This paper is published in TOPOI (2015). The final publication is available at Springer via
http://dx.doi.org/10.1007/s11245-015-9338-z

Sensorimotor theory and enactivism

Jan Degenaar [1,2] & J. Kevin O'Regan [1]

[1] Laboratoire Psychologie de la Perception, Université Paris Descartes, France.
28 Rue des Saints-Pères, 750006 Paris, France.
[2] Department of Philosophy, University of Antwerp, Belgium.
Email: degenaar.jan@gmail.com; jkevin.oregan@gmail.com.

Abstract. The sensorimotor theory of perceptual consciousness offers a form of enactivism in that it stresses patterns of interaction instead of any alleged internal representations of the environment. But how does it relate to forms of enactivism stressing the continuity between life and mind (and more particularly autopoiesis, autonomy, and valence)? We shall distinguish *sensorimotor enactivism*, which stresses perceptual capacities themselves, from *autopoietic enactivism*, which claims an essential connection between experience and autopoietic processes or associated background capacities. We show how autopoiesis, autonomous agency, and affective dimensions of experience may fit into sensorimotor enactivism, and we identify differences between this interpretation and autopoietic enactivism. By taking artificial consciousness as a case in point, we further sharpen the distinction between sensorimotor enactivism and autopoietic enactivism. We argue that sensorimotor enactivism forms a strong default position for an enactive account of perceptual consciousness.

Keywords. Consciousness; sensorimotor theory; enactivism; autopoiesis; artificial consciousness

1. Sensorimotor theory

Enactive approaches to consciousness emphasize patterns of interaction, and active engagement, instead of any alleged internal representations of the environment. But there are various forms of enactivism, raising questions about their interrelations. While some enactive accounts may be complementary and differ mainly in emphasis, others may manifest outright conflicting conceptions of consciousness. Indeed, in this paper we identify a deep divide between two classes of enactive theories. We are particularly concerned with the sensorimotor theory, and how it compares to other forms of enactivism.

Sensorimotor theory makes a constitutive claim: perceptual consciousness is constituted by the exercise of sensorimotor capacities (O'Regan & Noë 2001; O'Regan, 2011; Degenaar & O'Regan 2015).[1] Moreover the sensorimotor theory claims that the particular *quality* of conscious experience, e.g. what makes an experience the experience of red, lies in particular patterns of sensorimotor engagement. The theory thereby offers an enactive account of consciousness, rejecting accounts in terms of inner

[1] Perceptual consciousness refers to a subset of conscious experiences, including our sensory experience of the environment by vision, audition, touch, etc., and excluding experiences of thought, emotion, and feelings like hunger and pain.

representational models.[2] In this paper we discuss how sensorimotor theory relates to forms of enactivism stressing the continuity between life and mind, and more particularly autopoiesis, autonomy, and valence (e.g. Thompson 2007). We shall distinguish *sensorimotor enactivism*, which exclusively stresses perceptual capacities themselves, from *autopoietic enactivism*, which appeals to a broader range of processes or capacities by claiming a *necessary* connection between experience and autopoietic processes or associated background capacities. We shall discuss whether, given a focus on capacities, there are reasons to embrace autopoietic enactivism. But first we shall briefly summarize the sensorimotor theory.

1.1. Perceptual capacities

Perceptual consciousness, sensorimotor theory claims, lies in the exercise of perceptual capacities, not in any internal representation or dualistic 'je ne sais quoi' that may be thought to accompany them. Thus perceptual capacities can be compared to behavioral capacities such as the capacity to cycle. To be able to cycle means among other things that you can keep your balance on a bike despite irregularities in the road and the wind, that you can avoid obstacles and take a curve. The cycling is enacted, it unfolds in dynamic engagement with the environment, and it is unhelpful to describe cycling as putting to use any kind of internal representation.

In order to define perceptual capacities, we first define *sensing capacities*, which are to be understood as capacities for being attuned to aspects of the environment, that is as having mastery of the current sensorimotor dependencies linking possible actions and resulting changes in sensory stimulation. We use 'sensing capacities' here to refer to capacities that can be displayed even by simple artifacts, such as missile guidance systems, which can be said to 'sense' their target, albeit of course not in a way that implies perceptual experience. Sensing capacities are objective properties of a system that are expressed in the system's behavior: the fact that the missile guidance system is attuned to its target is shown by the fact that the system responds appropriately to movements of the target. (Note that while sensorimotor dependencies are defined proximally, in terms of bodily movement and sensory stimulation, attunement to sensorimotor dependencies is attunement to the world. Consider for example that a missile guidance system could make use of a continuously moving camera, so that movement of the target implies a change in the obtaining sensorimotor dynamics: clearly the systematic dependency of the sensorimotor dynamics on the environment forms the basis for the system's attunement to the environmental target.)

Perceptual capacities are the sensing capacities of perceivers: Through the use of sensory systems, perceivers can potentially make use of their sensing capacities in various environmental situations. In contrast to simple systems such as the missile guidance system, perceivers have the capacity to make use of sensing capacities in a broad range of situations, and to express their sensing capacities in a broad range of behaviors. (Note that this is not intended as a full-blown definition of perceivers, but merely as a definition of perceptual capacities – possible constraints on what counts as a perceiver will be discussed in the sections below. The point here is merely to define perceptual capacities without reference to consciousness, providing common ground for subsequent debate on perceptual

[2] As an enactive account, sensorimotor theory also contrasts with computational functionalism as familiar from the work of Clark and others (e.g. Clark 2006). An enactive sensorimotor account emphasizes capacities rather than committing to a computational account based on information processing. For criticism of a cognitivist/representationalist conception of sensorimotor theory, see Hutto (2005).

consciousness.) While it will surely be hard or even impossible to draw a precise line between perceivers and other systems, there are clear cases of perceivers such as human beings. From the perspective of sensorimotor theory, the perceptual capacities of human beings are no less objective than the sensing capacities of a missile guidance system. Perceptual capacities are indeed just sensing capacities, but sensing capacities that are being *put to use* by a perceiver. A perceptual capacity common in humans is for example the capacity to see the color of an object; this capacity may for example be displayed in the behavioral capacity to point at the colored object or to name its color.

Sensorimotor theory claims that perceptual capacities depend on the 'implicit grasp' of sensorimotor dependencies (e.g. O'Regan & Noë 2001; O'Regan, 2011; Degenaar & O'Regan 2015). This means that perceivers need not have personal level access to the way in which sensory stimulation depends on movement. (Similarly, when an agent has the skill to ride a bike, the agent does not thereby have personal level access to the details of the sensorimotor patterns he or she is thereby mastering.) Sensorimotor theory claims that the person's implicit attunement to the environment depends on attunement to sensorimotor dependencies.

The sensorimotor dependencies underlying sensing and perceptual capacities can be defined at various levels of abstraction. For example, some patterns are already present at short timescales, and concern immediate sensory consequences of movement (e.g. in the case of visual stability). Other patterns of sensorimotor dependency are only actualized in the extended exploration of objects (e.g. in case of our interacting with 3D objects in space) (Maye & Engel 2012). It is often far from evident which patterns we engage with, as can be seen for example in Philipona and O'Regan's (2006) analysis of some patterns relevant to color vision. An apparently 'simple' experience such as the color of an object may involve attunement to complex patterns of the way in which stimulation of photosensitive cells in our retina changes when the illumination changes, or when we sample differently-illuminated parts of an object. Similarly, while it is far from immediately evident how bodily activity changes olfactory stimulation, evidence suggests that olfaction depends on complex smell-specific sensorimotor profiles (see discussion in Cooke & Myin 2011). Note that the various levels of abstraction or complexity of sensorimotor patterns need not be thought of solely in terms of temporal scales: at one particular timescale relevant to a domain of experience, some patterns of sensorimotor dependencies may be relatively simple, while others are more complex (for example, for evidence suggesting that some color experiences depend on simpler sensorimotor profiles than others, see Philipona & O'Regan 2006). Furthermore, in some cases sensorimotor complexity may derive from differences in the complexity of relevant exploratory activity (as pointed out by Maye & Engel 2012), while in other cases it may derive from the fact that the same action has more variegated consequences within one perceptual modality than in another.

Note also that sensing and perceptual capacities develop and change gradually over time. To explain *which* sensorimotor dependencies we are attuned to, we must therefore consider the history of our sensorimotor interactions. Biases in action, as well as bodily and environmental constraints, create biases in the occurring sensorimotor dependencies, thereby affecting the sensorimotor dependencies that an agent, be it biological or artificial, has the opportunity to become attuned to. As a simple example, if you never were to move your eyes, you presumably would not have the opportunity to get attuned to the sensorimotor dependencies of eye-movements. Similarly, if you were to learn to estimate the speed of moving objects exclusively while tracking them with your eyes, you would presumably not get very skilled in estimating speed under different conditions. And given a tendency to

coordinate the movement of both eyes, you would have little occasion to fine-tune chameleon-style vision. Whenever there are systematic ways in which action depends on sensory stimulation (motor-sensory dependencies), or in which actions depend on other actions (motor-motor dependencies) this narrows the range of sensorimotor patterns forming the background of perceptual development.

1.2. Conscious experience

With the notions of sensing and perceptual capacities in mind, a further question can be asked about the conditions under which we can speak of *conscious perceptual experience*. Merely having a sensing capacity clearly doesn't imply consciousness (take for example the sensorimotor capacity of a missile guidance system). Further adding the requirement that the agent having the sensing capacity can be considered a *perceiver*, also does not guarantee that the experience should be conscious – after all, there are cases such as absent minded driving where a perceiver displays perceptual capacities without being fully conscious of the environment. A distinction can be made between conscious perception and unconscious perception (which is a way to say that also conscious perceivers have non-conscious perceptual capacities). What then are the criteria to guarantee that a perceiver demonstrating his sensing capacity should additionally be conscious? There may not be a clear line between conscious and not-conscious, so that consciousness 'comes in degrees'. Still it seems that full-blown consciousness additionally requires that the perceiver's sensitivity to the environment should potentially be useable in (rational) planning, thought, and in the case of typical human adults, verbal report (O'Regan & Noë 2001; O'Regan 2011). Perceptual *consciousness* involves additional constraints over and above perceptual *capacity*.

And here then lies the crux of the problem to be addressed in this paper: In humans, verbal report and the acting out of plans based on our experience *expresses* conscious experience. However, intuitively at least, adding a capacity for verbal report or planning to a machine doesn't suffice for us to be willing to ascribe consciousness to the machine, certainly not as long as the range of capacities remains severely restricted. It would seem that something else is required before we would speak of conscious experience; what could this be? Could it be simply *more* perceptual capacities and associated capacities to express these? Or do we need *non-perceptual* capacities such as the self-producing capacities of living organisms that (unlike the planning, thought, or speech referred to above) do not *express* the sensing capacities.

We may compare the case with the case of intelligent action. A machine having just one capacity cannot be said to act intelligently. It is only in the context of a broader range of capacities that the performance of an action can be more or less intelligent. In the case of intelligence, perhaps no more is needed than the flexible use of a broad set of capacities, so that no special 'intelligence-capacities' may have to be postulated to account for intelligence. In the case of perceptual consciousness, we may wonder whether some special 'consciousness-capacities', perhaps related to life, are required for sensing capacities to be capacities for conscious perceptual experience.

Note that 'consciousness-capacities' should not be understood in a dualist fashion. After all, a defining feature of the enactive approach is its emphasis on bodily engagement. A way to make the point is by contrast with the philosophical notion of a zombie, which is the idea that a physical duplicate of a conscious being could lack consciousness (Chalmers 1996). Enactive approaches reject this dualist notion: when a conscious being is duplicated in every structural, functional, and behavioral detail, it seems incomprehensible to deny that the duplicate is a conscious system, too (Thompson 2007: 231).

This rejection of dualism, however, does not in itself help to single out which capacities or bodily processes are the relevant ones for consciousness. It may be clear enough that if a conscious living being were to be duplicated, its perceptual consciousness would be duplicated. However, if a conscious living being is duplicated, except that the duplicate's toe nails do not grow, we still would say that perceptual consciousness is duplicated. Within enactive theorizing, the issue at stake is *which* bodily (physical/functional/behavioral) processes would have to be duplicated for an experiential duplicate. In other words, the issue concerns which aspects of bodily engagement are the crucial factors for perceptual consciousness.

A key issue within enactive theorizing concerns the possible role of life in a theory of consciousness. In the next section we shall introduce ideas about life and autopoiesis, and two kinds of enactivism making different claims concerning which are the relevant capacities for speaking of conscious experience.

2. Life and mind: autopoiesis in two kinds of enactivism

Within a broadly enactive approach, our conscious experience is considered to be an aspect of our lives as biological organisms (e.g. Jonas 1966; Varela, Thompson & Rosh 1991; Thompson 2007). Without pre-judging the possibility of artificial life (or the artificial reproduction of relevant features of living organisms), the idea of the enactive approach is that to understand human mental phenomena, is to understand the way our living bodies engage with the environment. Here we shall introduce autopoiesis, and consider the role it may have in two kinds of enactivism.

2.1. Autopoiesis

Perhaps the most basic motivation for turning to life for understanding consciousness is that structural features of life match with structural features of consciousness, leading to the idea that consciousness is a particular form of life, and that there is a deep continuity between life and phenomenal experience (Thompson 2007). Just as life is a form of organization in which the organism deals with its environment, consciousness is a form of organization in which the system deals with its surroundings. The idea then is not only that a philosophy of life comprise a philosophy of mind (Jonas 1966), one may suspect that a philosophy of mind and consciousness always is a philosophy of organic life.

In the enactive literature, life is typically understood as *autopoiesis* (Maturana & Varela 1980). Living organisms are spontaneously active, self-maintaining and self-producing systems, and it has been proposed that we should understand our phenomenology in light of this *self-creating* or 'autopoietic' organization. More particularly, an organism lives under precarious conditions – without its activity, its organization breaks down – and the organism must adapt to the environment in order to continue its autopoietic organization (Di Paolo 2005; 2009).

It has been suggested that living organisms display a certain 'concern' for their own persistence, which may be considered a precursor for more advanced forms of concern and purposefulness characteristic of our lives (Jonas 1966; Di Paolo 2005; Thompson 2007). While we may regard living systems as displaying purposeful behavior, aimed at self-preservation and reproduction, some have rejected interpretations of autopoiesis in terms of purpose, promoting a more neutral characterization of autopoietic organization (e.g. Maturana, in Maturana & Varela 1980). Be this as it may, the widely accepted idea (among enactive theorists) is that fundamental properties of our mental lives (including our concern with the environment) build on more basic principles of organization of living systems.

Thus, enactive theorists have linked (adaptive) autopoietic processes to (basic forms of) 'sense-making' (e.g. Di Paolo 2005; Thompson 2007). The reason is that in order to ensure their continued existence, autopoietic systems must successfully cope with their surroundings; the fact that they adequately respond to the environment may (in some basic sense) be interpreted in terms of the organisms 'making sense' of their environments.

The more sophisticated forms of human experience, including our most advanced perceptual capacities, can be considered as extending our basic adaptive autopoietic activity. Perceptual engagement forms a subset of all our intricate interactions with the environment, and clearly perception is a form of engagement with the environment which contributes to the organism's continued existence. While such more advanced capacities may to some extent have a life of their own (they cannot be reduced to the more basic capacities), and introduce their own normativity, they are constrained by the more basic autopoietic activity, which enables the more advanced forms of life (Jonas 1966; Di Paolo 2005). If this is right, understanding our adaptive autopoietic organization may be required for explaining how our perceptual capacities come into being.

Note that an appeal to autopoiesis is fully compatible with the emphasis on perceptual capacities in sensorimotor theory. One may for example simply claim that *in living beings*, sensorimotor skills depend on autopoietic processes. However, as we shall explain below, there can nevertheless be very different views on the precise role of autopoiesis in an enactive account of consciousness.

2.2. Sensorimotor enactivism and autopoietic enactivism

Enactive accounts of perceptual consciousness are focused on interactive capacities. From this perspective, the question is: which capacities? Which of the interactions of a system are relevant for the sensible ascription of perceptual consciousness? We distinguish two mutually exclusive kinds of enactivism, which we define as follows:

Sensorimotor enactivism puts only perceptual capacities center-stage, by claiming that perceptual consciousness can be understood without further appeal to factors outside the domain of perceptual interactions and their behavioral expressions. Autopoietic processes are then only relevant for perceptual experience in so far as they impact on the perceptual capacities and their behavioral expressions.

Autopoietic enactivism claims that there is a *necessary and constitutive* relation between conscious experience and autopoietic processes or associated background capacities. On this view, to explain perceptual consciousness we have to appeal to factors outside the domain of recognizably perceptual interactions (that is, the domain relating to the perceiver's sensing capacities as defined in section 1.1 above) and their expressions. By putting autopoietic processes center-stage, autopoietic enactivism potentially relegates perceptual capacities to a secondary role with respect to consciousness, because the presence of perceptual capacities is no criterion for ascribing consciousness.

Autopoietic enactivism, as we define it, asserts not only that autopoiesis is necessary for conscious experience, but also that it is necessary *constitutively*, rather than 'merely' being instrumentally relevant for enabling perceptual interactions (as an example of instrumental relevance, a scaffold might turn out to be necessary for building a particular type of stone bridge, but it is not an essential part of what these stone bridges are). According to autopoietic enactivism, sensing capacities are only relevant to

consciousness when embedded in the context of an autopoietic system. On this view, a lack of autopoiesis implies the absence of a conscious perceiver.

Note that the difference between the enactivisms concerns a *constitutive* question, in the sense of a question regarding what perceptual consciousness consists in. A comparison may be made with different views on intelligence. While the concept of intelligence is not connected to the capacity to lift heavy weights, we would claim that intelligence is a matter of bodily capacities. It is just not the heavy weight-lifting capacities that are the relevant ones. Similarly, we regard perceptual consciousness as *constituted* by the exercise of bodily capacities of engagement with the environment, and we discuss different views about which capacities are the relevant ones.[3]

Which of these enactivisms forms the best framework for understanding consciousness? A good way to proceed for answering this question, we suggest, is to take sensorimotor enactivism as the default position, and to investigate what if anything is missing. After all, sensorimotor enactivism appeals to a more restricted range of processes compared to autopoietic enactivism, so that we may ask what the latter has to offer that the former doesn't have. In other words, our question then is: *are there reasons for turning to autopoietic enactivism?*

2.3. Getting the contrast right: a focus on capacities

As said, all forms of enactivism, as we see it, focus on capacities of bodily interactions or engagements with the environment; the question is *which* interactive capacities are the relevant ones for perceptual consciousness. While the interactive nature of sensorimotor capacities is evident, the interactive nature of autopoiesis may be less obvious. Before addressing possible reasons for turning to autopoietic enactivism, we shall therefore briefly elaborate on autopoiesis as an interactive capacity, and we shall point out that this conception could potentially allow us to set aside philosophical issues associated with an appeal to interaction-independent inner processes in an account of consciousness.

Our analysis contrasts with a different distinction that may be made, between on the one hand a focus on (sensorimotor) interactions with the environment, and on the other hand a focus on inner (autopoietic) processes (cf. Ziemke 2007). From that perspective, an appeal to autopoiesis might look like the claim that we must peer inside a system before we can decide on the issue of consciousness. On our view, however, this would not do justice to the interactive nature of life, and to the potential relevance of autopoiesis to human experience. In enactivism as we see it, interactions are fundamental, and as Maturana writes: "living systems are units of interaction" (in Maturana & Varela 1980: 9). An

[3] There are two senses of *constitutive*. In our view, perceptual consciousness (like intelligent action) is constituted by the exercise of capacities, in the sense that this is what we take conscious perception (or intelligent action) to be, or to *consist in* (and thus it does not consist in, say, internal representations or ghostly processes accompanying our bodily activities). There is a different sense of constitution, meaning what it is 'materially constituted' by. The present paper does not concern the question whether our perceptual consciousness is 'materially constituted' by (a subset of) autopoietic processes (and of course in living systems, the material constitution is in constant flux). Instead it concerns the question whether the exercise of autopoietic capacities is part of what perceptual consciousness consists in, such that autopoietic capacities would be necessary for spelling out what perceptual consciousness is. If autopoiesis is constitutive of perceptual consciousness in this sense, then for a (natural or artificial) system to have perceptual consciousness it is logically required that the system be an autopoietic system – a claim embraced by autopoietic enactivism but rejected by sensorimotor enactivism.

autopoietic system needs to metabolize and continuously adapt to its environment in order to continue its activity. By its very definition autopoiesis cannot withdraw itself from the interactive context.

Given such an interactive conception of autopoiesis, it may be a mistake to accuse autopoietic enactivism of making a category mistake similar to the ones discussed by Ryle (1949). For example, Ryle argued that it would be mistaken to think that intelligent behavior is intelligence in virtue of hidden operations in an inner realm, as if the behavior and the intelligence are independent things (Ryle 1949). Similarly, one might argue, the presence or absence of conscious experience does not logically hinge on hidden properties of a system's inner organization. On this view, consciousness must necessarily at least potentially show in behavior. Whether or not a particular neural process is relevant to conscious experience depends not just on the process itself, but on the role of the process in governing behavior, on what happens next (Dennett 1991). Indeed, as Wittgenstein (1953) argued, criteria for the use of, say, the word 'sensation,' cannot lie in the privacy of a proposed radically subjective inner world, as in his famous 'beetle in a box' analogy of *Philosophical Investigations* §293, or the knob or wheel that is no part of the machine, *Philosophical Investigations* §270-271. Arguably, the fact that we can and do learn to use psychological concepts in interaction with others supports the notion that "An 'inner process' stands in need of outer criteria" (Wittgenstein 1953 §580).

These are of course highly controversial issues in their own right. In opposition to these views of Ryle, Wittgenstein, and Dennett, some theorists have defended a notion of conscious experience as something that can be systematically independent from the person's activity (e.g. Block 2007; Chalmers 1996). The idea then is that qualia may be absent or systematically different without any noticeable difference in the interaction of the system with its environment. Arguably, such accounts of consciousness that appeal to interaction-independent inner processes (or in the case of Chalmers, independent natural laws) are subject to the objection of 'semantic disengagement', implying that they rely on a notion of consciousness that lacks any significant use – it cannot be used to say anything substantial (Sloman 2010).

Addressing such an issue is to enter an argumentative territory that is very different from the one we aim to engage with here. Indeed, we mention these ideas primarily to emphasize that questions about the relevance of inner processes, independently of the interactive context, are not what differentiates the two kinds of enactivism we discuss here. Commitment to a 'beetle in a box' conception of consciousness is certainly not implied by the position we call autopoietic enactivism (and may even be in tension with it).

At first sight, our understanding of autopoiesis as an interactive process may seem in tension with the fact that the continuity between life and mind has sometimes been framed in terms of an 'interiority' characteristic of living organisms and their 'inner' psychological worlds. For example, in agreement with Varela, Thompson asserts that "autopoiesis is a condition of possibility for the emergence of interiority" (Thompson 2007: 79). However, Thompson does not claim here that the theory of autopoiesis concerns only *body*-internal processes, instead he is concerned with phenomenal interiority. As Thompson remarks, "this emergence of an inside is also the specification of an outside. Thus the dynamic emergence of an interiority can be more fully described as the dynamic co-emergence of interiority and exteriority" (Thompson 2007: 79). While Thompson regards autopoiesis as necessary for the 'interiority' of consciousness, he does not thereby imply that autopoiesis is possible without interaction with the

environment. Autopoiesis is always a process in the organism's interaction with the environment, and arguably it has to be if it is to be of relevance to perceptual experience.

Autopoietic interactions go beyond perceptual interactions: the set of autopoietic interactions includes metabolic interactions that are not included in the set of perceptual interactions (i.e. the interactions displaying the perceiver's sensing capacities). Are there reasons to include such extra-perceptual interactions in an account of perceptual consciousness? In the following sections, we shall critically examine potential reasons to embrace autopoietic enactivism by discussing views on autonomous agency, the affective dimension of experience, and preconditions for artificial consciousness. We shall start with autonomy. In the enactive literature stressing life-mind continuity, autonomy has often been connected with autopoiesis. So let's see if considerations concerning autonomous agency may offer reasons for broadening our view beyond the perceptual interactions as determinants of perceptual consciousness.

3. Autonomous agency

Conscious perceptual experience can be defined as the exercise of sensing capacities by a conscious perceiver. Given that conscious perceivers are autonomous agents, autonomy may be relevant to perceptual consciousness. Here we shall first turn to what we might call the action side of consciousness (autonomous behavior), after which we address a notion of autonomy closer to the heart of autopoietic enactivism.

We conscious human beings are evidently capable of acting on our own goals: clearly our behavior is not fully determined by the present environment, and the more autonomous we are, the less can we be manipulated into pursuing goals that do not fit us. Part of the motivation for an emphasis on life comes from an interest in autonomous agency and associated cognitive/behavioral capacities. The idea is that living organisms (or adaptive autopoietic systems) have the kind of organization allowing for robust forms of autonomy (Thompson 2007; Di Paolo 2005). We shall see how autonomy may be relevant not only to conscious agency, but to perceptual consciousness as well.

There are different conceptions of behavioral autonomy. In the enactive literature one often finds strong requirements for autonomy. For example, as Di Paolo and Iizuka put it, "In order for a system to generate its own laws it must be able to build itself *at some level of identity*. If a system 'has no say' in defining its own organization, then it is condemned to follow an externally given design like a laid down railtrack" (Di Paolo & Iizuka 2008: 410). The implication here is that, if the system's organization is not up to the system, the system is not fully organizationally autonomous. Consequently, one may then take the behavioral autonomy of the system to be restricted. By contrast, more deflationary conceptions of autonomous action could for example focus at the system's present capacities and behavioral tendencies, irrespective of the *origin* of these capacities and tendencies. From this perspective one might argue that even if a system's aims were to be fully determined by external factors in evolutionary or ontogenetic history, or by a designer, they do not therefore fail to be the system's own aims.

Under both readings, the notion of autonomous agency implies that the system's behavior must go beyond a stimulus-response determinism. When the behavior of the system is thus underdetermined by its present environment, we may (under certain conditions) speak of the system as having its own aims. (We say 'under certain conditions', because additional criteria may be added to ensure that the system's contribution goes beyond merely adding random noise, to give a more serious sense of the system

having its own aims.) The stronger notion of autonomous agency proposes additional criteria related to the origin of the system's tendencies. In both conceptions, autonomous agency implies that there are *systematic (non-random) patterns in behavior which are underdetermined by the present environment and that depend on the agent*.

3.1. Sensorimotor enactivism on autonomous behavior

Irrespective of which of the above notions of autonomous agency one favors, *sensorimotor enactivism* can offer a simple way to relate autonomous action to perceptual experience (we'll address a possible complication in subsection 3.2). For by focusing on autonomous behavior, the focus lies, as one might put it, at the action-side of conscious agency. Thus, in line with sensorimotor theory as introduced in section 1 above, we can then claim that autonomous behavior may be important to explain *which* sensorimotor dependencies we might get attuned to. If organizational principles of life, such as autopoiesis, are required to explain the patterns that occur in our behavior – the behavior of living organisms – they will thereby also be relevant to understanding which sensorimotor patterns occur during the development of the perceiver. Thus the organization of life may offer particular possibilities for the development of attunement to these sensorimotor dependencies.

Indeed, if living organisms tend towards certain behavioral patterns, this introduces systematic biases in the opportunity to become attuned to sensorimotor dependencies. Understanding patterns in our (autonomous) behavior will thereby become relevant for explaining our perceptual experience. We may compare the situation to a simple case: there are an astronomical number of physically possible combinations of activity of the muscles, but our biological organization and need for energy optimization, as well as other constraints determined by the necessity to survive, restrict the patterns of muscle contraction that we actualize. This puts constraints on the sensorimotor patterns that occur, and thus impacts at the level of *preconditions* for attunement to sensorimotor patterns. Similarly, whatever principles govern our autonomous agency, and wherever these come from, understanding autonomous agency may be required for understanding the development of perceptual experience. However – sensorimotor enactivism stresses – this does not imply that conscious perception is *constituted* by anything outside the domain of recognizably perceptual capacities.

In addition to such an indirect role of autonomous agency, influencing perception through the change of its developmental preconditions, sensorimotor enactivism could point out that in the case of living beings, aspects of the organization of autonomous living beings (such as their autopoietic organization) will no doubt be relevant for understanding *how* perceptual attunement develops. Sensorimotor enactivism does not doubt that the organization of action is relevant for understanding perception, since perception and action are strongly interdependent and both depend on sensorimotor dynamics (Hurley 1998; 2001), as stressed by sensorimotor enactivism. But note that this does not imply that autonomous living beings embody the *only* possible way for perceptual capacities to develop. Neither does it imply that perceptual consciousness consists of anything beyond the exercise of perceptual capacities themselves. Understanding the organization of life may then be relevant for explaining our perceptual consciousness, not because life is necessary for consciousness (which need not be the case), but simply because we are alive.

In short, sensorimotor enactivism can do justice to a role for autonomy in explaining the perceptual consciousness of autonomous, living beings. Importantly, it does so without jumping to the conclusions

of autopoitic enactivism. What this shows is that the possible importance of life and autonomy does not in itself support autopoietic enactivism.

3.2. Autonomy and autopoietic enactivism

There may be a complication for the above strategy. For in the enactive literature, autonomy often figures as a technical concept that cannot be equated with autonomous behavior. Autonomy is then used as a term referring to an organizational property of living systems. Briefly, in an autonomous system, the processes constituting the system depend on each other to form a unified, self-producing and self-maintaining system that determines a domain of possible interactions with the environment (Thompson 2007: 44; Varela 1979). A living system is autonomous in this sense. It is an adaptive autopoietic system, that is, a system generating and maintaining itself under precarious conditions (e.g. Thompson 2005; 2007; Di Paolo 2009). The idea is that an autonomous system to some extent creates its own laws or rules of engagement, in the sense that an organism 'enacts' its own lifeworld (or *Umwelt*) in a way that is not determined by the surrounding world. As Thompson puts it, based on its autopoietic organization, the animal "meets the environment on its own sensorimotor terms", and systems organized in this way "bring forth what counts as information for them" (Thompson 2005: 418). On this conception, we may not only speak of autonomous behavior, for there is already *autonomy implied in perception itself*, so that we can speak of the autonomy of perception.

The question is whether such 'perceptual autonomy' provides any reasons supporting autopoietic enactivism. One might attempt to defend autopoietic enactivism by claiming that a (strongly autonomous) agent with perceptual autonomy is necessary for conscious perceptual experience. But what would support such a claim? It seems to us that whether we consciously experience or not, does not *constitutively* depend on our history, so it cannot constitutively depend on whether we in some sense have previously self-created the terms of our sensorimotor engagement. If somehow we had emerged a minute ago – like Davidson's (1987) 'swampman' – this does not make our experiences less vivid, or less real.[4] Note that, in a swampman scenario, while lacking history, the creature would still have the capacity to adaptively engage with its present environment. However, autonomous adaptability of perception may be a lovely feature of biological experience, but that does not imply that such adaptability is *necessary* for experience. Suppose that – say due to some drugs – an organism loses its perceptual adaptability, in the sense that it can no longer develop new perceptual and behavioral capacities. It seems far from obvious that this would form a proper basis for denying that the organism has experience, as long as the organism keeps responding to its environment. (Note that the organism's sensing capacities will still be adaptive in the sense that it helps the system to maintain itself; this however is a feature that can easily apply to non-autopoietic robotic systems that have to survive in a hostile environment.) Thus, even if autopoiesis makes perception autonomous in some sense, this

[4] In a swampman scenario there is of course no ontogenetic explanation for experience, beyond an appeal to a cosmic accident; there's no history in which the agent had influence on its own development. Still, when we regard experiencing as the exercise of a bodily capacity, one may give a *constitutive explanation* – an account of what the exercise of the capacity *consists in* – of the newly emerged agent's experience. If you are uncomfortable with the far-fetched swampman scenario, note that the same would hold for robots: whether their bodily capacities are deliberately designed or emerged over a period of robotic development, we may give a constitutive explanation of capacities of interest, even if we have no knowledge of the history of the capacities. We address the issue of robotic consciousness in section 5 below.

would not support the view that perceptual consciousness necessarily requires the particular interactions associated with autopoiesis.

Attempting to connect experience more strongly with autonomy, Thompson (2005) redefines the challenge of *perceptual* phenomenality in terms of the 'body-body problem', that is the challenge of relating our experienced, lived body to our living body. A problem with this approach is that it already assumes that (our) experiences of the living body figure in an account of (our) perceptual experiences. The problem with this is that there seems no reason to suppose that in general conscious experience of the body is required for conscious experience of the world (e.g. we aim to understand the experience of colors, not the experience of being a bodily subject of the experience of color). A further question is whether an autopoietic notion of the body and bodily experience is necessary in such an account. As Thompson notes, sensorimotor capacities must be capacities of an agent or self that is embodying the capacities. Importantly, he subsequently claims that "agency and selfhood require that the system be autonomous" (2005: 417), which he spells out in autopoietic terms. His view of perceptual consciousness then builds on the autopoietic notion of selfhood, based on the idea that conscious experience (or 'first-person givenness') involves self-awareness for experiences to be 'phenomenally manifest as mine' (Thompson 2005: 420).

Suppose that perceptual consciousness indeed depends on autonomy and self-awareness. Would this imply that *autopoiesis* is co-constitutive of perceptual consciousness? We don't think so. Sensorimotor enactivism could argue that there are ways to incorporate autonomous agency and self-awareness in a sensorimotor account without going autopoietic. For example, Kiverstein (2007) suggests that a subjective first-person point of view, in which things are experienced as phenomenally given to the subject, depends on sensorimotor know-how and that this know-how is not necessarily dependent on the particular biological 'wetware' of living organisms. In line with such a sensorimotor enactivism, also Di Paolo and Iizuka write: "we suggest that there are ways of modelling and maybe even instantiating artificial autonomy that do not require building a fully autopoietic artificial system" (Di Paolo & Iizuka 2008: 411).

Our point here is not to defend such a sensorimotor account of autonomy or self-consciousness. We merely wish to stress that we should not presuppose that autopoiesis must be necessary for relevant forms of autonomous agency or self-consciousness. Thus, even if autonomous behavior and signs of self-consciousness are taken to be important criteria for attributing consciousness to a system, this does not in itself support the view that perceptual consciousness constitutively depends on autopoiesis. Perhaps an adaptive autopoietic organization provides the best way to make a system autonomous. Perhaps it could even be argued that autopoiesis consequently provides the best way to generate a system's self-awareness. But neither of this supports the view that there is no autonomy, no self-awareness, and no perceptual experience to be found in a system without autopoietic organization.

In short, in several ways autopoiesis may be relevant to human perception, but that does not support the claim that it is *necessary* for perceptual consciousness in general. Moreover, we might hold that autopoiesis is necessary for perception to be truly autonomous, but this does not mean it is necessary for perception to be conscious. We should not suppose that autopoiesis is a precondition for being a conscious perceiver. Without further reasons, considerations of autonomy do not seem to support the idea that perceptual consciousness is co-constituted by extra-perceptual capacities specifically associated with autopoiesis.

4. Affective dimensions of perception

Conscious perceptual experience has an affective side. It has been suggested that affective aspects of perception may necessitate an appeal to factors outside the range of perceptual sensorimotor capacities (e.g. Bower & Gallagher 2013; Stapleton 2013). Affect and valence have often been brought into connection with autopoiesis. Indeed, living organisms display a certain 'concern' for their own persistence, which may be considered a precursor for more advanced forms of concernfulness characteristic of our (mental) lives (Jonas 1966; Di Paolo 2005; Thompson 2007). Do considerations of affective dimensions of perception warrant favoring autopoietic enactivism over sensorimotor enactivism?

If sensorimotor enactivism can be applied to affective aspects of perception, it will become more obvious that it offers an account of our rich perceptual experiences. This is the more so since non-philosophers may commonly associate subjective experience with valence (Sytsma & Machery 2010). Conversely, if sensorimotor enactivism could not account for the affective, even if its account of other aspects of perceptual consciousness is not thereby disqualified, at best its range of application would be restricted. More severely, autopoietic enactivists may argue that failing to do justice to the affective dimension of perceptual experience would imply failing to account for conscious perceptual experience in general, claiming that without an affective component there is no concern with the environment at all. So let's see if the affective dimension of experience may provide support for autopoietic enactivism, or reasons to abandon sensorimotor enactivism.

4.1. Affect in sensorimotor enactivism

According to Bower and Gallagher, "To perceive a situation in light of sensory-motor contingencies is close to perceiving it impersonally, as *one*, that is, anyone sharing the standard perceptual skill set, would perceive it" (Bower & Gallagher 2013: 123). Would the affective dimension of conscious perception then warrant an appeal to factors outside the range of sensorimotor capacities? We argue this need not be the case.[5]

Contrary to Bower and Gallagher, we think that there is already much that can be captured in terms of one's perceptual engagement with sensorimotor contingencies. Even if people share the same sensorimotor capacities, there can be enormous variation in which sensorimotor dependencies they presently engage with. It is not just that there are individual differences in perceptual capacities: the very same situation can be perceived very differently even by people with identical perceptual capacities. Demonstrations of inattentional blindness are a clear case in point. To take a famous example, subjects engaged in counting the passing of a ball between members of teams in a basketball game can easily fail to notice someone in a gorilla suit walking through the basketball scene (Simons & Chabris 1999). This failure to notice an otherwise salient event depends on the subjects' present concerns (in this case the passing of the ball, and it depended on whether the subject was concerned

[5] There is much of Bower and Gallagher with which we agree, such as when they point out that "sensory-motor contingencies are of no avail to the perceiving agent without motivational pull in one direction or another or a sense of the pertinent affective contingencies" (Bower & Gallagher 2013: 108). What we object to is the suggestion that affective aspects of experience must form a problem for sensorimotor enactivism. On our view, *engagement* with sensorimotor dependencies (the exercise of perceptual capacities) may already capture motivational and affective aspects of experience. We are unsure whether Bower and Gallagher would reject sensorimotor enactivism, and we should emphasize that their paper is not set up as a defense autopoietic enactivism.

with the team of the color that matched the color of the gorilla suit). Thus the perceiver's present interest is expressed in the way in which particular aspects of the scene stand out at the expense of others.

Of course the question is how far this brings us regarding the *affective* aspects of perception. Without claiming to have anything like a worked-out theory about affective aspects of perception, we can still point out that the fact that we can define sensorimotor contingencies at arbitrary levels of abstraction or complexity should make us careful with claims that some perceptual phenomena must lie beyond the reach of a sensorimotor account. Consider for example a case of social perception, say the experience of someone's mood. When we know someone, we may be attuned to the ways to make the person smile, which certainly is a case of changing our sensory stimulation contingent upon our action. Our experience of the person's mood may then be viewed as partly consisting in our engagement with the ways in which the person may be expected to respond to what we do. Perhaps, then, the difference with other perceptual experiences is just a matter of which patterns are involved, for other experiences also depend on our engagement with the way in which sensory stimulation changes depending on our activity.

Arguably, this brief allusion to a sensorimotor way to think of the experience of someone's mood still doesn't capture much of the way in which we may be *affectively* engaged with others. After all, being attuned to social contingencies doesn't imply that we care about others. Still we think it would be too quick to conclude that sensorimotor contingencies "are too impoverished to account for social cognition" (Bower & Gallagher 2013: 127) – for the resources of sensorimotor enactivism are by no means exhausted.

A possibility for further enriching the account of perceptual experience offered by sensorimotor enactivism is to appeal to action tendencies. To start with a simple case outside the social domain, when I'm thirsty I tend to take a glass of water and drink it: rather than disinterestedly inspecting the glass, it 'solicits' or 'invites' action (cf. Dreyfus & Kelly 2007; Rietveld 2008; Withagen et al. 2012). Perhaps the fact that perception does not involve neutral contemplation of the scene may in part be captured by such inclinations to act with regard to the object of experience. Thus, when I am inclined to comfort a friend when she's sad, or when I tend to try to make her laugh when she's cheerful, these behavioral tendencies may be part of the affective aspects of my experience.[6] To the extent that this is the case, affective aspects of the experience lie within the domain of perceptual attunement and its expressions.

We agree with Bower and Gallagher's emphasis on "the individual significance of perceptual experience," and their claim that this "consists at least in part in one's perceptual interest, inclination, or investment, the sense of one's own stakes in a given situation" (Bower & Gallagher 2013: 123). We also agree that an account of the (affective) significance of perception is urgently needed. But we see no

[6] Such an appeal to action tendencies would go significantly beyond the core of the sensorimotor approach (and thus it should not be made lightly). For example, a sensorimotor account of what it is to see a straight line appeals to facts such as that sensory stimulation does not change if one moves one's eyes along the line (O'Regan & Noë 2001). However, attunement to such contingencies need not be expressed in any particular behaviors (behavioral expressions may be highly context-dependent), and sensorimotor theory certainly does not claim that one is inclined to move one's eyes across straight lines! Our suggesting here is merely that the option of extending sensorimotor theory to include action tendencies should not be dismissed prematurely, and we see no reason to presuppose that an appeal to action tendencies would imply a constitutive role for autopoiesis.

reason to exclude the possibility of spelling out perceptual interest and the subject's inclinations in terms of the exercise of perceptual sensorimotor capacities.

4.2. Autopoiesis-based affect without autopoietic enactivism

We certainly do not wish to deny that one might build an account of our intricate affective perceptual experiences as continuous with a more basic autopoietic normativity. We further acknowledge that an account of the capacities of (human) perceivers may for example appeal to the 'normativity' implied by the tendency of living systems to ensure the continuation of their existence. Indeed, spelling out the continuity between human experience and organic life is a highly valuable project. However, we are not convinced that an appeal to extra-perceptual capacities is *necessary* to account for the affective aspects of perceptual consciousness. What is at stake is not whether links can be made between our experience and life. The point is that autopoietic enactivism can only appeal to affective aspects of experience by providing a plausible case for the necessity of autopoiesis for affective perceptual experiences; without such an argument we should not suppose that autopoiesis is essential for affective experience.

In short, sensorimotor enactivism can claim that the affective dimension of perception lies in the perceptual engagement itself, and we presently see no reason to claim that affective aspects of experience must be co-constituted by extra-perceptual interactions (by interactions that lie outside the range of the exercise of sensing capacities). So also when taking into account affective aspects of perception, we suggest that sensorimotor enactivism remains a strong default position for a capacities-oriented enactivism.

5. Artificial consciousness

A good way to further bring into focus the contrast between autopoietic enactivism and sensorimotor enactivism is by considering the preconditions for artificial consciousness. Does artificial consciousness require artificial life or autopoiesis? And if so, why? Consider a missile guidance system that is attuned to sensorimotor dependencies, such as the way in which its 'view' of a target airplane changes when it moves, allowing it to home in on the airplane (O'Regan & Noë 2001). While the missile with guidance system can be said to exercise its sensorimotor capacities, we do not ascribe conscious experience to the system. What does the system lack?

5.1. What the missile guidance system lacks

While sensorimotor enactivism would claim that the system lacks the relevant perceptual capacities constituting conscious experience, autopoietic enactivism claims that the system lacks the autopoietic organization required for conscious experience.

As Thompson notes, the missile guidance system is not autonomous in the autopoietic sense: "It is not a self-producing and self-maintaining system that actively regulates its own boundary conditions so as to ensure its continued viability" (Thompson 2007: 261). He claims that *therefore* the system has "no genuine sensorimotor knowledge or mastery" (p. 261) but only ascribed knowledge, "merely attributed to the system by the observer" (p. 260).

Appealing to life here – instead of to the limitation in the missile guidance system's sensorimotor capacities – reveals commitment to autopoietic enactivism. Autopoietic enactivism attributes

consciousness and 'genuine' mental capacities exclusively to autopoietic systems, rejecting such attribution where the autopoiesis is lacking.

In line with Thompson's ideas, Noë (2004: 230) refers to a lack of a 'unified subject.' Although he doesn't want to restrict consciousness to living beings a priori (p. 231), he speculates that: "Living beings are already, by dint of being alive, *potentially conscious*. Living beings, that is, have just the sort of unity that makes it plausible that they might be able, for example, to "keep track" of patterns of sensorimotor covariation. A robot, in contrast, or a bit of "mere matter," seems to lack precisely that sort of unity" (Noë 2004: 230). This remark is strongly suggestive of autopoietic enactivism, as it differentiates between conscious and non-conscious systems on the basis of the presence or absence of organizational principles of life.

However, there are equally strong hints of sensorimotor enactivism in Noë's work, for example where he states that "where the sensorimotor repertoire is rigid and simple, there is no compelling reason to attribute mind or experience. But where the sensorimotor repertoire increases in complexity, the notion that there is also primitive awareness becomes more plausible" (Noë 2004: 229). In general, Noë's work is strongly focused on perceptual capacities, which can be reconciled with the emphasis on unity in the quote above, by simply insisting that the relevant unity must be found in the perceptual capacities of the system or the way in which these are expressed (e.g. in thought, speech and deliberate action). Of course questions would remain concerning the precise nature of this 'unity', but such a position could be compatible with sensorimotor enactivism.

As both the range of perceptual capacities, as well as the range of ways to express these capacities, are severely restricted in the case of the missile guidance system – as well as in (most) present-day robots – we do not have to appeal to anything beyond these perceptual capacities to distinguish us from such machines. But neither have we offered reasons to reject autopoietic enactivism. So the question remains which interactions are relevant for separating us from the machines. Which of autopoietic or perceptual sensorimotor repertoires better supports the sensible ascription of conscious experience?

5.2. Comparing life with non-living systems

When comparing living and non-living systems, it is important to make a fair comparison. It's not so informative to compare the extremely restricted capacities of a missile guidance system with the elaborate capacities of a living organism. To further sharpen the contrast between autopoietic enactivism and sensorimotor enactivism, let us therefore imagine two extreme cases.

First, consider the possibility of a living system that has only the sensorimotor capacities of the missile with its guidance system, plus some basic aspects of metabolism. Let us assume for the sake of argument that such a system is possible, and let us further suppose that its missile-guidance activity is properly integrated with the system's autopoietic organization. To make sense of such a scenario, we could for example assume that there are artificial restrictions to limit a living system to a very restricted set of capacities, or to put it the other way around, we could imagine to integrate life in a robot with a very restricted sensorimotor repertoire. Presumably we should give the system some time to develop, such that the autopietic activity gets integrated with the sensorimotor activity of the system: let's say that the artificial part of the system ensures that nutrient availability depends on success in missile guidance.

Second, consider an artificial system sharing an enormous range of capacities with human beings. Suppose the system could walk and talk, it could respond to the environment as well as reflect on its own experiences or 'experiences', and it would allow the application of a behavioral distinction between conscious experience and unconscious responses much like in the human case. It would respond to humans depending on their emotional state, and it would itself show behavioral signs that are in humans associated with emotions, affect, and motivations. Let's say it would even write papers on the mystery of its conscious experience. However, the system would run on batteries rather than on food, and it would not be self-producing, so that after some time it just breaks down due to wear and tear.

The question is whether it would be plausible to ascribe perceptual consciousness to these systems. Supposing for the sake of argument that the above scenarios make sense, autopoietic enactivism would have to assert experience in the first scenario, and deny it in the second scenario. Sensorimotor enactivism would deny experience in the first scenario, and assert in in the second one.

Of course one might be skeptical about the imagined possibilities, and we certainly should not presuppose that the sketched scenarios are genuine possibilities. Especially the second scenario may raise the suspicion that we might be imagining something impossible here. However, it is important to recognize what might motivate such skepticism. One might hypothesize that an autopoietic organization will be required to get an interesting range of capacities, such that the ascription of perceptual consciousness would become even remotely plausible. This may be so. But note that such a response can be fully in the spirit of sensorimotor enactivism. After all, the sensorimotor enactivist might insist that the perceptual capacities themselves are the ones that warrant ascription of perceptual consciousness, even if, as one might hypothesize, autopoiesis may be instrumentally required for the development of these capacities or for having them. If life turns out to be crucial for behavioral capacities associated with perceptual consciousness, this would therefore not support an appeal to factors outside the domain of perceptual interactions; it would provide no reasons to abandon sensorimotor enactivism. Skepticism about the possibility of the imagined scenarios does in no way tip the balance towards autopoietic enactivism.

While these considerations may make it difficult to differentiate between the different forms of enactivism, this certainly does not imply that the positions coincide. The enactivisms have clearly different conceptions of what constitutes conscious experience, and they appeal to different considerations for ascribing perceptual consciousness to a system. As a result, the enactivisms can have different reasons for appealing to life in an account of consciousness. For example, the autopoietic enactivist might insist that life provides an intrinsic unity or a self-producing metabolism that in themselves provide for consciousness. The sensorimotor enactivist on the other hand would claim that if life provides something special, it is doing so by providing the particular range of perceptual capacities, and not because life intrinsically adds some additional ingredient.

The different forms of enactivism embody different conceptions of what legitimates ascription of perceptual consciousness, and considerations of the scenarios sketched above can help to sharpen the contrast between these conceptions. Indeed, when we consider the scenarios above as potential possibilities, interesting consequences follow. When carefully considering the scenarios, in which we make a more fair comparison between an artificial system and a living system, we can note that the intuitive connection between life and perceptual consciousness breaks down. In the first scenario, the life integrated with the sensorimotor capacities of the system might be the life of a bacterium, and it

seems that there is no more reason to ascribe perceptual consciousness to the system than there is reason to ascribe consciousness to it at all. In fact, the ascription of perceptual consciousness in this case would come with a lack of any behavioral criteria for application of a conscious/not-conscious contrast. This would result in a concept of perceptual consciousness that becomes disconnected from the broad range of perceptual capacities that seem of fundamental relevance to the sensible ascription of perceptual consciousness. Also when considering the second scenario, imagining the absence of consciousness would bring us in an argumentative territory that we aimed to sidestep here (section 2.3), but to which we must now briefly return.

In section 2.3 we stressed our focus on capacities, rather than building our discussion on an internal-external dichotomy. We did so in an attempt to avoid a 'beetle in a box' conception of perceptual consciousness, according to which experience becomes conceptually disconnected from everything we say and do. We would like to emphasize once again that for autopoietic enactivism, just as in the case of sensorimotor enactivism, experience does not become conceptually divorced from the system's interaction with the environment. However, if we take the possibility of the second scenario seriously, autopoietic enactivism would impair the link between perceptual consciousness and some capacities of the system, including its verbal behavior. It is true, and also recognized by sensorimotor enactivism, that there already exist robots with what we may call 'communicative capacities' in a deflationary sense, and which we would not describe as conscious beings. But the point is that, as the level of capacities increases, the intuitive tendency to think of a system as 'merely' a machine would come under increasing pressure. When the sensing capacities of the system are sufficiently close to ours, and if the system can express these capacities in a broad range of behaviors, e.g. including verbal behavior, then at some point a behavioral distinction applies which corresponds to what in humans would be the distinction between conscious and unconscious responses, and we could identify apparent cases of verbal reports of experience. By claiming that the system still lacks perceptual consciousness, autopoietic enactivism would deny the necessary link between behavioral expressions of sensing capacities and perceptual consciousness.

Thus, if, e.g. in the admittedly artificial scenarios described above, autopoiesis and behavioral expressions of sensing capacities may come apart, autopoietic enactivism would choose autopoietic capacities as requirement for the ascription of consciousness. Considering the presence of artificial consciousness, autopoietic enactivism would have to say that we have to find out whether a system eats, or otherwise get the resources for self-production and self-maintenance, in order to find out whether the system is capable of conscious perception. In other words, while we do not necessarily have to 'peer into' the inner workings of the system, we would have to 'peer into' capacities other than the sensing capacities and the ways these are expressed to decide whether these capacities pertain to perceptual consciousness.

Whether or not this should be interpreted as a turn 'inward' may not be so relevant after all. For it seems that autopoietic enactivism runs the risk of committing to a notion of consciousness that is conceptually divorced from the capacities that we associate with perceptual consciousness. When we learn to use the concept of conscious perceptual experience, we do not appear to make use of knowledge of metabolism. If that is right, autopoietic enactivism would have to say that this practice fails to do justice to perceptual consciousness. Just as Block (2007) and Chalmers (1997) defend conceptions of consciousness that are divorced from the behavioral expressions of consciousness, autopoietic enactivism would then defend a conception of perceptual consciousness that becomes

relatively independent of behavioral expressions of sensing capacities. Autopoietic enactivism thereby threatens to distance us from the capacities that arguably should form the starting-point of theorizing about perceptual experience.

In sensorimotor enactivism, by contrast, the notion of conscious experience remains strongly linked to criteria for the ascription of perceptual capacities and their behavioral expression. As a result, considering the scenarios sketched above, it seems that sensorimotor enactivism would force a wedge between the notion of perceptual experience and autopoietic processes. We cannot claim to have shown that this is right. However, what we have shown, we think, is that the attempt to choose between the positions forces us into the direction of the argumentative territory we set forth to sidestep here, concerning the conceptual connections between perceptual consciousness and the associated behavior that forms the very basis for learning the meaning of this public concept.[7] Autopoietic enactivism may come in conflict with the idea that the 'inner processes' of perceptual consciousness stand in need of the outer criteria for perceptual capacities.

5.3. Does artificial consciousness require artificial life?

We have rejected the idea that comparing artificial systems with living systems provides argument in favor of autopoietic enactivism. Thus we would suggest to focus on perceptual capacities rather than presupposing that we must including autopoietic capacities in an account of perceptual consciousness. Perhaps artificial consciousness depends on artificial life or autopoiesis, not for the reasons to which autopoietic enactivism is committed, but because an interesting range of perceptual capacities may only result from an autopoietic system. But *this* is an empirical question which doesn't answer the question of which enactivist framework we should favor.

Our concern here is with different accounts of what constitutes perceptual consciousness, and we have considered potential reasons to turn to autopoietic enactivism about perceptual consciousness. The possible necessity of autopoiesis for certain forms of autonomy is not at stake here, and we need not even decide whether autopoiesis may be required for the development of an interesting range of perceptual capacities. When making a fair comparison between artificial systems and biological systems, we do not find reasons to reject sensorimotor enactivism about perceptual consciousness, but rather potential (though controversial) reasons to be suspicious of autopoietic enactivism. We would then suggest to focus on perceptual capacities as the central notion in an account of perceptual consciousness.

[7] A note on zombies and the living dead: One way to contrast sensorimotor enactivism and autopoietic enactivism is as follows, in the b-movie style familiar in the philosophy of mind. We could define 'zombies' (a different notion of zombie from the one of Chalmers 1996) as systems exercising the full range of behavioral and sensing capacities that we have, while lacking conscious experience. And we could define the 'living dead' as exercising the full range of behavioral and sensing capacities that we have, while lacking the autopoietic organization characteristic of life. Sensorimotor enactivism then rejects this notion of zombies, whether or not the notion of the living dead sketches a genuine possibility. Autopoietic enactivism can only reject the notion of zombies by rejecting the possibility of the living dead, for accepting the idea of the living dead would commit the autopoietic enactivist to the notion of a zombie. (On our view, in line with sensorimotor enactivism, the possibility of the living dead is an empirical issue regarding machine consciousness. However, we would claim that the idea of zombies is not an empirical issue, for we take this notion to be nonsensical.)

How our perceptual capacities relate to autopoiesis then remains an open question. Di Paolo holds a similar view regarding intentionality, by explicitly leaving it open how intentionality relates to metabolism, when he suggests that: "We may invest our robots not with *life*, but with the mechanisms for acquiring *a way of life*, that is, with habits. This may be enough for them to generate a natural intentionality, not based now on metabolism, but on the conservation of 'one' way of life as opposed to 'another one'" (Di Paolo 2003). Similarly, we suggest that how to relate our perceptual consciousness to our metabolism may remain an open question. The reason is that metabolic processes can be causally relevant without being an essential part of an account of what consciousness consists in.

Sensorimotor enactivism offers a framework for thinking about perceptual consciousness that applies both to living organisms, as well as potentially to non-autopoietic artificial systems. On this view autopoietic organization, although relevant in living organisms, is not itself essential for perceptual consciousness.

Conclusion

How does sensorimotor theory relate to other forms of enactivism stressing the continuity between mind and life (and autopoiesis, autonomy, and valence)? Two frameworks have been sketched. In *sensorimotor enactivism*, autopoiesis may be relevant for understanding how the development of perceptual capacities occurs, but it is the perceptual capacities themselves, and not the link with autopoeisis that enables or constitutes experience. In *autopoietic enactivism*, by contrast, our perceptual capacities are *only* relevant to conscious experience in virtue of a larger set of capacities characteristic of our autopoietic organization. In particular, genuine concernful conscious engagement, as characteristic of our affectively laden experience, is under this view thought to depend necessarily and essentially on autopoietic organization. In this paper we discussed and rejected possible reasons for favoring autopoietic enactivism relating to autonomy, affective aspects of perception, and the comparison of living and non-living systems. We further considered scenarios suggesting that autopoietic enactivism could potentially come in conflict with the idea that the ascription of perceptual consciousness should be grounded in publicly observable perceptual capacities. We conclude that without further reasons favoring autopoietic enactivism, sensorimotor enactivism forms the default capacities-oriented form of enactivism.

Acknowledgements

We thank Sanneke de Haan, Erik Myin, David Silverman and the audience at the 2014 AISB conference at Goldsmiths' University of London for helpful discussions and critique on an earlier version of this paper. We further thank our anonymous reviewers for their excellent comments and suggestions. The work was supported by ERC advanced grant 323674 "FEEL" of J. Kevin O'Regan.

References

Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, *30*(5), 481–548.

Bower, M., & Gallagher, S. (2013). Bodily affects as prenoetic elements in enactive perception. *Phenomenology and Mind*, *4*(1), 78–93.

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

Clark, A. (2006). Pressing the flesh: a tension in the study of the embodied, embedded mind?

Cooke, E., & Myin, E. (2011). Is trilled smell possible? How the structure of olfaction determines the phenomenology of smell. *Journal of Consciousness Studies*, *18*(11-12), 59–95.

Davidson, D. (1987). Knowing one's own mind. *Proceedings and Addresses of the American Philosophical Association*, *60*(3), 441.

Degenaar, J., & O'Regan, J.K. (2015). Sensorimotor theory of consciousness. *Scholarpedia* 10 (5): 4952.

Dennett, D. C. (1991). *Consciousness Explained*. Little, Brown and Company.

Di Paolo, E. A. (2003). Organismically-inspired robotics: homeostatic adaptation and teleology beyond the closed sensorimotor loop. In K. Murase & Asakura (Eds.), *Dynamical systems approaches to embodiment and sociality* (pp. 19–42). Adelaide: Advanced Knowledge International.

Di Paolo, E. A. (2005). Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the Cognitive Sciences*, *4*(4), 429–452.

Di Paolo, E. A. (2009). Extended life. *Topoi*, *28*(1), 9–21.

Di Paolo, E. A., & Iizuka, H. (2008). How (not) to model autonomous behaviour. *Bio Systems*, *91*(2), 409–423.

Dreyfus, H., & Kelly, S. D. (2007). Heterophenomenology: heavy-handed sleight-of-hand. *Phenomenology and the Cognitive Sciences*, *6*(1-2), 45–55.

Hurley, S. L. (1998). *Consciousness in Action*. Harvard University Press.

Hurley, S. L. (2001). Perception and action: alternative views. *Synthese*, *129*(1), 3–40.

Hutto, D. (2005). Knowing what? Radical versus conservative enactivism, *4*(4), 389–405.

Jonas, H. (1966). *The Phenomenon of Life: Toward a Philosophical Biology*. Northwestern University Press.

Kiverstein, J. (2007). Could a robot have a subjective point of view? *Journal of Consciousness Studies*, *14*(7), 127–139.

Maturana, H. R., & Varela, F.J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. Springer.

Maye, A., & Engel, A. K. (2012). Time scales of sensorimotor contingencies. In H. Zhang, A. Hussain, D. Liu, & Z. Wang (Eds.), *Advances in Brain Inspired Cognitive Systems* (pp. 240–249). Springer Berlin Heidelberg.

Noë, A. (2004). *Action in Perception*. MIT Press.

O'Regan, J. K. (2011). *Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness.* New York: Oxford University Press, USA.

O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, *24*(5), 883–917.

Philipona, D. L., & O'Regan, J. (2006). Color naming, unique hues, and hue cancellation predicted from singularities in reflection properties. *Visual Neuroscience*, *23*(3-4), 331–339.

Rietveld, E. (2008). Situated normativity: the normative aspect of embodied cognition in unreflective action. *Mind*, *117*(468), 973–1001.

Ryle, G. (1949). *The Concept of Mind*. University of Chicago Press.

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception*, *28*(9), 1059 – 1074.

Sloman, A. (2010). Phenomenal and access consciousness and the "hard" problem: a view from the designer stance. *International Journal of Machine Consciousness*, *2*(01), 117–169.

Stapleton, M. (2013). Steps to a "Properly Embodied" cognitive science. *Cognitive Systems Research*, 22–23, 1–11.

Sytsma, J., & Machery, E. (2010). Two conceptions of subjective experience. *Philosophical Studies*, *151*(2), 299–327.

Thompson, E. (2005). Sensorimotor subjectivity and the enactive approach to experience. *Phenomenology and the Cognitive Sciences*, *4*(4), 407–427.

Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press.

Varela, F.J. (1979). *Principles of Biological Autonomy*. New York: Elsevier.

Varela, F. J., Thompson, E., & Rosch, Eleanor. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.

Withagen, R., de Poel, H. J., Araújo, D., & Pepping, G.-J. (2012). Affordances can invite behavior: Reconsidering the relationship between affordances and agency. *New Ideas in Psychology*, *30*(2), 250–258.

Wittgenstein, L. (1953). *Philosophical Investigations*. Blackwell.

Ziemke, T. (2007). The embodied self: theories, hunches and robot models. *Journal of Consciousness Studies*, *14*(7), 167–179.