "mental disorders": a perspective from research on fear and anxiety. *Molecular Psychiatry, 27*(3), 1322-1330. https://doi.org/10.1038/s41380-021-01395-5

Taylor, K., & Alvarez, L. R. (2019). An Estimate of the Number of Animals Used for Scientific Purposes Worldwide in 2015. *Alternatives to Laboratory Animals, 47*(5-6), 196-213. https://doi.org/10.1177/0261192919899853

Tononi, G., Albantakis, L., Boly, M., Cirelli, C., and Koch, C. (2023). Only what exists can cause: an intrinsic view of free will. arXiv preprint. arXiv:2206.02069v3

Trujillo, C. A., Gao, R., Negraes, P. D., Gu, J., Buchanan, J., Preissl, S., Wang, A., Wu, W., Haddad, G. G., Chaim, I. A., Domissy, A., Vandenberghe, M., Devor, A., Yeo, G. W., Voytek, B., & Muotri, A. R. (2019). Complex Oscillatory Waves Emerging from Cortical Organoids Model Early Human Brain Network Development. *Cell Stem Cell, 25*(4), 558-569.e557. https://doi.org/10.1016/j.stem.2019.08.002

Tye, M. (2016). *Tense Bees and Shell-Shocked Crabs: Are Animals Conscious?* Oxford University Press.

Wilson, M.N., Thunemann, M., Liu, X. et al. (2022). Multimodal monitoring of human cortical organoids implanted in mice reveal functional connection with visual cortex. *Nature Communications*, *13*, 7945. https://doi.org/10.1038/s41467-022-35536-3

Yu, F., Jiang, Q.-j., Sun, X.-y., & Zhang, R.-w. (2014). A new case of complete primary cerebellar agenesis: clinical and imaging findings in a living patient. *Brain, 138*(6), e353-e353. https://doi.org/10.1093/brain/awu239

Żuradzki, T. (2021, 2021/01/02). Against the Precautionary Approach to Moral Status: The Case of Surrogates for Living Human Brains. *The American Journal of Bioethics, 21*(1), 53-56. https://doi.org/10.1080/15265161.2020.1845868