A Predictive and Illusionist perspective of human subjectivity[1]

Maria Luiza Iennaco[2]

Abstract:
The main philosophical perspectives on human subjectivity uphold a series of properties that make it scientifically intractable. Given the importance of understanding subjectiveness in such terms, this chapter aims to propose an alternative perspective that considers not only the phenomenology of subjective experiences but also the underlying processes and mechanisms. We will start from Sellarsian worldviews and the illusionist philosophical current to explain that all our perceptual experiences are oriented towards affordances. The processes and mechanisms involved in these experiences will be developed based on Predictive Processing theory. This theory will be used to explain not only external perception but also internal perception, which is key to understanding subjectivity. We will see that predictive subjectivity has some conceptual limitations, and we will seek to address them with Illusionism.

Keywords: subjectivity; predictive processing; interoception; illusionism; philosophy of cognitive sciences

Content:

1. Introduction... 1
1.1 The Intelectual Jungle.. 2
1.2 The Intelectual in the Jungle.. 4
2. Predictive Processing..7
2.1 Conditional and Counterfactual Epistemic Actions...10
2.2 Instrumental Actions and Interoceptive Inferences.. 12
3. I Predict, Therefore I Am..15
3.1 Hallucinating a Top Hat with Rabbits Inside.. 17
3.2 The Illusionist in the Jungle... 18
References...21

---

[1] This text was originally published in Portuguese: https://anpof.org.br/wlib/arqs/publicacoes/96.pdf .
DOI: https://doi.org/10.58942/eqs.95.02

[2] PhD student in Philosophy of Cognitive Sciences at the University of São Paulo, Brazil. ORCID: https://orcid.org/0000-0002-5407-4852

1. Introduction

In general terms, a subjective experience refers to an individual's perception of their world and themselves. Such perception is composed of their emotions, thoughts, opinions, and interpretations of internal and external information, situated in the context in which the organism finds itself. In fields such as philosophy and psychology, the study of subjective experiences has been crucial for understanding the various facets of an individual. Thus, these subjective experiences, or simply "subjectivities," have become central topics of study in the sciences and philosophy of consciousness. Indeed, one of the most widely accepted definitions of consciousness in analytical philosophy states that for an organism to be conscious, there must be something it is like to be that organism (Nagel, 1974). That is, subjectivity, in this context, is the core of consciousness. Consequently, perspectives that emphasize this centrality developed and became popular at the end of the last century, treating subjectivity as the qualitative aspect of consciousness, or simply qualia (quale, in the singular) (Carruthers, 2020a).

In this context, the subjectivity of consciousness would be experienced through sensations of what-it-is-like for an individual to perceive something (Frankish, 2012). These sensations are generally characterized as intrinsic, ineffable, instinctive, and directly apprehended in consciousness (Tye, 2021). Thus, subjectivity seems to be a phenomenon difficult to capture with scientific tools. It is something private to the one who has it and, for some reason, it is experienced in a very intense, real, and spontaneous way, even though it does not seem to be completely and clearly possible to talk about it (Carruthers, 2020a). One of the reasons for the scientific intractability of subjectivity consists in its various relatively consensual definitions and descriptions in philosophy and psychology which, unfortunately, are ambiguous, circular, and often non-operationalizable (i.e., not scientifically treatable) (Frankish, 2012).

Although this is not necessarily a problem for philosophy, it certainly is for other disciplines, since science is the main provider of knowledge and means to deal with natural phenomena and their fluctuations. Enabling the description of subjectivity in scientific terms, then, seems essential for dealing with mental disorders and disturbances, social and cultural issues, and also for our relationship with the world we live in, in the sense of understanding the limitations and advantages of our perspective towards the external environment and other beings that inhabit it. With this in mind, this chapter aims to propose an operationalizable way to explore subjectivity that takes into account not only the phenomenology of subjective experiences but also the underlying processes and mechanisms.

In this first section, we will talk about some biases that make the studies of subjectivity overly intellectualist and discuss why this characteristic is limiting and obsolete. Then we will present other assumptions, based on evolutionary and ecological theories of cognition, that avoid the intellectualist treatment of subjective experience and emphasize its importance for interaction with and survival in the world[3]. In the second section, we will develop the framework of Predictive Processing for the functioning of the brain and body, properly utilizing the theory of Active Inference. Predictive Processing and Active Inference are theoretical constructs based on physics, originating from computer science and cybernetics, which address cognition, action, and mental phenomena through a single process of error minimization.

As we will see, the interaction of various predictive models allows organisms like us to anticipate the effects that will be captured by their sensors and prepare for these effects, in order to maintain optimized body functioning and physiological regulation, which ensures their survival. Regarding physiological regulation, we will see a branch of these frameworks called “Interoceptive Inferences” that tells us a lot about subjectivity in predictive terms. In the third and final section, we will present a particular understanding of subjectivity, called “Controlled Hallucinations.” It was proposed by neuroscientist Anil Seth in 2021 and is the result of his work with Interoceptive Inferences and his philosophical reflections. Furthermore, although it is compatible with Predictive Processing, we will see that the idea of Controlled Hallucinations carries some problems that we will try to solve with the help of the illusionist philosophical current, proposed by Keith Frankish in 2017. Thus, we will conclude with a proposed perspective on subjective experience that avoids intellectualism and scientific intractability and provides possibilities for exploration and interdisciplinary work.

1.1 The Intelectual Jungle

Subjectivity, as described by philosopher Thomas Nagel (1974) in his text “What is it like to be a bat?”, consists of a unique experience of being a conscious organism. For example, the experience of a visual sensation, the experience of darkness and light, the sound of a flute, the smell of coffee, the emotion of fear, etc., are all united by the unique sensation of what it is like to be an organism, thus consisting of states of experience (Chalmers, 1995, p. 3). These subjective states of experience became particularly difficult to study after philosopher David Chalmers segmented the “Easy Problems” from the “Hard Problem” in the 1990s. The Hard Problem questions why our physical cognitive processes would generate such a subjective experience, whose forms through which we become aware of its existence are

[3] Notably, the “world” in question, instead of being the “planet Earth,” will be limited to the “niche,” that is, the portion of the planet Earth with which the organism interacts for most of its life. However, we will use the words “world” and “niche” as synonyms.

only first-person reports and introspection. As an illustration, Chalmers explains that we can empirically understand and analyze the mechanisms that allow us to see the redness of a tomato (i.e., an Easy Problem). However, the conscious experience, that is, the phenomenal sensation of experiencing the redness of the tomato, would not be within the scope of current sciences, and even if the computational processes by which we see such redness are meticulously explained, the subjective qualities involved would still lack elucidation.

This type of problem regarding the mind-body relationship, and the possibility that the mind is something irreducible to the physical and inaccessible to science, has a long history in philosophy and human cultures. Peter Carruthers (2020a, p. 201-203) argues that this idea exerts a strong intuitive appeal and creates several obstacles to advancing knowledge about the nature of the mind. A classic example is Cartesian dualism, which asserts that the mind is an immaterial substance independent of the body, aligning with many religious conceptions. According to this perspective, the mind would be the essential and distinctive element of humanity, while the body would be secondary and irrelevant (Carruthers, 2020a, p. 199).

The term “essence” usually refers to a type of internal core in each thing, which causes a manifestation of specific physical properties, as well as possible mental and/or intentional properties (such as desires, preferences, etc.) (Carruthers, 2020b, p. 231). In psychology, essentialism consists of an intuition that all living beings possess something like an immutable and innate essence, shared and maintained similarly by each individual of the same species. Furthermore, similarities between species could be determined by an “underlying nature,” which could not be directly observed (broadly speaking, there seems to be, ultimately, some property of a cat, for example, that makes it different from a lion and a jaguar, but at the same time, more similar to them than to a snake – that is, its essence) (White, 2021, p. 117). Essentialism is an intuition that develops during human childhood (see Gelman, Wellman, 1991; Gelman, 2003), and is therefore widely spread and well-established.

Assuming that there is a distinction between body and mind and that the mind is or constitutes the essence of human beings and existence, we arrive at a final differentiation that seems to complete this picture of human subjectivity: the “self” as opposed to others around us. An organism must distinguish itself from its environment, so that it can know, for example, what is food and what is not, where it can go, what temperature it can withstand, etc. Interestingly, we tend to try to infer or predict what happens in the minds of others. This is called the theory of mind in psychology and philosophy and refers to the ability we have to attribute (to ourselves and others) mental states – an essential ability for the success of human interactions (Carruthers, 2020a, p. 195-196). Thus, the way human subjectivity has been constructed has strongly influenced how we reflect on our individual experiences and, consequently, how we relate to the world.[4]

Dennett (1991) argues, however, that this framework of human subjectivity unnecessarily fosters an epistemic gap between "scientifically explainable processes" and "phenomenal subjectivity," the latter being something irreducible and immutable. The problem with this dichotomy is that subjective experiences become increasingly rooted in intellectualist ideals, that is, ideals that require a high level of reflection. Thus, only beings with reflective capacities would be able to access the subjectivity of their experiences (but only their own, since they are private and unique to each individual). This ends up distancing subjectivity from its own corporeal and non-reflective nature and limiting it to a portion of existing species (perhaps even only humans), in addition to making any scientific exploration impossible. And, as mentioned, science is an indispensable tool for understanding mental phenomena and body dynamics, so we need to find ways to study subjectivity through it.

1.2 The Intelectual in the Jungle

As we have seen, there are good reasons to adopt new ways of thinking about subjectivity, ways that consider scientific findings and foster interdisciplinarity. Here we will attempt to outline one of these ways, by first putting subjectivity into perspective, emphasizing the difference between and coexistence of "appearance" and "reality" or, as proposed by Wilfrid Sellars (1912-1989), the separation of "individuals' perspectives in the world" into two images (1963, p. 5).

On one side, the manifest image is characterized by the "encounter of someone with themselves" and consists of an empirical and categorical refinement of a context throughout an individual's life. The empirical refinement would come from the very experiences that individuals live in their world and the beliefs they hold about these experiences. In turn, the categorical refinement would operate on the processes of detailing what is perceived and lived in the world, that is, through the perception of what objects do and are – which would lead to progressive categorization throughout an individual's life (Sellars, 1963, p. 11).

On the other hand, the scientific image is characterized as an optimized idealization of the manifest image, which seeks to integrate the various existing scientific frameworks – each equipped with its own (and distinct) theories, levels of analysis, methodologies, instruments, etc. – into a single image that would be constantly aided by the feedback provided by our

[4] Furthermore, there are good reasons to expect that intuitions about mental states can inadvertently influence scientific thinking and philosophical reflections on the human mind (Carruthers, 2020a, P. 214; Lau, Michel, 2020).

experiences in the manifest world itself. Thus, although it arises from the manifest image, this scientific image would, in a way, be beyond it, thanks to the use of tools that enable the measurement and observation of "manifestly" imperceptible phenomena (Sellars, 1963, p. 21).

It is essential to note that the different world images coexist on a continuum, even though the only way to access the scientific image is through scientific theories, methods, and instruments. Thus, a chair, for example, is a chair to us in our manifest image, but at the same time, it is a collection of particles when observed through the "lenses" of the scientific image. Here, there would be no decisive moment when this collection of particles "becomes" a chair. Both are the two things at the same time. What changes is who or what observes them. Scientific knowledge about these particles is just one way, among possible others, of making sense of the world. However, through the lenses of the scientific image, human perception does not function solely for the knowledge and classification of objects. There would then be an underlying nature to the Sellarsian images, which describes the perception of the world as it best suits us (in terms of survival and experiences). Thus, we can say that our perception is oriented by the possibilities of interaction with the niche in which we find ourselves.

The possibilities for interaction were noted by ecological psychologist James Gibson in 1979, for which he assigned the technical term "affordances." Affordances are what the environment and the organism can offer interactively, indicating the existence of a necessary complementarity between the organism and its niche (1979, p. 197). For Gibson, then, perception would be an active process aimed at detecting affordances, rather than a passive one aimed at detecting and categorizing the real qualities or properties of our world. For example, something like a chair would not be initially perceived by us as, for instance, a wooden, brown, and old object (or a collection of atoms), but rather as something sitable. Furthermore, affordances consist of unique relationships (or dispositions) for each organism, given their capacities and their history of interactions with the environment. For example, my perception of the affordance of a building's "climbability" is different from that of a parkour practitioner. Thus, this Gibsonian insight indicates a possible reason for the existence and configuration of conscious perception as we experience it: to enable better interaction with our environment. After all, considering the complexity of the contexts in which we live, we need an embodied mental apparatus with equally complex, sophisticated, and optimized capacities to make sense of what is important to us in these environments and thus (survive) in them (Godfrey-Smith, 2002).

Finally, from the perspective of the so-called "Free Energy Principle"[5], the co-determination of organism and environment and the details of this important interaction become clear. According to this principle, living beings can be defined as thermodynamically open autopoietic systems that are "far from equilibrium." That is, systems that, through their interactions and transformations, continuously regenerate themselves, creating relational networks that produce their own constituent states, forming locally concrete units, always striving to achieve their thermodynamic equilibrium.[6] For the maintenance of this equilibrium, organisms need to use their interactions with the environment, guiding themselves in it in such a way as to revisit, over time, homeostatically favorable states, even in the face of the continuous and random influence of environmental fluctuations. And, in doing so, they modify their environment and their relationship with it, which also changes what they capture from the environment – in a process of co-determination that underlies the perception-action cycle.

The phenotype of each living organism determines a set of possible and desirable states for it, which increase its probability of survival in a given niche over time (for example, a fish has a higher probability of surviving in water than out of it, because of its phenotype). The greater the structural complexity[7] of the organism, the smaller its set of possible states (for example, a human has less possibility of surviving in very different environments than a bacterium). Therefore, the survival of the organism depends on reducing the possibility of leaving this set and entering an undesirable state (Iennaco, Maia, Sayeg, 2023).

A living organism tends to occupy certain attractor states throughout its life, which reflect a model of its niche (i.e., the part of the physical environment it can access) and its action within it (Sims 2021, p.XXXI). If it does not maintain these states over time, it loses its distinction from the environment and thermodynamically dissipates – that is, it dies. These models are merely implicit statistical expressions of the organism's form, phenotype, and tendencies of internal states. And, as they probabilistically map the indirect interaction

---

[5] The Free Energy Principle mathematically describes the ability of living organisms to interact with their environment. According to the Free Energy Principle, living organisms must seek homeostatically favorable states and avoid unfavorable states. The "free energy" (which names the principle) is a theoretical-informational measure that can be minimized by the organism in two ways, thus showing the discrepancy between an implicit model and reality – a difference that, in the long run, needs to be reduced for the organism to survive (Friston, 2010).

[6] Notably, this type of system never actually achieves a static equilibrium. What such systems manage to attain is called homeorhesis, that is, they maintain an optimal pattern of states over time – as long as they are alive. Throughout this text, the concept of "homeostasis" will be used as a synonym for "homeorhesis," that is, equilibrium over time (and not at a specific moment).

[7] In this context, structural complexity is defined in terms of nestings or hierarchies of a set of statistical interdependencies (among internal, external, and sensory states) known as the "Markov Blanket." The greater the number of its statistical interdependencies, the lower the degree of freedom of the system as a whole.

between internal and external states, they can be seen as sub-personal estimates by the organisms – or approximate inferences made by them – about the environmental causes of their sensory/intermediate states (Corcoran, Pezzulo, Hohwy, 2020). Thus, the continuous existence of an organism implies that it incorporates an implicit model of its environment, and any action that favors this existence is equivalent to obtaining new evidence for the model incorporated by it – that is, an update of its network of sub-personal estimates.

2. Predictive Processing

Predictive Processing describes the general processes and mechanisms involved in human cognition, emphasizing the action-perception cycle, and proposes the unification of perception, cognition, and action into a single inferential process based on probabilities. Broadly speaking, it illustrates the brain as a hierarchical and embodied machine that, in order to perceive and act in the world, produces predictions[8] about the causes of received sensory stimuli, updating them as necessary or seeking new evidence in the environment that can corroborate them. Specifically, the predictions, along with external information, are generated by generative models, which can be local (i.e., models of the organism or parts of it) or global (i.e., models of the organism situated in its niche), and are used to guide action. In this context, it is assumed that the brain does not have direct access to the external environment and, therefore, makes (predictive) inferences about it (Parr, Pezzulo, Friston, 2022).

The hierarchical architecture of the predictive brain is composed of layers with different levels of complexity in information processing, operating on distinct spatiotemporal scales (Hohwy, 2013, p. 28; Wiese, Metzinger, 2017, p. 6). Each layer has at least two functional units: those that deal with predictions and those that deal with prediction errors. Communication within and between each layer occurs in a bidirectional flow, that is, from bottom to top – from sensors to the cortex – and from top to bottom – from the cortex to the sensors – in the hierarchy (Friston, 2005, 2010). The units responsible for predictions send inhibitory signals in the descending pathway, and the error units send excitatory signals in the ascending pathway. The mechanics of this communication in neurotypical conditions balance the influence that cognition has on perception and, at the same time, the amount of information that will be assimilated by higher layers. In Predictive Processing, such

[8] "Prediction" and "Inference" will be used as synonyms in this text.

mechanics are calculated through an approximation of Bayes' Rule, ensuring the maintenance and optimization of generative models in a probabilistic and dynamic manner.[9]

In this context, predictions can be understood as hypotheses that the brain makes to try to discover the causes of the effects that the body's sensors have captured. They are composed of prior information or "priors," emotions, biases, etc., that the organism has acquired throughout its life, modulating expectations of capturing certain information and, therefore, the way the organism interacts with its niche (Dołęga, Dewhurst, 2021, P. 3; Clark, 2016). Notably, predictions do not necessarily concern future issues but rather information that has not yet been observed by the organism (Sprevak, 2021, p. 9).

On the other hand, the effects or information captured by the organism's sensors go through some selection processes to compose the so-called prediction errors. One of these processes, for example, is what differentiates relevant information from irrelevant information, or noise, which ends up teaching the sensors themselves what types of information they should capture (for example, an ant bite) and ignore (for example, the touch of clothes on the body). Another selection process is precision weighting. In it, the brain adjusts the extent to which sensory information will be modulated by predictions. If a prediction is very different from the evidence, it is necessary to increase the capture of new information so that the generative models are updated as coherently as possible with the external environment. This occurs when something contrasts with what we are used to or expect to happen, for example, when a pedestrian crosses at the wrong time and in front of your car. On the other hand, if the predictions are quite reliable, the processing of new information will be mitigated, for example, when we are driving alone on a straight highway (Piekarski, 2021, p. 8).

An important characteristic of Predictive Processing is the energy efficiency in information processing. The descending flow sends predictions to the lower layers of the hierarchy to inhibit or "explain away" redundant information propagated by the ascending flow, ensuring that only residual prediction errors are propagated to the upper layers via the ascending flow (CLARK, 2016, pp. 37-39). As a result, the organism first constructs a general essence of its external environment—from the components already expected and deduced by the predictions—before focusing on any details of these environments, in a kind of "forest first, trees later" cognition (FRISTON, 2005, p. 825).[10]

[9] For a description, application, and more information on Bayesian formality, see Parr, Pezzulo, Friston, 2022.
[10]Logically, if there is something more salient or extremely surprising in this scenario, directing attention, then the order would be: "forest first, King Kong second, trees later" (or something like that).

Another important characteristic of Predictive Processing is the particular way the action-perception cycle is treated. This cycle is well-known in cognitive sciences and phenomenology (HUSSERL, 1970, p. 161), positing that the organism needs to act in the world to perceive and needs to perceive to act in the world. However, in the predictive context, the action-perception cycle dissolves into a single process called "active inference."[11]. In this process, the (predictive) inferences are guided by and for action. Specifically, action updates the generative models by enabling the capture of new information, which, considering the dynamic nature of the external environment, makes it necessary for the organism to remain active in its environment to stay synchronized with it.

The main motivation of the predictive brain in maintaining the balance between new information from the external environment and predictions from the internal environment is the orientation and interaction of the organism in its niche. The more time an organism spends and acts in its niche, the more predictable the niche will become, and therefore, the better its orientation and interaction within it—more affordances. This logic ensures not only the organism's survival but also the best possible way to survive. Even if the niche is roughly predictable under typical conditions, it will provide specific new information to the organism at all times, so that it learns this information and updates its generative models over time. Note that this process also results in changes in the niche, both made by the organism itself and by the dynamics of natural conditions. This ends up creating inter-relational and co-determinant regularities between the organism and the niche, and these regularities will induce the organism to remain there throughout its life and allow it to respond in advance to potential disturbances and avoid stressful or unfavorable conditions for its existence (CORCORAN, PEZZULO, HOHWY, 2020, p.6). Therefore, the very survival of organisms with predictive brains ensures that their models are supervised by what they are trying to model—the external environment—and, at the same time, that they modulate the external environment itself to satisfy their favorable conditions of existence (FRISTON, 2005, p. 825; HOHWY, 2016, p. 47).

In general terms, the models discussed here were first conceptualized in cybernetics, in the Good Regulator Theorem proposed by Roger Conant and William Ashby in the 1970s. It postulates that every good regulator of a system must contain a model of the system

[11] It is important to note that "Active Inference" per se is a complete theoretical framework in the literature, often used synonymously with Predictive Processing. However, within the context of the main formulations of Predictive Processing, it is merely a form of prediction error minimization. For a robust explanation of Active Inference, see PARR, PEZZULO, FRISTON, 2022.

and its actions upon it (CONANT, ASHBY, 1970)[12]. In the context of Predictive Processing, generative models capture various distributions (or mappings) of information to generate new models that resemble and assimilate the originally captured information. Thus, they represent the (hidden) causes of the captured sensory information and generate predictions about what is expected to be observed in the external and internal environments. Additionally, organisms become capable of generating prediction error models throughout their lives, which are used to predict noise levels in the captured information and also act as priors that will modulate subsequent information captures (HOHWY, 2013, p. 194).

Therefore, it can be said that there are at least four types of generative models in the predictive brain: external models, internal models, error models, and system models. They generate predictions for three types of perception: exteroceptive, which is the perception of the external world; proprioceptive, which is the perception of the body in the external world; and interoceptive, which is the perception within the body. Additionally, these models are guided by two distinct but co-determinant functions: exploring the world and controlling/regulating the organism. It is worth noting that these models are not equivalent to mental representations as described by some philosophical proposals, as they are merely probabilistic distributions used as heuristics in the context of the predictive brain (CLARK, 2016). And, although this specific division is not entirely conventional in the literature on Predictive Processing, it, as described in the work of neuroscientist Anil Seth (2018, 2021), will be useful for understanding the functioning of an organism's internal environment and subjectivity.

2.1 Conditional and Counterfactual Epistemic Actions

The perception of an organism involves not only its niche and its body within that niche but also its various possibilities for action. This means that the predictive brain tries to explain the causes of the effects captured in the external environment and also how these effects would change under certain actions taken by the organism. When it is minimizing exteroceptive prediction errors—and thus explaining the effects of the captured causes—it also suppresses proprioceptive prediction errors by altering its own sensations through behavior. For example, to drink coffee, the body needs to perform a movement to pick up the coffee cup. In doing so, the predictive brain needs to calculate which action is the best way (just as it calculates which hypothesis is the best to explain a sensation) to achieve

[12] More specifically, the Good Regulator Theorem considers the entropy of the output variation of a controlled system, showing that entropy is minimized when there is a mapping of the system states to the regulator states, with this mapping being equivalent to a model of the system (CONANT, ASHBY, 1970).

the goal. Thus, the models generate predictions of the proprioceptive consequences of the hand's trajectory to the cup and from the hand with the cup to the mouth, which will be fulfilled by the movements, in this case, mainly of the trunk, arm, and hand (SETH, FRISTON, 2016, p.3). In other words, to drink the coffee, the predictive body generates a prediction of the bodily movements necessary to bring the cup to the mouth before actually performing the movement.

However, these predictions will only be fulfilled (and the cup brought to the mouth) if proprioceptive prediction errors are attenuated (to avoid providing evidence that the organism is not yet actually acting) already in the spinal cord. This means that proprioceptive predictions will replace descending motor commands and will be fulfilled by peripheral reflexes. This substitution occurs whenever there is a need to regulate the body. For example, imagine you are walking to the bakery. Your goal is the bakery, and your body is performing the necessary movements to get there. Unexpectedly, you need to, say, avoid a new hole in the sidewalk. At that moment, considering the context of Predictive Processing, the substitution of descending motor movements by proprioceptive predictions occurs, as you saw the hole and will avoid it.

In fact, much less is needed for such substitution to occur. This is because system generative models that control multimodal predictions need to keep the organism regulated, not only to continue walking but also to maintain heart rate, sugar levels, cell genesis, and so on. Thus, the interoceptive perception of autonomic, hormonal, visceral, and also immunological signals will be associated with the dynamics of the perception of the external world and the perception of the body in the world. Similarly, all types of predictions will be compared with all types of prediction errors by the system's generative models. In this, the notion of interoceptive predictions or inferences suggests that bodily states are regulated by autonomic reflexes, which are induced by the multimodal predictions of the system's generative models.

Two crucial implications of these substitutions are the self-fulfilling nature of proprioceptive and interoceptive predictions, as the predictive body seeks, through action, the information that will make its predictions come true. For example, it is predicted that body temperature will remain constant over time, which is precisely why the body itself contributes to making this happen (SETH, 2021, p.191). The other implication is the conditionality of predictions, in the sense that the predictive body relies on generative models to indicate which action, among various other actions, would best reduce prediction errors. As a result, in this context, such implications equate the predictions of the system's generative models to perception itself.

According to Seth and Tsakiris (2018), the facets of perception, in this context, can be understood through the distinction between epistemic inferences and instrumental (control-oriented) inferences. The former pertains to exploratory perception, with the brain trying to discover the causes of the effects captured in the internal and external environment. For Seth and Tsakiris, one of the tools available in this exploration is called “objecthood,” which refers to the phenomenological character of perceived objects composed of conditional and counterfactual predictions. For example, to recognize an object, the predictive brain makes several epistemic predictions about it, such as how it presents itself to the visual system and also how the information about this object would change given this (conditional) or that epistemic action (counterfactual). This idea is actually based on the theory proposed by O’Regan and Alva Noë in 2001 called the Sensorimotor Contingency Theory. It postulates that the phenomenology of objecthood comes from the domain of relevant sensorimotor contingencies. This means that what an organism subjectively experiences will depend on a practical domain of how its actions will change sensory information. Perception, then, emerges from the brain’s implicit knowledge of how actions and sensations are coupled, suggesting that all perceptual modalities ultimately consist of what the organism does (SETH, 2021, p. 127). Instrumental inferences, on the other hand, are not experienced in the same way as the phenomenology of objecthood, since we do not experience being a body in terms of organ arrangements. Let’s get into the details.

2.2 Instrumental Actions and Interoceptive Inferences

As we have seen, interoception is the perception of the body’s internal activity, whether it is physiological, hormonal, emotional, etc. Seth argues that his ideas of interoception developed from the theory of emotions proposed by William James (1884) and Carl Lange (1885). According to the authors, emotions are originated by the perception of changes in the body (and not the other way around). In 1962, Schachter and Singer showed that injections of adrenaline (which approximately provoke a state of physiological arousal) could generate anger or euphoria depending on the simultaneous context (an irritated or euphoric accomplice). Thus, they concluded that the emotional experience would be defined by the combination of physiological changes and cognitive evaluation, that is, that an emotion would be the interpretation of a bodily expression. Although it did not become popular, this logic continued to underpin the understanding of emotions and their neural bases, such as, for example, in the “Somatic Marker Hypothesis” (DAMASIO, 2000) and in the “Sentient Self” model (CRAIG, 2009), which relate them to the concept of interoceptive consciousness or interoceptive sensitivity (CRITCHLEY et al., 2004).

Thus, interoception, in the predictive context, is the result of inferences from internal generative models based on top-down predictions that provide a homeostatic set point for primary interoceptive afferent signals. When there is a prediction error, the sympathetic or parasympathetic effector systems are activated to ensure homeostasis or allostasis. For example, “blushing with shame” is a reflex response to the predicted interoceptive consequence of sympathetic vasodilation of the smooth muscle in the face. This mechanism of autonomic reflexes is similar to motor reflexes, which allow the contraction of striated muscle to be prescribed by equilibrium points defined by descending predictions to motor neurons in the spinal cord (SETH, FRISTON, 2016). The comparison between actual and predicted interoceptive responses seems to occur mainly in the anterior insular cortex (SETH, SUSUKI, CRITCHLEY, 2011; HAGGARD, 2008; TSAKIRIS, 2010).

Interestingly, the activity of the anterior insular cortex is also associated with subjective states of feeling and conscious presence. These subjective states can range from experiences of body ownership (Limanowski, Blankenburg, 2013) to the perception of a world from a point of view (Blanke, Metzinger, 2009). Additionally, these states can also be experienced through intention and agency (Friston, 2010), the expression of a continuous self over time (Scoville, Milner, 2000), and finally, a social self – often shaped by others’ perceptions of this self (Frith, Frith, 2012). According to Seth and Friston (2016), interoception is responsible for this range of states that structure the experience of ‘being’ and ‘having a body.’ Specifically in the context of Predictive Processing, these experiences emerge from top-down predictions of the multimodal causes of interoceptive afferents, being characterized as ‘interoceptive inferences.’ Thus, there is a circular causality, where predictions about bodily states engage autonomic reflexes via active inference, while interoceptive signals inform and update these predictions.

The mechanism of interoceptive inferences is control-oriented, as it is based on the high-precision priors that the system’s generative models possess. These priors allow interoceptive predictions to self-fulfill through active inference (as in the example of body temperature in the previous section), ensuring that the best interoceptive predictions are selected for the expected physiological fluctuations, thus guaranteeing physiological regulation. The predictive brain is then able to minimize its prior interoceptive predictions of the body’s conditions to ensure that physiological variables remain stable, ultimately implying that, as long as the organism is alive, its predictive brain will never update the prior that it is alive (Seth, 2021, p. 181).

Consequently, organisms like ours will have a relatively immutable bodily experience, given that strong predictions about stable bodily states generate an attenuation of the detection of fluctuations in these states so that only the most unexpected variations are perceived and promptly minimized[13]. It would be as if we could not perceive our physiological conditions changing when they actually do. An important implication of this perspective for our perception is that interoceptive priors allow organisms to make counterfactual inferences with high levels of certainty, which dangerously distances them from their reality by acting in the absence of direct sensory feedback. Furthermore, counterfactual inferences seem capable of generating great confusion regarding the relationship between internal patterns and inferences of external states – so that things like the redness of a tomato, for example, appear to us as real as the tomato itself (Clark, Friston, Wilkinson, 2019, p. 23, 29).

Indeed, just as conscious perception is the result of hypotheses that predictive brains formulate about the causes of effects captured by sensors, emotional or affective perception is the result of hypotheses that predictive bodies formulate about the causes of interoceptive effects. Moreover, thanks to the system's generative models that handle multimodal information, all perceptions are interconnected. This means that conscious perception will be influenced by emotional perception and vice versa. Therefore, the way someone sees the world is closely linked to how they feel at that moment, in a given context, and because of the context and that mood at that moment, they see the world in that way.

It is worth noting that interoceptive inferences should not be considered a generalization of Predictive Processing from exteroception. In fact, it is precisely the opposite: exteroceptive and proprioceptive content are consequences of the crucial imperative of physiological regulation. This implies that all "perceptual realms" are supported by inferential mechanisms with a functional, ontological, and phylogenetic basis in allostasis (Seth, Tsakiris, 2018, p.11). The imbalance of this cycle can lead to various mental disorders and significant loss of quality of life. Thus, to maintain the balance between the "perceptual realms," simple physiological changes can generate complex actions that, in turn, generate new possibilities for action, new samplings, and new states of the body. This implies, for social beings like us, a body-body integration within a body-environment integration that occurs at all times, at various levels of complexity, and whose main

[13] This perception of highly unexpected variables is unique and also plastic. For example, hunger is felt because the system that regulates nutritional levels has detected that they are low. Thus, the empty stomach will release the hormone ghrelin, which will stimulate the hypothalamus. Finally, the hypothalamus will generate the conscious motivation to eat (Cf. Albiero, 2011). Each part of this "path" will be different for each organism in detail. And, even though they are similar, a person can develop conscious control of this hunger (by following certain diets or fasting) and stop noticing it. Therefore, what is "highly unexpected" will depend on each organism in each context.

objective is the maintenance of our existence.

3. I Predict, Therefore I Am

In the previous section, we saw the bottom-up and top-down flows of information and how they optimize the generative models that compose perception. Specifically, bottom-up information brings the effects captured by sensors to update top-down information and maintain the best predictions for the context. These, on the other hand, are composed of information that the organism has acquired over time and, therefore, are more complex and stable. By adjusting top-down predictions to minimize bottom-up prediction errors, the brain's best perceptual hypotheses maintain their adherence to possible causes in the world. We also saw that interoceptive inferences guide generative models, which will act in the best possible way to regulate the body's physiology. The expression of these inferences is known to us as emotion and mood.

Seth (2021) proposes that, in the context of Predictive Processing, beyond emotions and mood, all expressions of subjectivity and identity in organisms like us compose the so-called "controlled hallucinations." Specifically, these controlled hallucinations are equivalent to the content of top-down multimodal predictions, as they are not capable of experiencing bottom-up information itself, but only interpretations of it (p.83). However, we have no reason to doubt that we perceive the world and reality (manifest) in a real way. The reason for this is relatively simple: even consciously perceiving only top-down predictions, our reality is real because we need to survive in it. Thus, we perceive the world not as it is (extra-perceptive), but as it is useful to us (in terms of affordances) to assist in guiding bodily actions that preserve the organism's physiological integrity. Therefore, phenomenological properties such as the redness of a tomato appear to us as objectively real in an externally existing environment, as this allows us to respond more quickly and effectively to its dynamics and fluctuations (p.138).

However, it is worth noting that Predictive Processing and controlled hallucinations are different things. The former is a theory about the mechanisms and processes by which the brain and body perform perception, cognition, and action. The latter concerns how the functioning of the brain and bodily mechanisms explains the phenomenological properties of conscious perception (p.107). Seth argues that describing perception in this way does not mean that anything goes (for example, if everything is a hallucination, we could jump out the window and fly), but rather that the way things in the world appear in perceptual experience is a construction of the predictive brain (p.93). Thus, our subjective experiences will never be identical to the corresponding occurrences of the external environment, as we will always be behind a sensory veil (Hohwy, 2013).

Furthermore, Seth explains how new experiences fit into the context of controlled hallucination, as they are entirely top-down. The point is that nothing is really new – whether in psychological, physiological, or phylogenetic terms. Everything we perceive for the first time is a new product of multiple characteristics we have encountered before, either as an individual or as a species. For example, seeing a tiger may be unprecedented for a Brazilian child, but she has already seen furry animals, other felines, quadruped mammals, etc., so everything she needed to perceive this tiger, she already had, just updated the details (perhaps she had never seen orange and black fur) (p.94).

The phenomenological part of perception, in this context, is neither localized nor based on objects or objecthood, but is part of the emotional and affective activity of organisms. Let's go back to the tiger. A child who sees a picture of a tiger for the first time might associate it with a cat and find it "cute." The cuteness of the tiger is not a property of the tiger, but rather an interpretation of the positive feeling that the child had upon seeing the tiger. Her behavior will then be one of laughter and joy, pointing to the picture and, say, saying she wants one for herself. Now, suppose this child goes to the zoo to see a live tiger and it comes very close to her, as if preparing to attack. Immediately, her joy will give way to fear. All the symptoms of the activated flight system will be subjectively felt as fear and interpreted negatively, generating the need to act, either by crying or moving away to reach safety and regulate the system again.

And, because these emotions are linked to the physiological system and are manifestations of interoceptive inferences, they are experienced in the most real way possible. The realization of interoceptive inferences, besides self-promoting the system's reactive patterns, tends to impact both exteroception and action selection, making us certain that the experience of subjective states such as the cuteness of a baby is as real to us as the baby itself (Dennett, 2017, p. 210). The system's generative models add the point of view to this mix, making these emotions not only as real as possible but also experienced from a specific point of view. Thus, we approach the figure of mentally complex individuals; individuals who, to maintain their balanced homeostasis and allostasis, feel and act based on these emotions, with their niche being their greatest influencer, or rather, their co-determinant.

Notably, for Seth, all experiences (hallucinatory or real) are always based on predictions and perceptual expectations for and within an organism's environment. When someone says they have a "hallucination," it would mean that the perceptual priors overloaded (for some reason or pathology) the networks that deal with sensory information, preventing the predictive brain from performing its usual check of environmental information against its predictions. Therefore, once the priors remain in balance with sensory evidence, the organism becomes more synchronized with its environment and stops hallucinating. The difference between perception and hallucination would then be determined only by this balance, which means that hallucination is a form of uncontrolled perception and perception is a form of controlled hallucination (Seth, 2021, p. 127).

3.1 Hallucinating a Top Hat with Rabbits Inside

Although controlled hallucinations are consistent with Predictive Processing and promisingly embodied, they do not escape a strong idealism, that is, a limitation to the mental itself. The reason, however, is not necessarily the uniqueness of top-down predictions in the creation of perception, but rather the claim that perception is a hallucination. The main definition of hallucination consists of a sensory perception that occurs in the absence of an external stimulus and usually arises from disturbances or in response to drugs. Even though hallucination is within the spectrum of perception (in pathological terms or altered states), it does not seem to be the best option to ground normative perception, even in predictive terms, since we are in sync with the external world. That is, the effects that our sensors capture from the environment are caused by the components of that environment, and not by our interpretations. The bottom-up flow, although not consciously accessible, is fundamental to conscious perception because it will precisely keep us in sync with the extra-perceptual world. Hallucinating is letting top-down predictions completely overshadow the bottom-up flow and simulate a capture of information that does not exist.

For illustration, the ecological disjunctivism proposed by philosopher Eros de Carvalho distinguishes perception from hallucination by stating that the former consists of a controlled experience of tuning into information, while the latter is passive and refractory to the organism's activities of exploration and tuning (Carvalho, 2023, p.171). Specifically, the act of perceiving is an embodied skill, dependent on action that is learned over time, whose function is to enable access to what exists. Therefore, for the author, the perceptual experience would be generated by the agent organism's knowledge about establishing a cognitive contact between itself and some object or event in its niche. Hallucination, on the other hand, would be the result of a failed attempt to establish this cognitive contact, with no knowledge of any object in the niche (Carvalho, 2021, p.298). This perspective clarifies the reasoning proposed at the beginning of this chapter. Perception (of the manifest image) is refined over time, actively and affordance-oriented, which means it depends on the interactions the organism has with the environment. Hallucination, on the other hand, has no sensory relation to this environment; it is an independent mental construct, usually triggered by some anomaly in the functioning of our embodied brains.

3.2 The Illusionist in the Jungle

Given the reasons to abandon the word "hallucination" in Seth's construct, we need a substitute, which here will be "illusion." An illusion consists of a sensory perception different from what the object or event in question really is. In other words, when we talk about perceptual illusions, we are not denying their relation to what supposedly exists in the external world, but denying that the manifest form through which we perceive it would be faithful to its (extra-perceptual) structuring; whereas perceptual hallucinations would occur independently of what exists in the external world. Therefore, adopting the word illusion means saying that conscious perception is not what it seems to be, which is entirely compatible with Seth's construct. In the author's words,

> "The entirety of perceptual experience is a neuronal fantasy that remains tied to the world through a continuous making and remaking of better perceptual hypotheses." (SETH, 2021, p. 87).

Thus, given the distinction between hallucination and illusion, what Seth seems to argue, then, is that "controlled illusions" are the manifest image of organisms like us, who use them to navigate the world and stay alive. This manifest image is not an absolute and universal reality, but the reality of this organism, as it experiences the world through a body. This subjective experience is composed of the interpretations that the predictive brain and body make of the very hypotheses that their models generated about the internal and the external, hypotheses that guide the action that will regulate the physiology and interactions of this organism with its niche.

This idea Seth proposed and that we slightly modified here has roots in the philosophical current of Illusionism. This current argues that phenomenal consciousness is an introspective illusion, that is, when we introspect about conscious states, we mistakenly represent them as having phenomenal properties independent of their relation to us, when in fact they do not. Seth acknowledges this current and says that, although he is sympathetic to the weak version of Illusionism (a version that does not eliminate the phenomenality of consciousness but only argues that it is not what it seems to be), he would not know if his theory would relate to such a current (Seth, 2021, p.282).

However, we have seen that the weak version of Illusionism[14] is entirely compatible with Seth's theory. Frankish (2017) states that subjectivity is formed by "quasi-phenomenal" physical properties of representations that, when introspected, are treated as if they were phenomenal (p.18). These quasi-phenomenal properties produce a specific effect in our acts of introspection and thus create the evident sensation of subjectivity. Therefore, when we see an object, our sensory system collects various data about it, which are processed and modified, so that if introspection accesses the caricatured representations we build about this object, they would "mistakenly" communicate (as phenomenal) its attributes, along with the perceptual content itself. As we have seen, quasi-phenomenal properties seem to have the same effect as interoceptive predictions, which, being highly precise and self-fulfilling, end up causing an intense sensation that the content of the prediction, once computed by the system's generative models, is real and fully corresponds to the external world.

The illusion itself would be to believe that, instead of understanding that what we experience as phenomenal is the result of our physical properties, as embodied beings, the illusion would be to suppose that this feeling, this qualitative impression, is a kind of capacity of the observed objects to influence us. That is, that the redness would be illusorily presented as an ability of the tomato's surface to affect us in some way (Cf. Leal-Toledo, De Vasconcelos, 2023, p. 235). Thus, and concomitantly with Predictive Processing, conscious perception and the subjectivity contained within it are a set of illusions created by the continuous cycles of updating generative models, being, therefore, a kind of dynamic patchwork of informational pieces, tuned to the environment but indirect to it. And, regardless of whether we know (or at least believe) that this perception is illusory, we do not stop perceiving it in this way (considering, of course, neuro-normative conditions).

In philosophical terms, Seth's theory reveals the experience of a "really-existing-self" and the complexity of the self as aspects of perceptual inference capable of helping us overcome Chalmers' Hard Problem. In the author's words,

> "the intuition of the Hard Problem that subjective experience is somehow separate from the rest of corporeal nature - a really existing immaterial internal observer looking at a material external world - turns out to be just another confusion between how things seem and how they are." (Seth, 2021, p.193).

The confusion between appearance and reality, as we have seen, has generated numerous debates about whether subjectivity, as it is experienced by us, is or is not an

[14] In a previous work, I explored, together with Professor Gustavo Leal Toledo, the strong version of Illusionism in greater detail. See Leal-Toledo, De Vasconcelos, 2023.

object of scientific study. Here we understand that, once the characteristics exposed in the first section are accepted (that is, highly intellectualist and private to the individual), subjectivity would not be such a thing.

However, once we adopt the idea of illusionism and understand that this subjectivity is nothing more than an interpretation of the effects of the body internally and externally situated, we turn to other (and perhaps better) questions, which do not hinder progress with a gap between what can be scientifically studied and what cannot (Dennett, 2016). New questions seek to understand phenomenality from more neutral and non-anthropocentric assumptions, such as the Illusion Problem (Frankish, 2017), which aims to understand how and why we have the illusion of subjective experiences, and the Real Problem (Seth, 2021), which aims to understand subjectivity from Predictive Processing. The importance of using the Predictive Processing framework lies in its growing applicability in multiple disciplines (Active Inference Institute, 2023, p.4) and its progressive empirical robustness (Cf. Hodson, et al., 2024), making it an excellent tool for scientifically based studies of human subjectivity.

In summary, we address here an illusionist and predictive perspective of human subjectivity, first by getting rid of intellectualist assumptions that foster the epistemic gap and adopting Sellarsian assumptions, which imply the coexistence of two ways of experiencing the world – the one we are accustomed to and the one dominated by scientific tools and methodologies. The concept of affordances was added to understand an active perception oriented towards interactions with the environment, which are co-determinants for it and the organism. Then, we formulated that the purpose of this perception is the survival of the organism in its niche, on various distinct scales, by minimizing the chances of it finding itself in unfavorable situations for its existence. To deepen the idea of minimization, we brought the idea of brain functioning based on Predictive Processing, whose main processes are the minimization of prediction errors and the updating of generative models. Then we saw that, in fact, this would apply to the whole body, since we can understand the role of action in this model updating through the same processes and mechanisms. Finally, the internal part of the body is crucial for subjectivity, as it is the effects captured by internal sensors that will generate emotions and affections. And, because interoceptive predictions are highly accurate and interpreted by the system's own model, they generate the sensation of a point of view, making emotions felt by "someone." In this, controlled hallucinations or illusions would be the interpretative "junction" of sensory and emotional models from a point of view.

References

ACTIVE INFERENCE INSTITUTE, et al. (2023). The Active Inference Institute and Active Inference Ecosystem. Zenodo. https://doi.org/10.5281/zenodo.8266281.

Albiero, K. (2011). Mecanismos Fisiológicos e Nutricionais na Regulação da Fome e Saciedade. Revista Uningá, 30(1).

Blanke, O., & Metzinger, T. (2009). Full-body illusions and minimal phenomenal selfhood. Trends in Cognitive Sciences, 13(1), 7-13.

Carruthers, P. (2020a). How Mindreading Might Mislead Cognitive Science. Journal of Consciousness Studies, 27(7-8), 195-219.

Carruthers, P. (2020b). Explaining the Empiricist Bias. Journal of Consciousness Studies, 27(7-8), 230-235.

Carvalho, E. (2021). An Ecological Approach to Disjunctivism. Synthese, 198, 285-306.

Carvalho, E. (2023). O disjuntivismo ecológico e o argumento causal. Trans/Form/Ação, 46, 147-174.

Clark, A. (2016). Surfing Uncertainty: Prediction, Action, and the Embodied Mind. Oxford: Oxford University Press.

Clark, A., Friston, K., & Wilkinson, S. (2019). Bayesing Qualia: consciousness as inference, not raw datum. Journal of Consciousness Studies, 26(9-10), 19-33.

Chalmers, D. J. (1995). Facing up to the problem of consciousness. Journal of Consciousness Studies, 2(3), 200-219. (Article)

Conant, R., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. International Journal of Systems Science, 1(2), 89-97.

Corcoran, A., Pezzulo, G., & Hohwy, J. (2020). From allostatic agents to counterfactual cognisers: active inference, biological regulation, and the origins of cognition. Biology & Philosophy, 35(32), 1-45.

Craig, A. D. (2009). How do you feel--now? The anterior insula and human awareness. Nature Reviews Neuroscience, 10(1), 59-70.

Critchley, H. D., Wiens, S., Roth, Y., & O'Doherty, J. P. (2004). Neural systems supporting interoceptive awareness. Nature Neuroscience, 7(2), 189-195.

Damasio, A. R. (2000). A second chance for emotion. In R. D. Lane & L. Nadel (Eds.), Cognitive neuroscience of emotion (pp. 12-23). Oxford University Press.

Dennett, D. C. (2017). From Bacteria to Bach and Back: The Evolution of Minds. Penguin Books Limited.

Dennett, D. C. (2016). Illusionism as the obvious default theory of consciousness. Journal of Consciousness Studies, 23(11-12), 65-72.

Dennett, D. C. (1991). Consciousness Explained. Londres: The Penguin Press.

Dołęga, K., & Dewhurst, J. (2021). Fame in the predictive brain: a deflationary approach to explaining consciousness in the prediction error minimization framework. Synthese, 178, 7781-7806.

Feldman, H., & Friston, K. J. (2010). Attention, uncertainty, and free-energy. Frontiers in Human Neuroscience, 4(215).

Frankish, K. (2012). Quining diet qualia. Consciousness and Cognition, 21(2), 667-676.

Frankish, K. (Ed.). (2017). Illusionism as a Theory of Consciousness. Exeter: Imprint Academic.

Friston, K. J. (2010). The free-energy principle: a unified brain theory? Nature Reviews Neuroscience, 11(2), 127-138.

Friston, K. J. (2009). The free-energy principle: A rough guide to the brain? Trends in Cognitive Sciences, 13(7), 293-301.

Friston, K. J. (2005). A theory of cortical responses. Philosophical Transactions of the Royal Society B: Biological Sciences, 360(1456), 815-836.

Friston, K. J., Parr, T., & De Vries, B. A. (2017). Network Neuroscience. The graphical brain: belief propagation and active inference. Network Neuroscience, 1(4), 389-409.

Frith, C. D., & Frith, U. (2012). Mechanisms of social cognition. Annual Review of Psychology, 63, 287-313.

Gelman, S. A. (2003). The Essential Child: Origins of Essentialism in Everyday Thought. Oxford: Oxford University Press.

Gelman, S. A., & Wellman, H. M. (1991). Insides and essence: Early understandings of the non-obvious. Cognition, 38(3), 213-244.

Gibson, J. J. (1979). The Ecological Approach to Visual Perception. Houghton, Mifflin and Company. (Book)

Godfrey-Smith, P. (2002). Environmental complexity and the evolution of cognition. In R. Sternberg & J. Kaufman (Eds.), The Evolution of Intelligence (pp. 233-249). Lawrence Erlbaum.

Haggard, P. (2008). Human volition: towards a neuroscience of will. Nature Reviews Neuroscience, 9, 934-946.

Hodson, R., Mehta, M., & Smith, R. (2024). The empirical status of predictive coding and active inference. Neuroscience & Biobehavioral Reviews, 157, 105473.

Hohwy, J. (2013). The Predictive Mind. Oxford: Oxford University Press.

Hohwy, J. (2016). The self-evidencing brain. Nous, 50(2), 259-285.

Hohwy, J., & Seth, A. K. (2020). Predictive processing as a systematic basis for identifying the neural correlates of consciousness. Philosophy and the Mind Sciences, 1(2), 1-34.

Husserl, E. (1970). The Crisis of European Sciences and Transcendental Phenomenology: An Introduction to Phenomenology, trans. D. Carr. Evanston, IL: Northwestern University Press.

Iennaco, M. L., Maia, T., & Sayeg, P. (2023). Processamento Preditivo: uma introdução à proposta de unificação da cognição humana. Principia, 27(3), 425-452.

James, W. (1884). What is emotion? In W. Dennis (Ed.), Readings in the history of psychology (pp. 290-303). Appleton-Century-Crofts.

Lau, H., & Michel, M. (2020). A socio-historical take on the meta-problem of consciousness. Journal of Consciousness Studies, 26(9-10).

Leal-Toledo, G., & Vasconcelos, M. L. (2023). O problema da ilusão: uma breve introdução e defesa da teoria ilusionista de Keith Frankish. In L. Peruzzo, K. Candiotto, & M. Karasinski (Eds.), Tendências Contemporâneas de Filosofia da Mente e Ciências Cognitivas (pp. 166-186). PUCPRESS.

Limanowski, J., & Blankenburg, F. (2013). Minimal self-models and the free energy principle. Frontiers in Human Neuroscience, 7, 547.

Nagel, T. (1974). What is it like to be a bat? Philosophical Review, 83(4), 435-450.

O'Regan, J. K., & Nöe, A. (2001). A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences, 24(5), 939-1031.

Parr, T., Pezzulo, G., & Friston, K. (2022). Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. 1st ed. The MIT Press.

Piekarsi, M. (2021). Understanding Predictive Processing: A Review. AVANT, 12(1).

Sellars, W. F. (1963). Science, perception, and reality. New York: The Humanities Press.

Seth, A. K. (2021). Being You. Faber & Faber.

Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. Trends in Cognitive Sciences, 17(11), 565-573.

Seth, A. K., & Friston, K. J. (2016). Active interoceptive inference and the emotional brain. Philosophical Transactions of the Royal Society B: Biological Sciences, 371(20160007), 1-10.

Seth, A. K., & Tsakiris, M. (2018). Being a Beast Machine: The Somatic Basis of Selfhood. Trends in Cognitive Sciences, 22(11), 969-981.

Seth, A. K., Suzuki, K., & Critchley, H. D. (2011). An interoceptive predictive coding model of conscious presence. Frontiers in Psychology, 2, 395.

Sims, M. (2021). Strong continuity of life and mind: the free energy framework, predictive processing and ecological psychology. Edinburgh: The University of Edinburgh.

Schachter, S., & Singer, J. E. (1962). Cognitive, Social, and Physiological Determinants of Emotional State. Psychological Review, 69(5), 379-399.

Scoville, W. B., & Milner, B. (2000). Loss of recent memory after bilateral hippocampal lesions. Journal of Neuropsychiatry & Clinical Neurosciences, 12(2), 103-113.

Sprevak, M. (2021). Predictive coding I: Introduction. [Preprint]. URL: http://philsci-archive.pitt.edu/id/eprint/19365.

Tsakiris, M. (2010). My body in the brain: A neurocognitive model of body-ownership. Neuropsychologia, 48(3), 703-712.

Tye, M. (2021). Qualia. In E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy. https://plato.stanford.edu/entries/qualia/

Wiese, W., & Metzinger, T. (2017). Vanilla PP for Philosophers: A Primer on Predictive

Processing. In T. Metzinger & W. Wiese (Eds.), Philosophy and Predictive Processing: 1 (pp. 3-25). MIND Group.

White, C. (2021). An Introduction to the Cognitive Science of Religion: Connecting Evolution, Brain, Cognition, and Culture. Routledge.