Consciousness as the dissipation of information

Abstract

Despite its growing appeal for the study of consciousness, the notion of entropy has yet to lead to widely supported new insights about the nature of phenomenal experience. Typically, entropy measures of brain activity are found to correlate with cognitive functions that are assumed to index consciousness. Taking a very different approach, this theoretical framework does not conflate consciousness with any function. It presents a series of premises to argue that consciousness is fundamentally characterized as inactionable perception, i.e. that does not give rise to macrophysical action. This is then fitted in a framework of perception and action as informational changes in a dynamical neural state space. In this model, inactionable perception naturally arises as the prediction-driven increase of concept-related entropy. This entails an increase of (Shannon) information while its efficacy to produce macrophysical effects decreases, which is here referred to as information dissipation, analogously to energy dissipation in thermodynamic systems. It results from inefficient sensorimotor coupling with the environment, which occurs when behavior is not fixed relative to the stimulus. Despite the posited inefficacy of conscious perception, it consists of action-specific information and can therefore be interpreted as potential behavior.

Starting from fundamental properties, this framework may provide a new and coherent conceptual basis for a fuller understanding of what consciousness is and how it relates to the physical world. Although many of its implications remain to be explored, it appears consistent with empirical findings, and prompts subtle reinterpretations of some classical results in perception research.

1 Introduction

Entropy has become a tool of significant interest in the field of consciousness research. Studies that measure entropy in the signals picked up by recording brain activity (e.g. with EEG) have demonstrated that it is related to different states or levels of consciousness (see e.g. Keshmiri 2020). High entropy of brain activity can be seen as reflecting a high amount of information carried on the underlying networks. It may also be interpreted as indicative of activity being more random or disorderly, or having more complexity. These notions are related through Shannon's mathematical theory of communication (1949), which states (very) basically that higher entropy signals reduce more uncertainty about their source, and thus contain more information. Hence, Shannon, or information-theoretic entropy (H) serves as a measure of information (Jaynes 1957). More formally, it reflects the probability distribution of possible configurations of elements, depending on statistical regularities among the elements (Attick 2011). H increases if the probability for each possible configuration to occur gets more similar (so there is more uncertainty about the outcome) or if the total number of possible configurations gets larger (also implying higher uncertainty; Shannon & Weaver 1949/1998). Because high entropy signals are more difficult to predict, they appear more random or disorderly and can be said to have a higher degree of complexity.

Because of the link between entropy of brain activity and conscious states, it has been suggested that a certain level of entropy might even be the crucial ingredient for consciousness to arise. For example, it has been asserted that consciousness "might be supposed to be an emergent property of a web of connections endowed with certain complexity" (Erra et al. 2017); that the brain's complexity is "its fundamental property that underlies the manifestation of […] consciousness" (Keshmiri 2020); and that "the maintenance of rich brain dynamics measured by entropy is a critical aspect of conscious awareness" (Castro et al. 2024). But such claims are not warranted, at least not if they are about phenomenal consciousness (Block 1995). Much of the work they are based on relies on distinguishing global conscious states, which are ultimately determined by a person's various cognitive abilities (Rankaduwa 2023). Many cognitive functions may of course stand in some relation to consciousness. But the nature of this relationship is not known. Thirty years after Chalmers' (1995) influential formulation of the hard problem of consciousness, there is still no good reason to assume that phenomenal consciousness can be reduced to cognitive functions. It is thus not clear whether the measured changes in entropy values of brain signals are causally related to consciousness itself or only to the associated cognitive functions.

Hence it is interesting that Carhart-Harris (2018) claims that his entropic brain theory is about the qualia or subjective quality of conscious states. The theory asserts that some conscious states, specifically the psychedelic

state, are endowed with a greater richness of experience, and that this is reflected by greater measures of entropy of brain activity. In normal wakeful consciousness, it is claimed, the brain has to reduce its entropy to function effectively. However, according to the author, *richness of experience* refers to its informational content and to how unpredictable it is. This suggests a measure of quantitative aspects of consciousness, not of its qualitative aspect. The empirical work supporting the entropic brain theory thus confirms the existence of a relation between the quantitative contents of experience and underlying neural activity. But this does not provide us with a new perspective on why this is so or how it might work.
Also Beshkar (2023) presents a view in which reducing entropy levels of brain activity is crucial for the emergence of normal waking consciousness. Here consciousness does not refer to a global state, and the entropy decrease does not pertain to the whole brain or specific subsystems. Rather, Beshkar expands on Gupta and Bahmer's (2019) description of entropy decrease along the cortical hierarchy as an important mechanism for perception. According to Beshkar, this reflects reduction of uncertainty about a stimulus. When it gets below a certain level, the observer is assumed to become conscious of the stimulus. Why this would be the case, though, remains unexplained. The framework I will propose in this paper largely shares the emphasis on the importance of entropy reduction through the cortical hierarchy as a crucial mechanism in the explanation of consciousness. But it strongly rejects the notion that this reduction (even if done "beyond classical limits"; Beshkar 2023) would be sufficient for consciousness to arise in any plausible way.

2 An information-based framework for perception, action and consciousness

For the relation between entropy of brain activity and phenomenal consciousness to gain real explanatory power, the first has to be directly related to distinctive, qualitative aspects of the latter, which is notoriously difficult given its infamous elusiveness. The purpose of this paper is to present a way in which this relationship might plausibly exist, and to explore some of its implications. This will be done on the basis of five propositions (P1 to P5) that will each be substantiated separately and then synthesized into a mechanism for the emergence of phenomenal consciousness. This will allow linking consciousness to the information-theoretic notion of entropy in a fundamental but (to the best of my knowledge) previously unexplored way. Importantly, consciousness will not be conflated with cognitive or behavioral functionality but instead (P1) it is argued to arise from perception that does not lead to action. This fits naturally in a view in which (P2) perception is the dynamical, continuous transformation of information from highly fluctuating raw stimulus information to a stable behavioral repertoire, and (P3) it is organized in concepts that consist of action-specific information. The transformation results from an interplay between two opposing processes: (P4) an increase of entropy in concepts required for perception, which (P5) has to be reduced for effective action. The reduction of entropy during perception may be less than fully efficient, which means that some perceptual information may remain inefficacious. I will refer to this as dissipation, analogously to energy transformations involving an increase of thermodynamic entropy (which is formally equivalent to information-theoretic entropy; Weilenmann et al. 2016). Such dissipation of perceptual information would occur when sensorimotor coupling with the environment is not direct, i.e. when a behavioral choice can be made, and thus coincides with the typical experience of conscious behavior. The possible implications of the proposed framework touch on some outstanding philosophical questions and may put several empirical results, some of which will be discussed in the discussion section, in a different light.

2.1 P1: Consciousness is inactionable perception

Consciousness has been notoriously hard to describe, particularly in regard to phenomenal experience – which this paper is concerned with. Authors on the topic often rely on an assumed common understanding of what it means to have first-person or subjective experience. As Nagel (1974) famously put it: for an organism to be conscious there is something "it is like" to be that organism. Wording of this kind is routinely used in an attempt to define consciousness (Farrell 2016). The deliberate vagueness of the phrase "what it is like" seems to convey the idea that as soon as we try to specify what it actually is like, our descriptions do not quite capture the idea of conscious experience anymore. This becomes clear for example when we imagine trying to describe in objective terms what it is like to someone who has never experienced "it" (e.g. color perception to someone who is congenially blind). It appears to be a hopeless task. This evasiveness to any kind of description is such a remarkable feat that it suggests an essential aspect of our conscious experience, which is often referred to as its ineffability (Edelman 2008, p. 87). That is, in Dennett's (1988/2021) words: "one can not say to another, no matter how eloquent one is and no matter how cooperative and imaginative one's audience is, exactly what way one is currently seeing, tasting, smelling and so forth." This is not simply the result of a loss of information, as Ji et al. have argued (2024). They claim that the ineffability of conscious experience is a consequence of computational processing between the cognitive states underlying our experiences and the linguistic messages that we use to express them. Not all low-level information is retained during this process, so by knowing the message, an amount of uncertainty about the underlying cognitive states remains. According to Ji et al. ineffability just refers to an inability to accurately describe one variable by another. But this is not what philosophers of mind mean with ineffability in the context of phenomenal consciousness. Their point is that during conscious experience information about its content seems very much available, but just not in a way that is reportable, which points to a qualitative rather than a quantitative aspect. This is what Chalmers (1996) meant when he wrote that "information (or at least some information) has two basic aspects, a physical aspect and a phenomenal aspect" (p. 23). Below, I will argue that this kind of ineffability can be seen as an aspect of an underlying, even more fundamental property of consciousness.

Such an underlying property of consciousness should not only account for ineffability. Other qualities have been mentioned in the literature, of which some can be considered more or less canonical. Apart from ineffability, Dennett (1988/2021) includes also privateness, intrinsicness and immediate apprehensibility as traditional properties. I argue that these can be reduced to a single quality. First we shall dispose of the property of immediate apprehensibility, which Dennett has convincingly argued to be illusory. For consciousness to be apprehensible in any way it would have to be internally observed in a "Cartesian theater", but there obviously is no such place in the brain, nor could there be (Dennett 1991/2007). The remaining traditional properties are more tenacious. Privateness, according to Dennett, refers to the unique evasiveness of the qualities of consciousness to any kind of interpersonal comparison, while intrinsicness implies that they are atomic and unanalyzable. The latter two properties, as well as ineffability, seem to point out important aspects of consciousness. But they may be more parsimoniously accounted for by postulating that consciousness can be more fundamentally characterized by the impossibility to conceptually subdivide a particular experience. Any description, comparison or analysis of what it is like to have a certain experience requires the conceptual availability of at least some components that make out that experience. The lack thereof would thus entail ineffability, privateness and intrinsicness. The non-conceptual nature of our perceptual experience has also been argued for by others, often on the basis that its details are too fine-grained to be captured by our conceptual capacities (Wright 2003). Even though our perceptual experience seems undeniably conceptually structured (Shieber 2010), the granularity of this structure appears to be too coarse to pick out the differences that make out the experiential instances of the concepts in our possession. This line of reasoning, characterizing conscious experience as non-conceptual, can be further pursued if we assume that concepts ultimately emerge from interpersonal communication. That is to say that concepts (at least the ones that are not fully determined by physical regularities) are based on (implicit) conventions between people (Feldman-Barrett 2017/2018). An impossibility to conceptualize conscious experience then implies the impossibility to be communicated, i.e. non-communicability. Still we have not arrived at a foundational property of consciousness, since communication can be achieved by any arbitrary means of behavior. This depends exclusively on the interpretation by the receiver. Any brain state that results in action can in principle result in communication. But it is also plausibly the case that any communication must involve some action (doing nothing might sometimes count as communication, but this requires resisting environmental perturbations, which might be considered an action as well). The a priori indistinguishableness of action and communication echoes Watzlawick's communication axiom that humans are necessarily always engaged in behavior and that all behavior is communication (Watzlawick et al. 1967/2011). It follows that brain states that can not lead to action and brain states that can not lead to communication are the same states. Therefore, instead of relying on arbitrary communicative inefficacy, I take the impossibility to lead to action in general, i.e. inactionability, as a general constitutive property of conscious states.

Another broadly accepted property of consciousness is that it is always *about* something. This something can be said to be the object of our perception, if we consider perception not only the processing of information from the external world, but also of information from the body (Witt & Riley 2014) and even from other brain processes as a source of internally generated perception (Fuster 2003). Since consciousness always occurs during perception in this broad sense, I propose a definition of consciousness as inactionable perception.

Now we have arrived at a conclusion that seems to contradict something that many believe to be an obvious truth: the causal efficacy of consciousness (as in the belief that I raise my arm because I consciously decided to do so). Of course others have disputed this convincingly on empirical grounds (see e.g. Halligan & Oakley 2021). However, philosophical arguments for an epiphenomenal view often rely on the putative non-physical nature of consciousness (see e.g. Block 1995; Jackson 1986/2021), which as Mathieson (2024) points out, is a premature assumption. In my proposal, on the contrary, the inactionability of consciousness is not argued for on

the basis of non-physicality, but derived from its phenomenal properties. Non-physicality, conversely, does not follow from the inactionability of consciousness either, because the latter is only a statement about its effects, not about its constitution. As a physical process then, it seems clear that conscious perception must have at least microphysical effects, e.g. in the sense of changing neuronal potentials. Therefore, the inactionability of consciousness must be specified as inefficacy on a macrophysical level. Indeed, this is where behavior takes place as it is the scale of motions that are considered under our control (Sturgeon 1998).
If my reasoning so far has been correct, the mind-body problem comes to stand in a different light. The question then is not how can a seemingly non-physical mind have macrophysical effects, but under what conditions does the physical process of perception fail to have macrophysical effects. This is addressed in the remainder of this paper.

2.2 P2: Perception is the continuous transformation of raw stimulus information into a behavioral repertoire

It has been argued that a clear perspective on consciousness has been obstructed by traditional ways to view the mind, in which perception, cognition and action are seen as sequential steps (Hurley 2001). More contemporary approaches consider these processes as much more intertwined, and there is a large body of evidence supporting the view that action and perception are not separate processes. This includes experimental results showing that perception is often action-specific, such as when our perceived spatial environment scales with our possibilities to act in it (Witt, & Riley 2014). Clark (2014) writes that it has become increasingly clear that “Perception is not a general purpose system, delivering input to a distinct and independent action system” and that “the products of perception […] may instead constitute direct recipes for acting and intervening” (p. 105). According to Hurley (2001; p.15), perception and action can be seen as forming a continuum in which they are not just instrumentally but constitutively interdependent or co-constituted, reflecting "co-dependence on the same […] systems of input-output relations and feedback."
In the literature a reservation is often made that the intimate link between action and perception as described above can not be the case for all perception. Some perception would just be “for the sake of perception” (Wilson 2002), a point that is related to the putative existence of a dorsal “how path” and ventral “what path” of neural sensory processing. But even if “what-processing” is more involved in identifying patterns and objects, it may still be seen as ultimately serving action, specifically what Turvey (1977) called “identifying reactions”, such as naming a perceived object. As Barsalou (2012; p. 241) puts it: “Even when naming is implicit (i.e. subvocal), this can be viewed as the production of a word instance, grounded in a motor and auditory simulation.” Such covert behavior may only lead to memory formation to be used later for arbitrary actions. But this does not effectively argue against co-constitution of perception-action. Memory-based behavior may be seen as the result of information stored during previous interactions, providing input for ongoing perception on a later occasion as part of an internal and external stimulus array.
My proposal embraces the continuous view on which perception and action are intimately intertwined and co-constituted by the same networks. It posits that the underlying system can not be separated in a meaningful way in an action part and a perception part. Such an integrated system can be characterized as subserving essentially one process: the transformation of raw stimulus data into an executable form that can be used for action.

2.3 P3: Perception is organized in concepts that consist of action-specific information

In this section I aim to describe the perceptual transformation process, as introduced above, in informational terms – that is, referring to the states of its underlying neural networks. Because the perception-action process is continuous, I consider the relevant neural networks as directly connecting input patterns (originating from sensors or other brain systems, such as during imagining) to action-specific motor patterns. These networks do not just passively respond to the world, but adaptively change their response properties (Roelfsema & Holtmaat 2018; Pakan 2018) based on interactions with the environment. This way the conceptual system acquires knowledge about the world (see Barsalou 2012).
At this point it is useful to adopt a dynamical perspective of neural processing, in which the perception-action process can be seen as a trajectory through a space of possible brain states (Clark 1994). This multidimensional space represents the brain’s possible states, where each of the many dimensions spans the possible values of a relevant variable of one of the information-carrying units (e.g. neurons). This space can be divided in subspaces, or regions, containing states that are involved in the expression or experience of the same object of perception.

This pertains not only to physical objects, but also to more abstract things such as emotions, or even the self. Such a continuously adapting region of neural state space that is selective for a given object of perception may be said to underly our ability to perceive indefinite variations of that object, and thus to function as what Turvey (1977) called a perception concept. A further distinction with separate action concepts, supposedly underlying indefinite variations of actions, is superfluous on an integrated view of perception and action, and I will simply speak of concepts, underlying perception as well as action.

If this is correct so far, we may now define concepts as the dynamic coupling between a motor plan for a specific action and a set of internal and external stimulus patterns. This coupling consists of informational pathways that adapt based on interaction with the environment. This depiction of concepts fits with theories that propose that conceptual knowledge is grounded in the brain's modality-specific systems (e.g. Barsalou 2003). Extensive evidence exists for this claim (Martin & Caramazza 2003) and as of recent evidence, newly learned conceptual knowledge could be traced back even to primary visual cortex (Garagnani et al. 2021). Such a grounded concept, then, is neither an abstract entity, nor some central representation for other systems to draw on. Instead, it is instantiated every time it is addressed (Feldman-Barret 2017/2018). Concept-specific information is carried by shared networks, therefore concepts are not structurally discernable but should be seen from a dynamical perspective as regions of the neural state space that together form a concept space.

Now that concepts have been defined in dynamical terms, a measure of their informational content can be introduced. Importantly, the neural networks underlying a concept are organized in many hierarchical levels of processing, providing input to each other through synaptic connections. According to information theory, the total amount of information in such a dynamical hierarchical system depends on its possible states, and on the probabilities of the transitions between them. The latter in turn depend on the distribution of the possible states over the processing levels. In a more equal distribution (with all other things kept constant) the probabilities of the possible state transitions are more equal to each other. According to Shannon's formalism, this increases the entropy of the system and the information it can produce.

2.4 P4: Perception requires an increase of entropy in concepts

The informational value of a concept is defined in the previous section by the number and probability distribution of its possible state transitions. Those are the state transitions that will lead to the expression or experience of that concept, i.e. its perception, by activating a state on a subsequent level. This implies that higher concept levels are receptive to specific patterns coming from lower levels. As Sayre (1976/2015, p. 154) puts it, using an analogy with a radio receiver: "[…] the information gets through to the higher level only if the receiver is properly tuned." Indeed, the tuning of neural responses to certain information appears to be a general principle of nervous system organization (O'Donnell & Nolan 2011). However, since the amount of ways any object of perception can present itself in terms of stimulus patterns is extremely large, we may assume that not all possible states are predefined. An instance of perception, then, must entail an increase in the number of possible states and transitions, and therefore of the concept entropy. This comes down to an expansion of the concept region, which can be accomplished by top-down prediction of new stimulus patterns. Because state transitions in a concept as defined here occur stochastically and their probability depends on previous states on lower levels, it can be represented as a Markov process (Figure 1). Here it is assumed that each concept state is in principle equally able to activate all states on the next level, and that the probability of each transition is determined by the number of possible transitions per state and thus by the distribution of states over the processing levels. Top-down expansion of states during perception is assumed to be linear across levels.

Figure 1
Expansion of a concept region

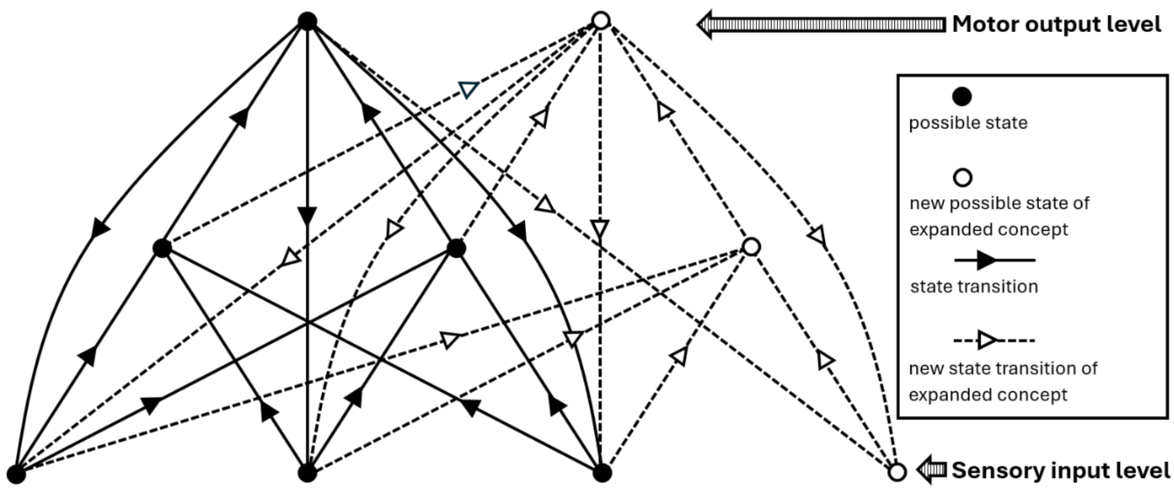

Fig. 1 The state space of a simple example concept (i.e. concept region), organized in three hierarchical levels (here represented as a Markov chain, omitting the probabilities for clarity). Perception of new stimulus information involves the top-down inclusion of predicted new states (open discs) in the concept region. The information of the concept, measured as Shannon entropy (H) depends on the number of possible transitions between states and on the distribution of their probabilities. Note that in this model feedforward transitions from new states to old states do not occur, because the old states do not yet accept the new input.

Contemporary views of perception indeed emphasize the role of prediction in interpreting what is being perceived. Models adhering to a predictive processing scheme essentially describe perception as a continuous attempt to match incoming sensory inputs with top-down expectations or predictions (Clark 2013). On a unified perception-action view, the goal is not to interpret what is perceived, but to determine what behavior to execute. Given the discussion above, this can be stated more broadly as which concept should be activated. This would come down to the increase of the possible states of a selected concept to include predicted stimulus patterns. The extent to which this view fits with existing models of predictive processing requires further scrutiny. Its reliance on the top-down increase of possible states, suggests that it might be particularly suited to a form of backward constraint relaxation, which has recently been proposed to be an important mechanism in biologically plausible learning algorithms (Song et al. 2024; Millidge 2020).

2.5 P5: The increasing entropy in concepts has to be reduced for effective action

Aside from the notable exception of reflexes, our behavior is not innate but results from learning. Learning can be defined as a process in which a response pattern converges to a specific possibility, which can be expressed as minimization of entropy within a given behavioral category (Watanabe 1981). Therefore, action requires a relatively low level of entropy and thus a constrained number of possible states, reflecting the essential invariance of action plans (Turvey 1977).

The notion of stable low-variance action representations is firmly embedded in neurophysiological evidence. Several motor areas, including the primary motor cortex (M1), have been demonstrated to encode the abstract intentions of movements, invariant to motion details (Kadmon Harpaz 2014). A particular category of actions, which is perhaps the most potent in this view on consciousness, consists of vocal actions. The importance of their invariance is obvious for the purpose of minimizing noise in communication. Relatively invariant activation patterns appear indeed to underlie vocal actions, as human speech has been found to be driven by action-specific and sparse functional activations of the mouth motor cortex (Kern et al. 2019).

Previously I have casted action as the behavioral outcome of the transformation of information during perception (P2). Now this outcome has been further specified as response patterns with (relatively) small entropy. However, according to P4, perception also requires an increase of entropy in concepts. It then follows that for a continuous perceptual process to result in effective action, a reduction of entropy must occur. This comes down to condensing of information, and it is in agreement with Gupta and Bahmer (2019, p. 11) who say that "Robust increases in mutual information [reflecting a decrease in entropy] between successive circuits, from sensory areas to premotor areas result in movements that produce successful interaction with the environment", based on studies of neural synchrony during movement. In my framework, this means that during the process of transformation of perceptual information into action, a concept maintains a disparity between the (relatively large) entropy of the stimulus patterns and the (relatively small) entropy of the behavioral output. The size of this disparity is inversely related to the overall entropy of the concept, because it determines the probability distribution of the possible state transitions (see P3).

The model predicts the transition from passive perception to behavior initiation to be accompanied by an entropy decrease of activity patterns within concept-related portions of the motor cortex. Empirical evidence seems indeed to point in that direction, as e.g. Crone et al. (1998) found that movement execution is correlated with somatotopically focused bursts of gamma activity in contralateral sensorimotor cortex. This happens just before and during a movement, while more dispersed and less specific synchronizations in lower frequency bands disappear as soon as movements are executed (Van Wijk 2012). Although the entropy in these signals was not measured, a spatially more localized activity pattern indicates a more constrained number of states, and thus a lower entropy of activity in the motor region. A striking transition between passive perception and active engagement has been elucidated on the level of individual neurons by single cell recordings in cats. Putrino et al. (2010) compared neuronal activity in cat primary motor cortex (M1) during behavior production and quiet sitting. The results show some widespread coincident firing between pairs of neurons in M1 during the latter condition. This pattern of interactions changes remarkably during task performance. Then activity in M1 shows increased synchronization of a more specific set of mostly task-related (and plausibly predominantly inhibitory) neurons. Also widely distributed motor-related activations in humans reveal a lower variance during movement compared to rest (Ottenhoff et al. 2024). These results point at spatially and temporally more restricted activity patterns during action, in accordance with a reduction of the number of possible states of motor representations, and thus a lower entropy.

2.6 Synthesis: inactionable information dissipates and constitutes consciousness

Above, perception has been framed as a continuous transformation of highly fluctuating raw stimulus information into a stable behavioral repertoire, which requires an interplay between opposing informational processes. I have argued that perception requires prediction of new possible concept-related states to allow new information to flow through the concept. This implies a top-down expansion of any concept region that is engaged in perception of new information, and thus an increase of its entropy. Because behavioral outcomes require low variability to be effective in the world, also entropy reduction is needed, which I refer to as the condensing of information. Interaction with the environment thus requires a well-balanced interplay between expansion and condensing of concept regions during perception. While expansion in this model allows for novel stimulus patterns to be predictively coupled to behavior, condensing is required to instigate its execution.

Obviously perception does not always result in action. The framework I have presented so far implies that the outcome of the perception-action transformation is related to its efficiency. Under P3 we have seen that low entropy of a concept depends on a sharp distribution of the possible concept states over its levels. This indicates an efficient perception-action transformation because the involved concept is activated in response to many different variations within a stimulus category (i.e. it is highly sensitive). At the same time its effects have low variability so that the resulting action is consistently distinct from responses to other categories (i.e. it is highly specific). My framework suggests that the efficiency of the perception-action transformation, and thus the concept entropy, would depend at least in part on the directness of the sensorimotor coupling between a creature and its environment. With direct coupling I mean to say that the organism immediately responds to any perturbation of its system. The information from the environment is then fully transformed into corresponding actions. Instead, when the coupling is indirect it involves a choice between several options, and action may be delayed while perception continues. This means that the expanding drive of a concept outweighs the condensing of its information, thereby increasing its informational content. Under such circumstances, an increasing amount of information may have a decreasing effectiveness because responses lack the specificity to have macrophysical effects. A part of the stimulus input to the concept then dissipates as inactionable states. As mentioned in the introduction, the use of the term dissipation in the context of information processing resembles its use in thermodynamics, where it refers to an increase of system entropy indicating a decrease of the usefulness of the energy in the system.

Since under P1 I suggested inactionable perception as constitutive of consciousness, the portion of the informational flow that is not transformed into action during perception fits the bill to be identified as conscious experience. This information, although inactionable, still exists within the confines of a concept region. If this is correct, the implication would be that when we consciously experience something, this reflects the availability of a corresponding action in our repertoire, that we may produce if we sufficiently reduce entropy within the appropriate concept. This would refute an intuitive view of consciousness as a reflection of the world as it presents itself. Instead consciousness may be more aptly described as potential behavior that has not manifested. Consciousness and behavior, then, may be best seen not as subsequent stages of processing, but as different outcomes of the same perceptual transformation process depending on its efficiency.

3 Discussion

The link between consciousness and entropy of brain activity has been established in a growing body of empirical work. This has converged on a common theme: states of consciousness are associated with the entropy of brain-derived signals, where an increased level of consciousness – as in wakeful or even psychedelic states – correlates with higher entropy values. This is in accordance with my proposal, which gives a new theoretical account of how and why consciousness is related to the entropy of brain activity. Contrary to other entropy-centered accounts, this is based on an operationalization of consciousness that does not depend on measures of cognitive capacities or on a quantitative measure of the contents of consciousness. Instead, it presents a model of perception-action in which an increase of entropy reflects the qualitative difference between conscious and unconscious processing.

This proposal places perception and action in a unified entropy-based framework, in which raw stimulus information is transformed gradually into behavior. Featuring prominently in this transformation are two opposing entropic processes, postulated to operate throughout the levels of the neural hierarchy. A top-down mechanism increases possible states and thereby the entropy of a concept, driven by prediction of novel stimulus patterns. The opposing mechanism reduces entropy to produce stable and accurate low-variance output activations that have macrophysical effects on the environment in the form of behavior. This framework is combined with a definition of consciousness as inactionable perception. Consciousness, then, emerges as the result of inefficient perception-action transformation within concepts, resulting in inactionable states: efference patterns that have no macrophysical effects.

Clearly, many possible implications and underlying assumptions, both philosophical and neurophysiological, have yet to be worked out and empirically verified. Below, I propose some first considerations.

3.1 Philosophical considerations

If corroborated, this proposal might provide some much-needed constraints on the problem of consciousness. It may not directly elucidate how or why conscious experience has the phenomenal qualities that it has, but it might at least improve our understanding of what it is and on what physical processes it depends. At the core of the proposal is the claim that conscious experience consists of perception that does not lead to behavior, i.e. it cannot have direct macrophysical effects on the world. This is not a premise, but it is based on arguments centered around the non-conceptuality of conscious experience. In this sense, it lends credence to the idea that consciousness is an epiphenomenon (see e.g. Kim 1998/2021), but it does so without assuming that it is non-physical and by stipulating that its inefficacy is restricted to macrophysical effects. As such it is not susceptible to the concerns of Mathieson (2024), who warns for too readily an acceptance of epiphenomenalism, rooted in unwarranted assumptions of non-physicality and non-functionality of phenomenal consciousness. Notably, my proposed operationalization of consciousness seems to lead to a kind of reversal of the mind-body problem. Instead of the inconsistency of a putative non-physical mind having behavioral effects (Westphal 2016), the interface between mind and body now involves (seemingly less problematically) a physical process of perception that has (under certain circumstances) no macrophysical effects.

Of course, the idea of mental causation is highly intuitive and is not easily abandoned. My proposal provides a clear suggestion as to why this intuition is so pervasive. Indeed, when we engage in behavior that we would typically think of as conscious, we always first have an experience of its possibility. But rather than endowing this experience with causal agency, this proposal explains how experience is the potentiality of behavior, which merely precedes the actual behavior that we would typically call conscious. A corollary is that a distinction between access and phenomenal consciousness (Block 1995) appears illusory. My proposal accounts for Block's idea of access consciousness as the availability of its content for utilization in reasoning or action, as well as his notion of phenomenal consciousness as subjective, ineffable experience. But it also implies that there would be no way in which these differ from each other, since they both are explained as the potentiality of behavior. This seems to be contradicted by Amir et al. (2023), who claim to have found evidence for phenomenal consciousness without access consciousness by demonstrating retrospective experience of a previously unreportable sound. My model might have difficulty explaining such a dissociation of consciousness types. However, it leaves room for an alternative explanation in which a subthreshold stimulus leaves a memory trace and only reaches the threshold for reportability when it is retrieved.

My framework suggests that Block's division might be a reflection of the qualitatively different results of neural information processing as either conscious perception or action, depending on the efficiency of the process. Such dichotomy of modes of information processing resonates with Chalmer's (1996) conclusion that "Experience arises by virtue of its status as one aspect of information". In accordance with the idea that consciousness is not to be observed or apprehended (Dennett 1991/2007), my model suggests that the experiential aspect of information, i.e. the dissipation of information, is itself the conscious experience. And if the dissipation of information is a potential action, then conscious experience is the potentiality of the action itself.

Yet, if information dissipation indeed turns out to be Chalmer's "extra ingredient" in the explanation of consciousness, this in itself is still not a satisfactory answer to the hard problem of why we have experiences at all (Chalmers 1996). We may however have usefully reframed the problem as the question of "Why does dissipation of information in concepts during perception give rise to experience?" I believe this to be a much better specified question. Moreover, the present proposal may provide at least a partial mechanism for how consciousness can arise from perception. If indeed this can be understood in terms of information and entropy, it may ultimately provide a new perspective on the relation of consciousness and the physical realm.

3.2 Theoretical considerations

More than the entropy-based accounts mentioned in the introduction, my proposal seems to have much in common with the sensorimotor model of consciousness. Both share an emphasis on the interconnectedness of action and perception and make claims to the effect of perception being determined by what we are ready to do (Noë 2004). However, there are important differences that are interesting for this discussion, especially because some of the evidence employed in support of the sensorimotor model (SMM) may be better explained by the current proposal. SMM claims that consciousness arises from dynamic patterns of sensorimotor contingency (Hurley & Noe 2003), which are defined as the regularities in how sensory stimulation depends on the activity of the perceiver (Wilson & Foglia 2018). According to SMM, neural activity is only one part of the extended sensorimotor dynamics, and therefore is not sufficient for conscious experience. But, as Clark (p. 218) points out, in reference to Block (2005 in Clark 2014), "nothing in the evidence makes a case for [this] claim." Like SMM, my proposal states that consciousness requires knowledge of regularities in the interactions between the organism and its environment, but only as far as they are embodied in the neural circuitry sustaining concepts. What is proposed here to be constitutive of consciousness is the transformation of stimulus information into action, and specifically a lack thereof. This transformation depends on the informational pathways that have been installed through training and tuning the system during interaction with the environment. The degree to which this transformation leads to condensed action information is then what determines whether it gives rise to consciousness or not. The point where the information may become irreversibly inactionable is at the interface with the motor plant, where perceptual information arguably reaches its most invariant form (after which instantaneous, online adjustments based on visual and proprioceptive feedback may be implemented by subcortical and spinal processes; Kasuga et al. 2022).

This suggests that consciousness arises specifically from the output of the neural system in its efference signals. The action that could be potentially instigated if condensing of information were more efficient, then, determines the quality of the experience. In contrast, SMM distributes consciousness over the extended sensorimotor loops, which would make conscious experience very sensitive to the idiosyncratic details of human embodiment (Clark 2014). Interestingly, experiments with optical rearrangement devices (ORDs) which are usually considered to support SMM, seem more suggestive of path-independence of conscious perception. In such experiments subjects are wearing prisms that distort their visual field in a systematic way. Results show that subjects, despite the optical shift, can reach normal proficiency in spatial tasks, after a period of adaptation involving active engagement with the environment. This is accompanied by normalization of their qualitative, subjective experience (e.g Taylor 1962, in Hurley and Noë 2003). What is thought to happen in these cases is that during active engagement with the environment, the prism-wearing subject acquires new sensorimotor contingencies. Sensorimotor relations become predictable again, which assumedly gradually normalizes the visual experience. However, during this adaptation, sensorimotor contingencies clearly do not return to what they were before the device was applied, because visual input remains different compared to before application of the device. So, it is not clear why visual experience should become *the same* as before, since what the subject adapts to are new, different contingencies. This discrepancy is avoided in my proposal, which can reconcile altered sensory input after ORD-adaptation with familiar experience by suggesting that the experience is determined by potential action which depends on efference patterns. If this is correct, then it is not surprising that even thoroughly distorted vision yields a familiar conscious spatial experience after adaptation. Efference patterns driving effective (potential) actions, after all, are determined by the actual relations with the environment, which have not changed.

My proposal does not deny that action and perception form a continuous loop arcing through the non-neural body and the environment. This is for instance also a central tenet of the ecological approach of perception (Gibson 1986/2015), which stresses that perception should be studied during movement, and can not be seen as a series of snapshots. My proposal is in agreement with this view on perception as intimately related to behavior. However, it suggests that conscious experience represents a mode of perception where the loop of action and perception is not continuous and closed but discontinuous and open, and more information is received than is used for action. In complex creatures such as ourselves, where usually no a priori relation exists between sensory input and the required action to achieve a goal, this state of informational dissipation coincides with situations in which a decision can be made about if and how we should respond to current stimulation. This resonates with the intuition that a decision out of free will always coincides with some level of conscious experience of the situation and of our possible responses.

3.3 Empirical considerations

My proposal makes several claims that may allow empirical testing. In the proposed framework, concept expansion is counteracted by an opposing condensing mechanism required for action. When condensing falls short, information dissipation is claimed to result in inactionable states and to give rise to consciousness. The extreme implication of this seems to be that conscious experience of stimulus information should disappear when its associated action is produced. However, this would require a complete transformation of stimulus information into condensed action information, while everyday reasonably effective transformations may more likely produce a mere relative drop in consciousness. Indeed, from informal observation it seems that when actions are tightly coupled to sensory input, such as while keeping one's balance when riding a bike or playing fast-paced sports, we seem less conscious of the sensorimotor interactions involved. This inverse relationship between action production and consciousness has experimental support, as skill performance appears to deteriorate when people try to exert conscious control over automatic actions (Toner & Moran 2011).

It may be possible to get more direct measurements of the predicted drop in conscious experience upon condensing of stimulus information. In the attentional blink (AB) paradigm, stimuli are presented in rapid succession and subjects are instructed to report target stimuli among a stream of distractors. Typically, accuracy in reporting a second target is negatively affected by a previous target shown shortly before (the AB effect). The underlying mechanism is subject of debate, but the paradigm may be useful to tap into the informational dynamics predicted by this proposal. For example Tang et al. (2022) show that the AB effect is absent when both targets have the same representation. According to my model, this is to be expected if information about the first target stimulus is still in relatively condensed state (because of covert activation of its concept) at presentation of the second identical target, facilitating the associated action. Other AB experiments have measured not only behavioral responses but also ratings of subjective visibility. Pincham et al. (2016) showed that the deficiency in reporting and in subjective experiencing of the second (different) target depends differentially on the duration of the inter-target lag. When targets appear in very close temporal succession they found relatively low reporting deficiency, but high deficiency in subjective visibility of the second target. The authors have dubbed this effect the experiential blink. In light of the present proposal, it would be interesting to see if increasing overlap in perceptual features between targets would lead to an increase of behavioral reporting accuracy with a decrease of subjective experience on certain inter-target lags. This might lend credence to the idea that a drop in consciousness (experiential blink) may result from the relevant information being in a condensed action-ready state.

The claim that conscious perception reflects potential action may lead to other testable hypotheses. Support for this claim is already widely available, coming from research showing that perception is action-specific. Many studies have for example demonstrated that perceived spatial properties of an object depend on the perceiver's ability to perform an action on the object, such as reaching, grasping or climbing (Witt & Riley 2014). Such effects could also be hypothesized to be obtained by non-spatial, language related behavior. Whether this can be confirmed remains an open question, but it is clear that language influences perception at several levels (Vulchanova et al. 2019).

In general, in this proposal potential action (as opposed to actual action) pertains to efference patterns with insufficiently condensed information. Consciousness, then, should relate to their high-entropy, non-condensed state, a claim that may be susceptible to scrutiny by neuroimaging experiments. Covert motor activity has indeed been found during passive perception (Bundt et al. 2015). Motor activity has also been shown to influence subjective ratings of visual experience, as TMS-induced M1 activation can improve perceptual awareness (Hobot et al. 2020).

The above-mentioned experimental results are not exclusively explained by the present proposal. Moreover, without a doubt other empirical findings may support, or instead pose a challenge for this proposal. A further exploration of its implications, together with a more thorough review of the literature, and the execution of specifically designed experiments, will be needed to provide more conclusive evidence.

4 Conclusion

The ideas presented in this paper cast the inactionability of conscious experience as its fundamental property. By placing this in a dynamical, informational framework of perception and action, consciousness emerges not as a reflection of what happens to us, but as an array of behavioral potentialities, structured in concepts consisting of sensorimotor couplings, embodied in neural networks. A picture emerges of consciousness as efferent activity patterns that are not specific enough to have macrophysical effects on the world. Since these patterns are part of the range of possible patterns of action-based concepts, they constitute potential actions. The model implies that this is a result of inefficient transformation of stimulus information into action, resulting from flexible coupling between organism and environment. Although negating the macrophysical efficacy of consciousness may seem counterintuitive, the proposal predicts conscious experience to occur under the same conditions as would be expected on a causally effective view of consciousness. Those conditions obtain when we engage in behavior that we think of as deliberate, because it seems not predetermined but instead requires a choice.

I believe this new view offers promising directions for further exploration of ways to embed consciousness in the framework of physical law. Its theoretical implications should be further examined and tested against existing literature and carefully designed experiments. Only this will tell if the proposal may develop into a viable theory of consciousness.

References

Amir, Yoni, Z., Yaniv Assaf, Yossi Yovel, Liad Mudrik. 2023. Experiencing without knowing? Empirical evidence for phenomenal consciousness without access. *Cognition*. https://doi.org/10.1016/j.cognition.2023.105529

Attick, Joseph. 2011. Could information theory provide an ecological theory of sensory processing? *Network*. https://doi.org/10.3109/0954898X.2011.638888

Barsalou, Lawrence W., W. Kyle Simmons, Aron K. Barbey, and Christine D. Wilson. 2003. Grounding conceptual knowledge in modality-specific systems. *Trends in cognitive sciences*. https://doi.org/10.1016/s1364-6613(02)00029-3

Barsalou, Lawrence W. 2012. The human conceptual system. In *The cambridge handbook of psycholinguistics*, eds. Michael J. Spivey, Ken McRae, and Marc F. Joanisse, 239-258. Cambridge University Press. https://doi.org/10.1017/CBO9781139029377

Beshkar, Majid. 2023. The QBIT theory of consciousness: Entropy and qualia. *Integrative Psychological and Behavioral Science*. https://doi.org/10.1007/s12124-022-09684-6

Block, Ned. 1995. On a confusion about a function of consciousness. *Behavioral and brain sciences*. https://doi.org/10.1017/S0140525X00038188

Bundt, Carsten, Lara Bardi, Elger L. Abrahamse, Marcel Brass, Wim Notebaert. 2015. It wasn't me! Motor activation from irrelevant spatial information in the absence of a response. *Frontiers in Human Neuroscience*. https://doi.org/10.3389/fnhum.2015.00539

Carhart-Harris, Robin L. 2018. The entropic brain: Revisited, *Neuropharmacology*. https://doi.org/10.1016/j.neuropharm.2018.03.010

Castro, Pablo, Andrea Luppi, Enzo Tagliazucchi, Yonatan S. Perl, Lorina Naci, Adrian M. Owen, Jacobo D. Sitt, Alain Destexhe, and Rodrigo Cofré. 2024. Dynamical structure-function correlations provide robust and generalizable signatures of consciousness in humans *Communications Biology*. https://doi.org/10.1038/s42003-024-06858-3

Chalmers, David J. 1995. Facing up to the problem of consciousness. *Journal of consciousness studies* 2(3): 200-19, 1995.

Clark, Andy, Josefa Toribio. 1994. Doing without representing? *Synthese*. https://doi.org/10.1007/BF01063896

Clark, Andy. 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*. https://doi.org/10.1017/S0140525X12000477

Clark, Andy. 2014. *Mindware: An introduction to the philosophy of cognitive science. 2nd ed.* New York: Oxford University Press.

Crone, Nathan E., Diana L. Miglioretti, Brian A. Gordon, and Ronald P. Lesser. 1998. Functional mapping of human sensorimotor cortex with electrocorticographic spectral analysis. II. Event-related synchronization in the gamma band. *Brain*. https://doi.org/10.1093/brain/121.12.2301

Dennett, Daniel C. 2007. *Het bewustzijn verklaard* (Ton Maas & Frits Smeets, trans.). Amsterdam: Olympus. (Original work published 1991)

Dennett, Daniel C. 2021. Quining qualia. In *Philosophy of mind: Classical and contemporary readings* (2nd ed.), ed. David J. Chalmers, 199 - 219. New York: Oxford University Press. (Reprinted from *Consciousness in contemporary science*, eds. Anthony E. Marcel, and Edoardo Bisiach, 1988, New York: Oxford University Press)

Edelman, Shimon. 2008. *Computing the mind: How the mind really works*. New York: Oxford University Press.

Erra, Ramon G., Diego M. Mateos, Richard Wennberg, and Jose L. Perez Velazquez. 2017. Towards a statistical mechanics of consciousness: Maximization of number of connections is associated with conscious awareness. Preprint at arXiv. https://doi.org/10.48550/arXiv.1606.00821

Farrell, Jonathan. 2016. 'What it is Like' Talk is not Technical Talk. *Journal of consciousness Studies* 23(9-10): 50-65.

Feldman Barrett, Lisa. 2018. *How emotions are made: The secret life of the brain*. London: Pan Books. (Original work published 2017)

Fuster, Joaquín M. 2003. *Cortex and mind: unifying cognition*. New York: Oxford University Press.

Garagnani, Max, Evgeniya Kirilina, Friedemann Pulvermüller. 2021. Semantic grounding of novel spoken words in the primary visual cortex. *Frontiers in human neuroscience*. https://doi.org/10.3389/fnhum.2021.581847

Gibson, James J. 2015. *The ecological approach to visual perception*. New York: Psychology Press. (Original work published 1986)

Gupta, Daya S. & Andreas Bahmer. 2019. Increase in mutual information during interaction with the environment contributes to perception. *Entropy*. https://doi.org/10.3390/e21040365

Halligan, Peter W. & David A. Oakley. 2021. Giving up on consciousness as the ghost in the machine. *Frontiers in psychology.* https://doi.org/10.3389/fpsyg.2021.571460

Hurley, Susan. 2001. Perception and action: alternative views. *Synthese*. https://doi.org/10.1023/A:1012643006930

Hurley, Susan, & Alva Noë. 2003. Neural plasticity and consciousness. *Biology and Philosophy* 18: 131-168.

Hobot, Justyna, Marcin Koculak, Borysław Paulewicz, Kristian Sandberg, and Michał Wierzchoń. 2020. Transcranial magnetic stimulation-induced motor cortex activity influences visual awareness judgments. *Frontiers in neuroscience*. https://doi.org/10.3389/fnins.2020.580712

Jackson, Frank. 2021. Epiphenomenal qualia. In *Philosophy of mind: Classical and contemporary readings* (2nd ed.), ed. David J. Chalmers, 283 - 290. New York: Oxford University Press. (Reprinted from Epiphenomenal qualia, 1982, *Philosophical Quarterly* 32: 127-36. Addendum excerpted from What Mary didn't know, 1986, *Journal of Philosophy*, 83: 291-95).

Jaynes, Edwin, T. 1957. Information theory and statistical mechanics. II. *The Physical Review* 108(2): 171-190.

Ji, Xu, Eric Elmoznino, George Deane, Axel Constant, Guillaume Dumas, Guillaume Lajoie, Jonathan Simon, and Yoshua Bengio. 2024. Sources of richness and ineffability for phenomenally conscious states. *Neuroscience of consciousness*. https://doi.org/10.1093/nc/niae001

Kadmon Harpaz, Naama, Tamar Flash, and Ilan Dinstein. 2014. Scale-invariant movement encoding in the human motor system. *Neuron*. https://doi.org/10.1016/j.neuron.2013.10.058

Kasuga, Shoko, Frédéric Crevecoeur, Kevin P. Cross, Parsa Balalaie, and Stephen H. Scott. 2022. Integration of proprioceptive and visual feedback during online control of reaching. *Journal of neurophysiology*. https://doi.org/10.1152/jn.00639.2020

Kim, Jaegwon. 2021. The many problems of mental causation. In *Philosophy of mind: Classical and contemporary readings* (2nd ed.), ed. David J. Chalmers, pp. 137 - 146. New York: Oxford University Press. (Reprinted from *Mind in a physical world: An essay on the mind-body problem and mental causation*, pp. 29-47, 1998, Cambridge: MIT Press)

Keshmiri, Soheil. 2020. Entropy and the Brain: An Overview. *Entropy*. https://doi.org/10.3390/e22090917

Kern, Markus, Sina Bert, Olga Glanz, Andreas Schulze-Bonhage, & Tonio Ball. 2019. Human motor cortex relies on sparse and action-specific activation during laughing, smiling and speech production. *Communications biology*. https://doi.org/10.1038/s42003-019-0360-3

Martin, Alex, and Alfonso Caramazza. 2003. Neuropsychological and neuroimaging perspectives on conceptual knowledge: an introduction. *Cognitive neuropsychology*. https://doi.org/10.1080/02643290342000050

Mathieson, Darryl. 2024. Consciousness, causation, and confusion. *Review of Philosophy and Psychology*. https://doi.org/10.1007/s13164-024-00762-9

Millidge, Beren, Alexander Tschantz, Anil K. Seth, and Christopher L. Buckley. 2020. Activation relaxation: A local dynamical approximation to backpropagation in the brain. Preprint at arXiv. https://doi.org/10.48550/arXiv.2009.05359

Nagel, Thomas. 1974. What is it like to be a bat? *The Philosophical Review*. https://doi.org/10.2307/2183914

Noë, Alva. 2004. *Action in Perception*. Cambridge, London: MIT Press.

O'Donnell, Cian, and Matthew F. Nolan. 2011. Tuning of synaptic responses: an organizing principle for optimization of neural circuits. *Trends in neurosciences*. https://doi.org/10.1016/j.tins.2010.10.003

Ottenhoff, Maarten C., Maxime Verwoert, Sophocles Goulis, Louis Wagner, Johannes P. van Dijk, Pieter L. Kubben, and Christian Herff. 2024. Global motor dynamics: Invariant neural representations of motor behavior in distributed brain-wide recordings. *Journal of neural engineering*. https://doi.org/10.1088/1741-2552/ad851c

Pakan, Janelle MP, Valerio Francioni, and Nathalie L. Rochefort. 2018. Action and learning shape the activity of neuronal circuits in the visual cortex. *Current Opinion in Neurobiology*. https://doi.org/10.1016/j.conb.2018.04.020

Pincham, Hannah L., Howard Bowman, and Denes Szucs. 2016. The experiential blink: Mapping the cost of working memory encoding onto conscious perception in the attentional blink. *Cortex*. https://doi.org/10.1016/j.cortex.2016.04.007.

Putrino, David, Emery N. Brown, Frank L. Mastaglia, and Soumya Ghosh. 2010. Differential involvement of excitatory and inhibitory neurons of cat motor cortex in coincident spike activity related to behavioral context. *The journal of neuroscience*. https://doi.org/10.1523/JNEUROSCI.0770-10.2010

Rankaduwa, Sidath, and Adrian M. Owen. 2023. Psychedelics, entropic brain theory, and the taxonomy of conscious states: a summary of debates and perspectives. *Neuroscience of consciousness*. https://doi.org/10.1093/nc/niad001

Roelfsema, Pieter, and Anthony Holtmaat. 2018. Control of synaptic plasticity in deep cortical networks. *Nature Reviews Neuroscience*. https://doi.org/10.1038/nrn.2018.6

Sayre, Kenneth, M. (2015). *Cybernetics and the philosophy of mind*. Oxford, New York: Routledge. (Original work published in 1976)

Shannon, Claude E. 1948. A mathematical theory of communication. *Bell system technical journal*. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Shannon, Claude E., and Warren Weaver. 1998. *The mathematical theory of communication*. Urbana, Chicago: University of Illinois Press. (Original work published in 1949)

Shieber, Joseph. 2010. On the possibility of conceptually structured experience: demonstrative concepts and fineness of grain. *Inquiry*. https://doi.org/10.1080/0020174X.2010.493371

Song, Yuhang, Beren Millidge, Tommaso Salvatori, Thomas Lukasiewicz, Zhenghua Xu, and Rafal Bogacz. 2024. Inferring neural activity before plasticity as a foundation for learning beyond backpropagation. *Nature Neuroscience*. https://doi.org/10.1038/s41593-023-01514-1

Sturgeon, Scott (1998). Physicalism and overdetermination. *Mind* 107(426):411-432.

Tang, Matthew F., Kimron L. Shapiro, James T. Enns, Troy A.W. Visser, Jason B. Mattingley, and Ehsan Arabzadeh. 2022. Visual awareness during the attentional blink is determined by representational similarity. Preprint at bioRxiv. https://doi.org/10.1101/2022.10.25.513789

Toner, John, & Aidan Moran. (2011). The effects of conscious processing on golf putting proficiency and kinematics. *Journal of sports sciences*. https://doi.org/10.1080/02640414.2011.553964

Turvey, Michael T. 1977. Preliminaries to a theory of action with reference to vision. In *Perceiving, Acting and Knowing: Toward an Ecological Psychology*, eds. Robert E., Shaw, and John D. Bransford, 211-267. London: Routledge. https://doi.org/10.4324/9781315467931

Van Wijk, Bernadette C. M., Peter J. Beek, and Andreas Daffertshofer. 2012. Neural synchrony within the motor system: what have we learned so far? *Frontiers in human neuroscience.* https://doi.org/10.3389/fnhum.2012.00252

Vulchanova, Mila, Valentin Vulchanov, Isabella Fritz, and Evelyn A. Milburn. 2019. Language and perception: Introduction to the special Issue 'Speakers and Listeners in the Visual World'. *Journal of cultural cognitive science*. https://doi.org/10.1007/s41809-019-00047-z

Watzlawick, Paul, Janet Beavin Bavelas, and Don D. Jackson. 2011. *Pragmatics of human communication: A study of interactional patterns, pathologies, and paradoxes*. New York, London: W.W. Norton & Company. (Original work published 1967)

Watanabe, Satosi. 1981. Pattern recognition as a quest for minimum entropy. *Pattern recognition* 13(5): 381-387.

Weilenmann, Mirjam, Lea Kraemer, Philippe Faist, and Renato Renner. 2016. Axiomatic relation between thermodynamic and information-theoretic entropies. *Physical review letters*. https://doi.org/10.1103/PhysRevLett.117.260601

Westphal, Jonathan. 2016. *The mind-body problem*. Cambridge, London: MIT Press.

Wilson, Margaret. 2002. Six views of embodied cognition, *Psychonomic bulletin & review*. https://doi.org/10.3758/BF03196322

Wilson, Robert A., and Lucia Foglia. 2018. Embodied cognition. The Stanford encyclopedia of philosophy (Winter 2018 edition), ed. Edward N. Zalta. https://plato.stanford.edu/archives/win2018/entries/embodied-cognition

Witt, Jessica K., Michael A. Riley. 2014. Discovering your inner Gibson: Reconciling action-specific and ecological approaches to perception–action. *Psychonomic bulletin & review*. https://doi.org/10.3758/s13423-014-0623-4

Wright, Wayne (2003). McDowell, demonstrative concepts, and nonconceptual representational content. *Disputation* 14 (14): 1-16.