AVANT, Vol. XV, No. 1, ISSN 2082-6710
https://doi.org/10.26913/ava1202401

We don't need a Science of Animal Consciousness. On the Unexpected Insights gained from Walter Veit's *A Philosophy for the Science of Animal Consciousnes*

Maja Kittel (f. Białek)
Institute of Philosophy, University of Białystok
majasbialek@gmail.com

Received 2024-03-26; accepted 2024-05-01; published Online First 2024-07-05.

Abstract

This paper is a critical commentary on Walter Veit's book *A philosophy for the science of animal consciousness*. My goal is to show that although Veit succeeds in presenting a compelling account of animal consciousness, he may have unintentionally undermined the purpose of such science. I argue that, despite the author's claims, his theory is not as empirically grounded as he makes it out to be. Paradoxically, some of Veit's arguments against his opponents seem to be double-edged. Still, I also argue that this is a necessary cost, as an alternative would be to succumb to a particular form of metatheoretical anthropomorphism. In effect, we end up with a surprising conclusion that we need many sciences of different consciousnesses, which seems at odds with Veit's Darwinian continuity assumption. However, this only goes to show that the book is well worth discussing.

Keywords: animal consciousness, Walter Veit, animal cognition, evolution of consciousness, science of consciousness, philosophy of mind, theories of consciousness, consciousness

1. Introduction

The explosion of interest in non-human animal consciousness has resulted in a deluge of literature, scientific and popular, which unfortunately not only fuels fervent scientific discussions, but also leads to many misunderstandings and conceptual chaos. Veit tempts us with a straight, clear path through the quagmire: a coherent story of animal consciousness based on the work done by evolutionary cognitive scientists, comparative psychologists and philosophers. This path also extends into the future, offering inspiration and direction for researchers. In this paper, I will show what delights lie ahead, if we succumb to Veit's temptation, and acknowledge the many metatheoretical virtues of his approach. However, there is a catch. Veit does not deliver exactly what he promises. We are certainly left with a better understanding of certain aspects of consciousness, but his account is far more speculative and much less grounded in empirical science than he makes us believe. Moreover, it raises rather unexpected questions about the status of our explanations and the science of animal consciousness in general.

2. An outline of Veit's argument in context

2.1. Metatheoretical choices

Veit's goal is to lay the philosophical groundwork for "the science of animal consciousness", which requires making several metatheoretical decisions, first of which concerns our views on the scope and status of this science, and especially its relation to the science of human consciousness. Although the author abstains from making any sweeping and probably unnecessary judgments about the advances the science of animal consciousness has made so far, it is obvious why his undertaking seems both worthwhile and timely. The science of animal consciousness is still in its nascent stage, operating without any unified, interdisciplinary paradigms or theories—it is more of a common dream shared by a diverse group of scientists exploring different fields. This dream of understanding consciousness is still elusive even as far as human animals go, despite our unique knowledge of our own first-person experience of the phenomenon (or perhaps because of it). Above all, we lack a comprehensive, universally accepted conceptual framework. The key term 'consciousness' is being used in many senses, by researchers with vastly different backgrounds and methodologies, while at the same time it remains deeply rooted in our folkpsychological outlook on the world with all the connotations this outlook brings. This is partly why even the topic of human consciousness seems hopelessly complex. In fact, it sometimes feels more complex that the theories concerning non-human animal consciousness. It may be this is because so much more detailed work has been done on human consciousness, both empirical and theoretical, and, hence, so much more conundrums have come to light, but it also

may be a sign of a deeper difference between the two phenomena. The questions about non-human consciousness probably sound less convoluted because they are grounded in a clear ethical and practical context, which is, no doubt, also important for Veit. We are usually much more down to earth when discussing non-human animal consciousness—we often ask yes or no questions (are certain animals conscious or not?) or seek specific answers pertaining directly to the animals' wellbeing, for example if they experience pain or what are the limitations of their senses.

However, moving beyond those practical questions is tricky. We still don't have many ideas how to investigate possible more advanced forms of reportable conscious experiences in animals that don't communicate with us in a way that would be both complex enough and understandable enough. Some philosophers would even claim that because of this, we are justified in claiming that either only humans with advanced conceptual capacities can be ascribed consciousness or that only such humans are endowed with a particularly "high" or "special" form of consciousness, incomparable with anything we could find in other sentient creatures. This line of thinking would lead to expecting a discontinuity in the science of animal consciousness—we would need two such sciences, one for whatever other animals have (and we probably share), and an additional one for those advanced, human forms. We would also go against Darwin, who argued for full evolutionary continuity, and probably accept some form of human exceptionalism.

Veit's explicit choice to follow Darwin (and Donald Griffin) is his first metatheoretical move. By framing it as an alternative, I don't wish to imply that I find both choices equally valid; in fact, for reasons of space, I will not even be discussing why the Darwinian position is far more plausible. I only do it to point out that Veit's decision to follow Darwin brings not only direct theoretical commitments of the theory of evolution, but also powerful metatheoretical obligations. We will be expecting a "continuity" approach to the science of consciousness, fostering such theories that can explain the evolutionary path of this ability that would also involve the forms it takes in humans. Our explanations of the simplest mechanisms of consciousness in non-human animals should shed light on what happens in humans as well and vice versa.

Veit's second metatheoretical move is a consequence of his first: as a Darwinian, Veit focuses on a fascinating question which surprisingly has been receiving too little attention: what is consciousness for? Moreover, our Darwinian approach offers us a constraint on possible answers straight away—they must explain consciousness's evolutionary role, they must show how it helps organisms to survive and reproduce; in short: they must be teleonomic. The requirement of pointing out the adaptive function is what Veit calls "a Darwinian bottleneck" for theories of consciousness and, as is so often the case with clear requirements in science, it makes the job of explaining consciousness easier in many ways. It drives a more focused approach to the other puzzling questions

about consciousness, most notably the ones about its origins and development in non-human animals. Moreover, a coherent, "continuous" picture would provide a strong, general scientific framework for the ethical debates. Veit doesn't stop at asking questions too; he offers a distinct view about how consciousness emerged, one I will now reconstruct and discuss.

2.2. Veit's story of animal consciousness

The main thesis of the book—the Pathological Complexity Thesis—claims that "the function of consciousness is to enable the agent to respond to pathological complexity" (p. 2). The notion of pathological consciousness is building on Godfrey-Smith's concept of "environmental complexity" (Godfrey-Smith, 1996), and let's keep in mind that Godfrey-Smith uses it to define cognition, as it will be relevant in the following sections. Veit's pathological complexity is essentially the measure of the Darwinian trade-off problem animals face in their struggle to increase fitness.

> The more degrees of freedom there are in the behavioural option-space (to put it simply: how many alterative actions and organism can take), the higher the pathological complexity of this fitness-maximization problem, since organisms have to make sure to make the right decisions at the right time, and this depends on their current state as well as that of the environment (p. 21).

The notion of pathological complexity is designed to serve as a functionalist bridge over the gap between externalist and internalist approaches in animal and human consciousness research, the former focusing on representations and sensory experience, and the latter on self-awareness. This is because, as a measure of an organism's strategy, it varies depending on the different life-histories and lifestyles. "Pathological complexity" is practically synonymous to "teleonomic complexity" and "life-history complexity"—Veit chooses to use the term "pathological" to evoke the concept of "health". In order to properly study consciousness, we need to distinguish healthy and pathological traits variations in its evolutionary history (just like ethologists need to distinguish healthy and pathological behaviours in what they observe). Veit proposes also to link those traits to specific incarnations of "phenomenological complexity".

Phenomenological complexity is the complexity of consciousness' forms or dimensions. Veit's creates a modified version of Birch et al.'s (2020) dimensional model of consciousness which included perceptual richness (of vision and of touch), evaluative richness, temporality, selfhood, and unity. Veit's simpler model only comprises five dimensions, which are: diachronic and synchronic unity of experience, sensory experience, evaluative experience, and experience of self. This is Veit's crucial step, because the heart of his book is the quest to find which of the five is the "heart" or "origin" of consciousness. Veit dubs

this the "war of five dimensions" and declares he can determine the winner by means of reverse-engineering. The author also brings up De Waal's idea of self-awareness as an onion whose heart could be treated as the "most minimal" form of consciousness (De Waal, 2019). The reverse engineering process essentially comes down to peeling the onion layer by layer until we get to its heart.

Veit begins with an easy rejection of both diachronic and synchronic unity of experience as plausible candidates. His arguments are a mixture of empirical cases and common sense, the gist being that we can easily conceive of conscious experiences without either diachronic or synchronic unity and/or find examples of individuals with split brains who are not devoid of other forms of consciousness. Notably, although Veit does not really engage in this book with the plethora of theories of consciousness on the market, in this part he specifically reviews the Integrated Information Theory (pp. 47-48), because it posits that synchronic unity is a necessary and defining feature of consciousness. He rejects it as a "competitor to the pathological complexity thesis" not based on its own intrinsic weaknesses, but chiefly because it doesn't pass the Darwinian bottleneck: it does not explain how the unity of experiences helps organisms survive.

Veit proceeds to reject the "self" dimension based on a critical discussion of a broadly construed "autopoietic approach", which he dubs another competitor to the pathological complexity thesis at "multiple levels of interest" (p. 53). This is because the autopoietic view, internalist, focused on self-organization, self-production, and autonomy, opposes an externalist, Darwinian/functionalist account of life, as is openly declared by Thompson (2010). The disagreement goes as deep as the basic idea of the chief goal of an organism. Veit simply rejects the autopoietic approach as "mistaken" (p. 54), reiterating the Darwinian thesis that the chief goal is reproduction—and not autonomy or self-organization. He does not undertake to defend Darwinism as such, although he does point to certain misinterpretations of Darwinism by the proponents of autopoiesis. Veit only briefly outlines his reasons for adopting the Darwinian stance on the main goal of living beings, which relies on the observation that organisms often choose maximizing the number of viable offspring over self-preservation.

However, there is also some agreement between Veit's pathological complexity thesis and the autopoietic approach. Both share similar deeper assumptions about the wide-spread character of consciousness, with a view to find its simplest, most basic forms, and a version of a strong life-mind continuity thesis. Veit acknowledges that the autopoietic approach makes room even for the most basic forms of bodily self-awareness, but points to two difficulties. First, he shares Godfrey-Smith's (2020) fear that the autopoietic approach paints consciousness as "an automatic feature of just being a living organism located in the world" (p. 55). This is mainly because it does not define any important func-

tion of consciousness and, a fortiori, does not explain it's adaptative value. Consciousness comes "for free" and, according to Veit, given how the autopoietic approach "neglects" the role of environmental feedback, it is difficult to see how it can gradually develop. Veit's other point against the autopoietic approach is that it asserts self-awareness as the most basic dimension of consciousness, one that constitutes the organism as an agent.

Veit ultimately rejects the possibility that selfhood could be the most basic dimension of consciousness in the same way he dealt with the temporal unity dimension: with a mix of speculation, empirical examples, and common sense.

He uses rhetoric rather than hard arguments here; for example, when criticizing De Waal's gradualist view of self-awareness, he states that it is "surely too strong" a claim to ascribe a minimal form of a self-concept to corals, anemones, and sponges (p. 57). When he recounts Feinberg and Mallat's (2016) distinction between interoceptive, exteroceptive, and affective sides of experiences, he surmises that "it does not appear hard to imagine that it [self-consciousness] simply constitutes a combination of these capacities" (pp. 57-58). Both arguments are presented rather as an appeal to common sense than an elaborated point, so they are only convincing if we share all the intuitions and assumptions behind them. I can easily agree that an anemone doesn't have a "self-concept"— but not necessarily, that it is an argument against self-consciousness being the key dimension. Although a "self-concept" might seem to be something cognitively advanced, too advanced for a "plant-like" anemone, a basic sense of self, a single, perhaps just a defined point of view which can have nothing to do with possessing any concepts or mental capacities, does not.

True, Veit does engage critically to some extent with the autopoietic tradition, the source of the idea of a basic self (and self-maintenance), but he dismisses the tradition as a whole, on the basis of metatheoretical objections. He doesn't really point out any specific reasons why the idea of a very basic self being the basis of any dimension of consciousness should be rejected. Veit claims that if we start from "a feeling of self", we could not explain how the evaluative dimension of consciousness could evolve from it, but this is, again, rather a metatheoretical intuition than a fact. To judge whether this would truly be problematic, we would have to have a much clearer notion of what counts as "explaining" the evolution of certain dimension. If we stick to intuitions, it stands to reason that nothing could ever be evaluated if there is no coherent point of view with respect to which it could be evaluated. Why shouldn't we, therefore, try harder to explain the evaluative dimension as secondary to the "self-consciousness" one? Especially that some theories, both phenomenological and empirical, suggest that a basic self-awareness is both a constitutive element of consciousness and an evolutionary milestone.

Veit is, of course, aware of this. Still, he opposes Godfrey-Smith's (2020) claim that the ability to distinguish self from the rest of the world was a necessary (and very early) step in the evolution of complex bodies. The argument is that this would not amount to any kind of self-experience, which is probably a fair point. We must always be careful not to confuse experiential forms of self-consciousness with our theoretical constructs concerning pre-reflexive self-awareness and concepts of self as the simplest possible coherent points of view, the latter being something more akin to Kant's notion of the transcendental unity of apperception than to anything empirically observable. However, Veit assumes that in order to prove that the self-awareness dimension of consciousness is the "dimension of origin", we would be committed to show that its most basic form is a "feeling". This assumption is unfounded, especially given that his dimension of choice, the evaluative dimension, also does not emerge at its very beginning as a full-blown experience. Moreover, we have no clarity as to the relations between the phenomenological and autopoietic ideas concerning forms of basic self-awareness, which contributes to the overall haziness of this dimension and complicates any attempts to assess Veit's arguments.

Here comes the twist: although the previous paragraphs might have sounded as if I were pushing self-consciousness as consciousness's "dimension of origin", this has never been my goal. I'm not sure such claim could be supported in a more decisive way than Veit's is. And this is exactly what I wanted to show. Veit promises us a "war" with a clear "winner" (and even "spoils"), but what he delivers is more like trench warfare. The dimensions are clenched in a bloody stalemate, with troops claiming and reclaiming the same scrape of land over and over again, and only our patriotic propaganda keeps convincing us that our side has made significant advances. The tired soldiers often doubt if there is any sense in all this fighting.

This becomes even more apparent when Veit tackles the last candidate opposing his champion: sensory experience. Veit begins, rather strikingly, with an argument grounded in experimental studies by Sytsma and Machery who investigated the folk notion of consciousness (2010). Veit paints a perhaps overly simplified picture of both the results of those studies and the way they support his claims that evaluation is more of a "core" feature of consciousness than phenomenal experiences. He is a bit hasty to announce decisively that laypeople's judgments are a "radical inversion" of the way philosophers think about the relation between evaluative and perceptual states. A more cautious interpretation would be that laypeople differentiate between evaluative/affective and perceptual states and are more reluctant to ascribe the former to non-human beings, such as robots. Please note that, yet again, the same point could lead to a directly opposite conclusion: maybe laypeople perceive perceptual consciousness as something "simpler" to achieve than evaluative states, so it must be also more basic? Regardless of our interpretation, Veit is aware

that an appeal to folk-psychology can only inspire, and never settle a philosophical discussion. Still, I think he is right in bringing up reasons to be wary of classic philosophical assumptions and the view of phenomenal experience embedded in mainstream philosophy of mind certainly requires more critical consideration.

Veit can effectively show how that deeply embedded representationalist and externalist ideas about perception are still permeating our view of consciousness without sufficient justification. The crux of argument here is in line with the critique offered by the "internalist" autopoietic approach. Phenomenal consciousness is not a stream of purely qualitative, internal experiences which are independent of the agent's situation, wants or needs, in short: devoid of intrinsic hedonic valence. The starting point of classic, philosophical and rather abstract approach to perception is how humans create rich, detailed representations of the world, which is very different from the reality of simple organisms fighting for survival and reproduction. The Darwinian perspective here is refreshing: we need to explain what purpose would be served by such computationally heavy, costly representations—and the proponents of embodied cognition often argue that this question has no satisfactory answer. It stands to reason that the need for perceptual experience is not the most basic need of any organism. A living being needs, first and foremost, to guide its actions, understood here as moving through option space, without implying any intentionality.

Although Veit's arguments for eliminating the other dimensions are sometimes wonky, he has a strong positive case for the evaluative dimension. The Darwinian, action-oriented perspective suggests that the primary function of consciousness must be to aid organisms in making choices. In order to act, an organism must somehow judge what is good and what is bad for its continued existence. It needs valence.

> Origins of valence are better conceived as the origins of fuzzy action imperatives, which arose out of something like the vague discomfort Romanes describes, but which then evolved to have richer discriminatory capacities between different states (p. 64).

It is a coherent and compelling vision of the evolutionary path of consciousness: it starts from "fuzzy imperatives" and only then develops increasingly detailed, sensory representations, gradually involving more and more self-awareness, temporal unity, and mental time travels depending on the needs and demands arising during the organisms' diverse life histories. This theory certainly passes the "Darwinian bottleneck", but it also has two additional virtues I would like to discuss.

First, we get a good glimpse of how evaluative states gradually become something much more advanced than "fuzzy action imperatives". Relying on Browning (2020) and the groundwork laid by McCleery, 1977, McFarland & Sibly, 1975, and McNamara & Houston, 1986, Veit points out that our concept of how animals choose different actions based on various trade-offs, desires, wants, needs, and fears, requires us to explain how organisms compute all those different incentives and deterrents. They need a "common currency of evaluation" something that according to Browning and Veit can be supplied by hedonic valence. The means of evaluation have to evolve as the organisms get more and more complex and their landscape of affordances grows richer and bigger, driving the development of all the other dimensions of consciousness. However, even in the later stages of this process, when organisms boast complex capabilities in all five dimensions of consciousness, the evaluative side clearly remains at their core, still serving as the common currency despite the increase in types of money in circulation. Importantly, we should not confuse the later complex role of being the common currency with the humble beginnings of this dimension of consciousness, identifiable even in the simplest animals, which is a point contention between Veit and Cabanac (pp. 75-76).

The second virtue of this account is that consciousness doesn't come "for free", it isn't just a byproduct of life. We can trace the evolutionary continuity of consciousness, but this path also has a clear beginning—a threshold in pathological complexity that triggers the need for evaluative powers.

> An explosion in pathological complexity through higher degrees of freedom comes to be dealt with through a major evolutionary transition towards a hedonic mode of evaluative agency ie what I call Benthamite creatures. (p. 69)

We can even locate this explosion in pathological complexity in time—not surprisingly and in line with other researchers such as Godfrey-Smith, Veit points his finger at the Cambrian explosion ca. 541 million years BC. According to the author, this was actually a second explosion in pathological complexity, the first one, the Avalon explosion, having ended in failure. The creatures of the White Sea couldn't solve the action selection problem, Veit claims—for that, you need to develop hedonic valence. Hedonic valence is the first step to reinforcement learning which can be understood as what ties the developing ability to subjectively evaluate to the objective, environmentally determined, sometimes hidden fitness of a given behaviour. At this stage of the book, it seems as if there was only one hurdle left on the way to a science of animal consciousness—this hurdle, however, is much higher than it seems at first glance.

3. How Veit's argumentation undermines the status of his claims

The hurdle I'm referring to is Peter Godfrey-Smith's thesis that the sensory and evaluative varieties of experience are not only separate, but thoroughly independent, and may have also evolved independently. Veit is, of course, happy to embrace the independence of the evaluative dimension, this is in line with his account. He discusses in detail Godfrey-Smith's example: sea slugs, which purportedly have evaluative experiences despite very poor sensory capabilities. The other half of Godfrey-Smith's thesis, that sensory experience also evolved independently of evaluative experience, if true, would destroy the very heart of Veit's theory. Veit, therefore, feels forced to deal with the empirical example provided by Godfrey-Smith: arthropods. Godfrey-Smith recounts experiments proving that insects possess formidable sensory capacities but, in contrast with slugs, seem to be indifferent even to grave bodily harm. It has often been observed that many insects ignore even a loss of a limb. Veit attacks this interpretation of the experiments from almost every angle, pointing to such facts as ambiguous results (some insects do seem to tend their wounds, it has also been shown that bees avoid heat) and methodological and ecological conundrums (insects, contrary to other animals, can't grow back any lost limbs, so it makes no sense to care for them, we are not sure if insects even can react to morphine). Veit even presents a way of refitting Godfrey-Smith's view of the experiments to ultimately match his theory. If we still stand by the interpretation that arthropods enjoy sensory experience without evaluative capacities, we could claim that they had such capacities, but lost them in the process of evolution, and now rely on a different dimension of consciousness. In this, we are supported by what happens on the ontogenetic level: the larvae seem to possess less sensory abilities but behave more as if they had some evaluative sense. It seems that Veit has every angle covered… but his victory is Pyrrhic.

The multi-level attack from all flanks definitely succeeds in shedding doubt on Godfrey-Smith's interpretation of arthropod research. The way Veit achieves this, though, points to a much deeper problem his own theory might be facing as well. It is, simply speaking, too easy. To begin with the most obvious: adding the caveat that certain organisms seem to be relying solely on the non-evaluative dimensions of consciousness not because they don't have the evaluative dimension but because they "lost" it would make Veit's theory practically infallible. It is difficult to imagine any potentially falsificatory empirical finding that could not be explained away in an analogous fashion. I'm not suggesting that such a general account of animal consciousness should be strictly empirically falsifiable in a straightforward way. It is always a sign of weakness, however, if a theory allows for "too much", and it significantly limits its explanatory power. A much better strategy would be for Veit to dig his heels in and unequivocally reject Godfrey-Smith's treatment of empirical facts. Unfortunately, though, despite presenting so many counterarguments, he couldn't ever have achieved that.

We simply can't really tell whose interpretation of insect behaviour is correct. We have no criteria of judging with any level of conviction what suffices as proof of possessing evaluative consciousness or of not possessing this dimension. Having said as much, I need to add that Veit's intuitions about the general significance of tending wounds in animals built like insects strike me as valid. In general, the debate about experiences of negative valence in invertebrates seems to be shifting towards an approach that is perhaps less cautious in terms of Morgan's Canon, but more cautious from an ethical perspective. We shouldn't be too hasty in assuming that animals don't feel pain just because they don't exhibit enough behavioural similarities to us. But this is a good ethical and political strategy, not one for a "science" of consciousness. It seems that the category of "evaluative experience" can't be applied to the simplest cases with any certainty, the concept is too broad. The observations, on the other hand, are simply too ambiguous, because too many empirical, methodological and theoretical links are still lacking. We can't be sure what we are observing because we can't, at this stage, build detailed enough hypotheses.

Ironically, Veit's meticulous analysis of the case of slugs, which was supposed to aid his cause, also reveals this same problem. He reminds us not to identify negative valence with pain (this would be anthropomorphizing), which suggests that we should not think of slugs' experience as something phenomenally conscious. At the same time, he quotes Godfrey-Smith's striking idea to imagine slugs as bigger and faster, so that experience in these animals would become "almost inescapable", and brings up the disability paradox to underline our inability to assess the subjective experience of others. Moreover, he admits that certain slugs' behaviour suggests that they also have certain awareness of self and diachronic unity, although, of course, those are "not necessary". Could we truly ever distinguish on the basis of coarse-grained, ambiguous observations of slugs whether they exhibit an awareness of self, experience of pain, just simple negative valence, diachronic unity, and which of those dimensions of consciousness, if any, are more basic than other? My point is, that without much stronger, possibly mechanistic theories behind all those terms, and clear operationalizations, those are not valid empirical questions.

The problem with Veit's "war" is, therefore, not that his assaults are too weak or that any of the opponents' defenses are too strong. It is rather that there is no clear battlefield at all. Given how Veit frames the debate about the origins of consciousness, we can think of two ways of settling it: the empirical route and the conceptual route. The key problem is that we are stuck in a vicious circle here: we don't have clear enough empirical evidence to fuel our conceptual musings, and it might be impossible to get such evidence because our concepts are too hazy. Worse still, this is probably not our fault—just an objective difficulty of the subject matter. If our goal is to investigate the very beginnings of consciousness, it is only to be expected that our empirical evidence will be as murky as the primordial soup. The dimensions are inevitably intertwined

and difficult to disentangle precisely because we aim to catch them at a stage when they are just being born and perhaps begin to differentiate.

Unfortunately, this line of criticism should cut even deeper: at the level of simple animals and minimal cognition, it is not only dimensions of consciousness that become difficult to distinguish in any kind of empirical research. As I mentioned in a previous section of this paper, Godfrey Smith used an analogous definition of complexity to define cognition, not consciousness. Cognition in its minimal forms often comes down to reinforcement learning. The ability to move through the space of options so as to increase fitness is how many would define intelligence. Some philosophers and other animal researchers find it suspiciously easy to describe all and any empirical findings, often employing such terms loosely and freely. This is also something Veit is aware of, especially given that he also intends to shed light on plant sentience (Terrill & Veit, 2024), which is a field riddled with conceptual difficulties and worries about overextending many cognitive concepts (Solé, 2023; Ten Cate, 2023, but see also Białek, 2024). The phenomena we are empirically investigating in plants and simple animals are just too far removed from both folk-psychological and traditional, philosophical ways of understanding such concepts as sentience, consciousness or cognition to serve as strong bases for empirical arguments. Empirical observations can and must fuel the theoretical debates, of course, but we can't hope to build a strong theoretical framework from the bottom-up. Slugs and insects can't show us whether they possess sensory or evaluative consciousness if we don't have first a very good idea about how to operationalize those notions and a strong naturalistic, mechanistic theory behind them. We are simply not sure what we are observing—and, crucially, what we are not observing.

It is tempting, therefore, to switch entirely to the conceptual mode and treat such proposals as Veit's as proposals of theoretical frameworks, an attempt to reconstruct the logical relations between certain notions pertaining to consciousness. This is what the title—"a philosophy for the science of animal consciousness"—suggests. The claim about the evaluative dimension being the first to evolve would now mean that it is the logical essence of our concept of consciousness and that all the other dimensions are built on it. The status of the theory would be such that it is a conceptual framework to be used and filled in by naturalistic accounts, and that only then it could one day generate predictions that would be precise enough to be tested empirically. However, it is not ready to do so yet. It fits well with the paleontological and experimental findings, but we can't really say it explains them at this stage. It certainly deepens our understanding of the evolutionary story of consciousness and proposes a novel way of ordering our data. Essentially, it's a philosophical analysis of the structure of the concept of "consciousness" fueled by the theory of evolution. What it has to offer, for example, is the neat (though sadly illegible in its

printed form—see p. 111) diagram representing the profiles of conscious experiences in different animals using the five-dimensions model. It is very helpful in ordering our thoughts about the empirical research. It is not predicted by the theory, though—just gerrymandered to fit the data, as they are interpreted. This is a new, attractive way of saying things we already know, and a useful tool for visualizing new empirical discoveries.

I have to say, though, that if we view Veit's account through this lens, some of his arguments don't really match the status of his claims. Especially the elimination process in the "war of the five dimensions" relies, as I have pointed out, on empirically grounded intuitions, not on decisive, conceptual arguments. For example, from a conceptual point of view, without relying on commonsensical assumptions about anemones, it is particularly difficult to defend the claim that the evaluative dimension precedes the self-consciousness dimension. And although this might sound as a damming criticism of Veit on my part, I actually think it is a cost he had to pay for avoiding a much more dangerous trap. Had he decided to build a purely conceptual construct (albeit grounded in a healthy dose of Darwinian evolutionism) with a view to use it in animal studies, relying on conceptual arguments, he would inevitably succumb to what I have defined as a case of bad "Kantian anthropomorphism" (Białek 2023, 2024).

4. Veit's theory as a way for dealing with Kantian anthropomorphism

Kantian—or transcendental—anthropomorphism is not the regular kind that is so often debated in animal studies, the practice of using certain concepts labeled as "human" to describe non-human behaviour. Transcendental anthropomorphism is metatheoretical. It determines how we build those concepts themselves, what we tend to naturally associate, what we consider a satisfactory account of certain phenomenon. Unavoidably, we do this from a human perspective, we apply the same metatheoretical principles that have made sense in human science studying humans. For example, in the case of emotion research, despite all the conceptual chaos, we have a general idea of all the components that should be involved in a satisfactory theory: certain physiological reactions, subjective experience, behavioural signs. We have even been trying (so far, without success) to build a model of how basic human emotions (whatever those are…) are expressed physiologically in different species. The very idea that such are the components of emotions, to say nothing of assuming the existence of concrete types of emotions across species, is anthropomorphic. It simply doesn't have to be that way, especially if we consider the long, branched story of evolution.

Veit is clearly aware of this pitfall. He points out that identifying negative valence with pain is anthropomorphic, a claim he reiterates in (2024). True, it is anthropomorphic, in exactly the way I have described. Veit astutely criticizes Ginsburg and Jablonka (Ginsburg et al., 2019) for defining the hallmarks of

minimal conscious experience in a human-centric manner (p. 78). This is important because Ginsburg and Jablonka took pains to avoid anthropomorphism and they explicitly made sure that their theory does make sure human-specific abilities such as language necessary for consciousness. They did not avoid transcendental anthropomorphism, though.

Of course, as I have also argued before, transcendental anthropomorphism can't be avoided—but it can be tamed, and Veit seems to be wonderfully apt at this. It seems that the author really took to heart Daniel Dennett's (1991) warning against "Mistaking a failure of imagination for an insight into necessity" which he quotes on p. 25. Taming transcendental anthropomorphism does require boosting our philosophical imagination, just like moving beyond our transcendentally ingrained Newtonian structures of space and time did. First and foremost, however, it requires making those structures that seem so necessary explicit; only then we can imagine what would happen if we let them go.

This is why I value so highly Veit's attempt at constructing a model of consciousness which is truly malleable and ready to be fitted to whatever we learn about the evolutionary history. Too often in comparative cognitive science do we take the stiff forms of human cognitive capacities and try to mold other animals into them. The crucial step is, I believe, the idea shared also by Birch and Godfrey-Smith, that there exist independent dimensions to conscious experience which can occur and evolve independently. This provides us with the necessary elasticity. Veit's and Birch's graphs of consciousness profiles that are so different from ours may feel gerrymandered, but they are also extremely valuable because of the way they open our minds to a widely different understanding of how consciousness my present in other creatures. Godfrey-Smith's arguments about slugs and insects may be vague, but the thought that sensory and evaluative dimensions are quite independent (even if we are not sure about empirical proofs of full double dissociation) is a thought well worth investigating. Paradoxically, many of the apparent weaknesses of Veit's theory and argumentation that I have been pointing out here may turn out to work to our advantage after all. We don't want our philosophy for the science of animal consciousness to restrict us – on the contrary, we want it to free us of from preconceived ideas and unfounded assumptions.

5. Conclusions

It seems that Veit's philosophy for the science of animal consciousness may lead us in unexpected directions, especially if we dig deeper into how it is construed and justified. It is a great conceptual framework firmly grounded in sound metatheoretical choices. Veit has a knack for asking tantalizing questions – especially those concerning the origin story of different dimensions of consciousness. Surprisingly though, if we look closer at both the conceptual

and empirical arguments backing Veit's answers, it turns out that not only are they not as convincing as we would like, but also that they couldn't possibly be better. This in itself is an informative takeaway message from studying this book. There is more, however.

As I argue in the latter parts of this paper, Veit's framework itself may serve a fantastic purpose: getting rid of many of the preconceived ideas that still permeate both philosophical and psychological investigations into non-human consciousness. It is malleable and flexible, allowing us to adopt fresh perspectives.

Of course, there is also a drawback here. After we let go of so many assumptions and ideas, we are inevitably left with a very broad and vague notion. Moreover, as Darwinians, we know that not only does consciousness come in vastly different profiles, but it also evolves via so many separate paths. Veit openly likens his idea of consciousness to movement—it is a very general term describing a huge variety of behaviours. Although he presents it as a virtue of his theory, it can also be worrying. The concept of movement, if we truly take into account all its forms from flying, swimming, and running to bacterial phototaxis, isn't really interesting, not from the point of view of animal studies. We don't really have a science of animal movement in general; it would have to be too absurdly diversified. Not many interesting things can be said about the evolution of all and any movement on planet Earth. We do study certain animal movements from a kinematic point of view, but this would be analogous to studying specific consciousness profiles from a particular perspective. An optimistic view would be, therefore, that in the future we will simply do just that: define particular profiles of consciousness, as different from one another as flying is different from swimming, and study their evolutionary paths and any other interesting properties.

This doesn't bode well for integrating the science of animal consciousness with anything we have learned about human consciousness so far, though. I think it is not a coincidence that, barring the IIT, Veit doesn't really touch upon any of the existing neuropsychological and philosophical accounts aimed at explaining human consciousness. On his view, our consciousness is like flying. To understand how flying works, we need a detailed study of how our wings and feathers are built, but all this has little to do with how seals move through water. And this is why, after all this effort and despite all the metatheoretical assumptions of continuity, we could end up with many sciences of consciousness instead of one. But perhaps this is fine. Perhaps we simply don't need any unified science of animal consciousness. And if a philosophy of it can tell us this much, then, I conclude, it is extremely intriguing.

References

Białek, M. (2023). The New Anthropomorphism Debate and Researching Non-Human Animal Emotions: A Kantian Approach. Annals of Philosophy, 71(3), 205–229. https://doi.org/10.18290/rf23713.10

Białek, M. (2024). Extending the Concept of Cognition and Meta-Theoretical Anthropomorphism. JoLMA, 2, JournalArticle_15404. https://doi.org/10.30687/Jolma/2723-9640/2023/02/007

Birch, J., Schnell, A. K., & Clayton, N. S. (2020). Dimensions of Animal Consciousness. Trends in Cognitive Sciences, 24(10), 789–801. https://doi.org/10.1016/j.tics.2020.07.007

Browning, H. (2020). If I Could Talk to the Animals: Measuring Subjective Animal Welfare. manuscript.

De Waal, F. B. M. (2019). Fish, mirrors, and a gradualist perspective on self-awareness. PLOS Biology, 17(2), e3000112. https://doi.org/10.1371/journal.pbio.3000112

Dennett, D. C. (1991). Consciousness explained (1. paperback ed). Little, Brown.

Feinberg, T. E., & Mallatt, J. M. (2016). The Ancient Origins of Consciousness: How the Brain Created Experience. The MIT Press. https://doi.org/10.7551/mitpress/10714.001.0001

Ginsburg, S., Jablonka, E., & Zeligowski, A. (2019). The evolution of the sensitive soul: Learning and the origins of consciousness. The MIT press.

Godfrey-Smith, P. (1996). Complexity and the Function of Mind in Nature (1st ed.). Cambridge University Press. https://doi.org/10.1017/CBO9781139172714

Godfrey-Smith, P. (2020). Metazoa: Animal life and the birth of the mind (First edition). Farrar, Straus and Giroux.

McCleery, R. H. (1977). On satiation curves. Animal Behaviour, 25, 1005–1015. https://doi.org/10.1016/0003-3472(77)90052-5

McFarland, D., & Sibly. (1975). The behavioural final common path. Philosophical Transactions of the Royal Society B: Biological Sciences, 270(5), 265–293.

McNamara, J., & Houston, A. (1986). The Common Currency for Behavioral Decisions. The American Naturalist, 127(3), 358–378.

Solé, R. V. (2023). Do plants have the cognitive complexity for sentience? Animal Sentience, 8(33). https://doi.org/10.51291/2377-7478.1810

Sytsma, J., & Machery, E. (2010). Two conceptions of subjective experience. Philosophical Studies, 151(2), 299–327. https://doi.org/10.1007/s11098-009-9439-x

Ten Cate, C. (2023). Plant sentience: A hypothesis based on shaky premises. Animal Sentience, 8(33). https://doi.org/10.51291/2377-7478.1795

Terrill, E. C., & Veit, W. (2024). Better to be a Pig Dissatisfied than a Plant Satisfied. Journal of Agricultural and Environmental Ethics, 37(1), 4. https://doi.org/10.1007/s10806-024-09922-4

Thompson, E. (2010). Mind in life: Biology, phenomenology, and the sciences of mind (First Harvard University Press paperback edition). The Belknap Press of Harvard University Press.

Maja Kittel (f. Białek) works at the University of Białystok's Faculty of Philosophy. She is a philosopher of mind focusing on two topics: the philosophy of psychiatry and animal minds. Her published work primarily concerns the nature of delusions and their epistemic status, the concept of cognition, and anthropomorphism in Animal Studies. She is also a member of the Academic Network for Environmental Humanities.