



A mathematical model of embodied consciousness



David Rudrauf^{a,*}, Daniel Bennequin^b, Isabela Granic^c, Gregory Landini^d, Karl Friston^e,
Kenneth Williford^f

^a CISA/FAPSE/Campus Biotech, University of Geneva, Geneva, Switzerland

^b Department of Mathematics, University of Paris 7, Paris, France

^c Department of Behavioral Science, Radboud University, Nijmegen, Netherlands

^d Department of Philosophy, University of Iowa, Iowa City, USA

^e Wellcome Trust Center for Neuroimaging, University College London, London, UK

^f Department of Philosophy, University of Texas, Arlington, USA

ARTICLE INFO

Article history:

Received 21 November 2016

Revised 21 March 2017

Accepted 23 May 2017

Available online 26 May 2017

Keywords:

Consciousness

Mathematical model

Projective geometry

Active inference

Embodiment

Self-Consciousness

Perspective taking

Free energy minimization

Point of view

Predictive coding

Bayesian approaches to cognition

Cognitive resilience

Spatial phenomenology

Subjective experience

ABSTRACT

We introduce a mathematical model of embodied consciousness, the Projective Consciousness Model (PCM), which is based on the hypothesis that the spatial field of consciousness (FoC) is structured by a projective geometry and under the control of a process of active inference. The FoC in the PCM combines multisensory evidence with prior beliefs in memory and frames them by selecting points of view and perspectives according to preferences. The choice of projective frames governs how expectations are transformed by consciousness. Violations of expectation are encoded as free energy. Free energy minimization drives perspective taking, and controls the switch between perception, imagination and action. In the PCM, consciousness functions as an algorithm for the maximization of resilience, using projective perspective taking and imagination in order to escape local minima of free energy. The PCM can account for a variety of psychological phenomena: the characteristic spatial phenomenology of subjective experience, the distinctions and integral relationships between perception, imagination and action, the role of affective processes in intentionality, but also perceptual phenomena such as the dynamics of bistable figures and body swap illusions in virtual reality. It relates phenomenology to function, showing the computational advantages of consciousness. It suggests that changes of brain states from unconscious to conscious reflect the action of projective transformations and suggests specific neurophenomenological hypotheses about the brain, guidelines for designing artificial systems, and formal principles for psychology.

© 2017 Elsevier Ltd. All rights reserved.

1. Introduction

1.1. Basic rationale and contextualization

We introduce a mathematical model of embodied consciousness, the Projective Consciousness Model (PCM), that identifies core aspects of the phenomenology of consciousness with specific computational mechanisms. The model integrates active inference and projective geometry to define a Field of Consciousness (FoC). The FoC operates as a mechanism of arbitration based on the anticipation of the consequences of action. Its architecture embeds mechanisms of cognitive and affective appraisal and reappraisal that exploit both local perceptual inference and non-local projec-

tive imagination in order to orient attention and motivate action. The overall algorithm yields a model that optimizes resilience to stressors and adverse conditions. Such optimization of resilience has been hypothesized to be a core function for embodied consciousness (Rudrauf, 2014). The PCM explains a broad range of perceptual, cognitive, and affective phenomena and leads to straightforward neurocomputational hypotheses with a broad application potential.

The PCM is based on an interdisciplinary synthesis of multiple traditions and concepts: Phenomenology (Husserl, 1913; Merleau-Ponty, 1945; Williford, 2006; Williford et al., 2012), Francisco Varela's neurophenomenology program and reflection on embodiment and autonomy (Damasio and Rudrauf, 2006; Rudrauf et al., 2003; Varela, 1979, 1999), Karl Friston's theory of active inference based on free energy minimization as a model of embodied autonomous agents (Friston, 2010), Alain Berthoz's emphasis on the sense of space as a multisensory, supramodal

* Corresponding author.

E-mail address: david.rudrauf@unige.ch (D. Rudrauf).

phenomenon (Berthoz, 1997); projective geometry, from the geometry of art and architecture during the Renaissance period to the theory of abstract cameras and the latest advances in computer graphics (Greenberg, 1993; Hartley and Zisserman, 2003), virtual reality, as it brings new insights about self-consciousness (Blanke and Metzinger, 2009), but also past and current models of consciousness (Apps and Tsakiris, 2014; Baars, 2005; Baars and Franklin, 2009; Clark, 2009; Crick and Koch, 2003; Dehaene et al., 2006; Hohwy, 2014; Limanowski and Blankenburg, 2013; Seth, 2014; 2015; Tononi, 2004), neuropsychology, computational neuroscience (Paradiso and Rudrauf, 2012; Philippi et al., 2012; Rudrauf, 2014) and the affective sciences (Damasio, 1999; Gray, 1990; Rudrauf et al., 2008; Scherer, 2009).

Our aim is not to try to solve the so-called hard problem of consciousness (Chalmers, 1995), but to focus on the structural phenomenology of experience and its representation in an overall computational framework capable of explaining complex behaviors, decision-making and action, based on internal parameters. We aim at establishing, based on general mathematical and computational principles, how we, as conscious, embodied systems, perceive, imagine and act in the 3-dimensional spatial world, from a first-person perspective, through a structural coupling with our environment, in order to cope with physical and social situations and maintain our autonomy by optimizing resilience (Friston, 2010; Rudrauf, 2014; Rudrauf et al., 2003; Varela, 1979). In agreement with other models, we conceive of the core of consciousness as an integrative global workspace (Baars, 2005; Baars and Franklin, 2009; Dehaene et al., 1998), selected through evolution for adaptive representation and control in an embodied context. We assume that consciousness participates in the orientation of cognition and action towards concrete and abstract objects in accordance with preferences; this can be expressed by the claim that consciousness is structured by an “intentional function” or intentionality (Freeman, 1999). More generally, conscious agents have to solve the problem of covertly or overtly (by explicit bodily actions) bringing into focus various parts of their surroundings for information integration via sensory evidence and simulation in order to navigate the world. Such intentionality can manifest at a psychological level, for example, as the so-called “attentional beam” (Baars, 2005). Our framework is also compatible with the general hypothesis that consciousness has evolved as a mechanism that optimizes information integration (Oizumi et al., 2014; Tononi, 2004), even though we do not formally relate our model to this framework here. Furthermore, we do not attempt to elucidate the relationship between so-called “access consciousness” and “phenomenal consciousness” (see Block, 1995) in the context of our model. Our view is that (1) phenomenal character (or qualitative character, as it is also sometimes called, e.g., of a tomato on a table, presenting the features of its exposed surface to conscious perception) normally correlates with (2) related intentional content (e.g., a tomato’s-being-on-the-table) and (3) accessibility (in the global workspace) for further actions or processing (e.g., I can report what I see and how it looks from my position, decide to pick up the tomato depending on my current narrative). These components may be dissociable to some extent, but one can think of our model as framing them as integrated in the overall structure of human conscious experience.

We place an emphasis on the embodied nature of consciousness because we see embodiment as a strong determinant of its overall structure and dynamics (Varela, 1979). However, our model does not reduce consciousness to embodiment in a strong sense, as the workspace relies on multiple layers of abstraction to represent and control the body, and the model is compatible with multiple forms of embodiment. Nevertheless, embodied cognition constitutes a framework within which a theory of consciousness can be derived in an ecological way. Autonomous embodied cogni-

tive systems can be described as situated agents that operate from a certain location in a specific environment under particular systemic constraints (Varela, 1979). They can further be described as performing perceptual and active inference in order to accumulate and deploy adaptive knowledge about themselves and the world through their history of interactions with it (Friston, 2010; Varela, 1979). In the framework of active inference, they carry out cycles of perception, prediction, and action according to prior beliefs about the self and the world held in their memory. These priors are constantly being updated and optimized, increasing the predictive power of internal models of physical and social exchanges with the world (e.g., energy efficiency, group acceptance). The internal models are used to appraise sensory evidence in relation to expectations and preferences and to govern action, including behaviors of approach and avoidance (Elliot, 2006; Gray, 1990; Rudrauf, 2014; Sander et al., 2003). The optimization of this overall process of appraisal will increase the agent’s likelihood of remaining within its domain of viability (homeostasis) and well-being, and help it cope with adverse events by maximizing its resilience (Kalisch et al., 2015; Rudrauf, 2014).

One of the pivotal considerations for the derivation of the PCM is that for such embodied systems in general, there is an evolutionary advantage of developing an integrative cognition of space in order to represent, simulate, appraise and control spatially distributed information and the consequences of actions. Geometries and spaces are important components of the distribution of information the agents’ brains need to determine in order to model relative changes in the internal and external environments. This need is so ingrained that sensory organs (e.g., the vestibular and visual systems) have evolved so as to embed implicit geometrical assumptions, and spatial cognition has undergone selective pressure throughout evolution to incorporate models of complex spatial relationships between agents in the context of the assessment of personally relevant outcomes (e.g., in prey and predator dynamics and luring strategies) (Cronin, 2005). The need for making inferences about space strongly suggests the existence of internal predictive models operating over sensory experience and prior beliefs in memory through which the world is represented and interpreted from a situated perspective.

Predictive models about spatial information can be expected across the hierarchy of neurocognitive processing, including non-conscious levels. But we hypothesize that, at the tip of the iceberg, consciousness relies on an integrative predictive model attributing a global 3-dimensional spatial geometry to multimodal sensory information and memory traces as they access the conscious workspace. We also hypothesize that this model serves as a support for the mapping and appraisal of preferences and personally relevant information. We expect that such a spatial model offers a parsimonious, minimally complex mechanism, which is consistent with basic principles of approximate Bayesian inference (and free energy minimization) but also with Occam’s Razor, which governs (statistical or generative) models that optimally generalize to new data (Friston, 2010). We hypothesize that such global internal model of space operates in a supramodal manner by dynamically integrating and unifying information from multiple sensory modalities through affective arbitration (Rudrauf, 2014).

Of course, there is much more to consciousness than spatiality and its valuation. In a normal human individual, consciousness extends to incorporate differentiated and sophisticated levels of representation (Damasio, 1999): from perception to language and mathematical thought. And lived spatiality must be thought of as integrated with the flow of lived time (Varela, 1999). But, to begin formulating a mathematical model of consciousness, it is sound from a phenomenological, functional, and practical standpoint to consider core consciousness (Damasio, 1999) from the perspective

of embodied spatial cognition in the context of active inference, which incorporates complex cognitive dynamics.

Our approach directly connects to a fast-growing literature relating the free energy principle (Friston, 2010) to perceptual inference, active inference and (embodied) conscious experience. The framework of free energy driven perceptual and active inference has been used to shed light on conscious processing, including the structure and unity of consciousness, aspects of the first-person perspective, self-models, embodiment, interoception, body illusions, and bistability (Apps and Tsakiris, 2014; Clark, 2009; Hohwy, 2014; Limanowski and Blankenburg, 2013; Seth, 2014, 2015). The PCM advances previous formulations of active inference (in relation to consciousness) by featuring a mathematical theory, based on projective geometry, of the form, structure and dynamics of spatial, perceptual, and imaginary experience. In particular, it accommodates counterfactual or “as if” multi-perspectival transformations of the first-person perspective in conscious experience. It specifies a core mechanism unifying perception, imagination and action for the optimization of the appraisal of action outcomes and thus maximizes the overall resilience of an embodied cognitive system imbued with it. It offers an explicit, formal, computable and integrative account of embodied consciousness, which links components of the affective sciences, and offers a basis for testing hypotheses about psychological mechanisms quantitatively.

Our argument follows three basic steps:

First, we assume that conscious processing entails a process of inference, in particular, active inference under the free energy principle. It relates the conscious intentional experience of the world with the products of typically unconscious inference based on sensory evidence.

The second step follows from the first assumption: to make inferences or predictions any sentient (or conscious) creature must possess a generative model of its sensorium. In other words, to engage in active inference it is necessary to generate predictions about what would happen if one did such-and-such. This is necessary for testing internal hypotheses against sensory evidence and updating one’s beliefs (or behavior) accordingly. This fits comfortably with the Bayesian brain hypothesis (Clark, 2013; Knill and Pouget, 2004; Yuille and Kersten, 2006) and more fundamental arguments such as the Good Regulator Theorem (from early formulations of self-organization) (Conant and Ross Ashby, 1970; Seth, 2014).

Finally, and most importantly, we posit projective geometry as a necessary form of the generative model that must be deployed by any conscious organism that navigates a three-dimensional world. Put simply, to produce generalizable and veridical predictions, the internal models used by our brain must recapitulate the formal geometry of the world we are trying to make predictions about (from situated standpoints). In this setting, the Projective Consciousness Model (PCM) regards projective geometry (in three dimensions) as the optimal geometry for situated inferences to the best explanation regarding sensory impressions.

In what follows, we will see that the ability to operate on or within a projective geometry carries with it all the necessary equipment for the sort of active inference characteristic of beings like ourselves, in particular, the ability to entertain fictive outcomes and multiple perspectives on the composition and dynamics of objects causing (or that could cause) sensory input. We focus here on the projective aspects of generative models and the pre-eminent role of projective geometry (as opposed to active inference and free energy per se). We suppose that such a geometry must (at some level) be entailed by the generative models used by the brain (see the section below on neurocomputational hypotheses).

1.2. Phenomenological postulates: a mode of access with a characteristic spatial phenomenology based on perspective taking and driven by the mapping of personally relevant information

We specify how spatiality is consciously experienced, including one’s own body and oneself as observer (Blanke and Metzinger, 2009). Our goal is to infer from this analysis, moving from what is represented to what carries out the representing, key computational properties of consciousness that can account for the generation of its characteristic spatial phenomenology. This will later allow us to make hypotheses about its neural implementation.

To ground the model, we use the following well-motivated postulates, which distill centuries of discussion about the structure of subjectivity and consciousness. The postulates feature an integrative subjective experience with a characteristic spatial phenomenology that is accessed from a specific perspective (a “point of view”). They were identified prior to any particular mathematical model from a synthesis of psychology, cognitive neuroscience, phenomenology, and philosophy of mind. Though not uncontroversial, we accept them here as a plausible starting point from which to draw out important consequences. (A similar method can be found in Oizumi et al., 2014.)

1.2.1. A 3-dimensional lived space

Investigations into the phenomenology of consciousness have emphasized that consciousness is grounded in a subjective standpoint and a space of sorts, using the basic schematics of a subject-object opposition and a relation between a center and a periphery (Kant, 1999; Husserl, 1913; Merleau-Ponty, 1945; Varela, 1999). Spatiality is a pervasive aspect of subjective experience. When one closes one’s eyes in a silent room, under normal conditions, one experiences a “lived space”, to adopt a term of art commonly used in the Phenomenological tradition, the sense of a 3-dimensional surrounding space that is invariant even as the state of the outside world constantly changes. We find such a space in both waking consciousness and in dream. It seems central to conscious experience in general, to the way it represents a world in three dimensions, in which we locate ourselves and from which we perceive, imagine, and plan action (Peirce, 1955; Gurwitsch, 1964; Berthoz, 1997; Blanke and Metzinger, 2009; Koenderink and van Doorn, 2008; Legrand et al., 2007; Lenggenhager et al., 2009). In this context, the particular “way” such a space appears is considered as an immanent structure of the subjective conscious experience we “live through”, and to which we have phenomenological access; it is not a feature of the world represented (Williford, 2015) (it only secondarily becomes itself an object of representation, something to which our attention can be drawn, just as one does not normally notice one’s hand in hammering, though one can). As a first approximation, lived space can be said to feature, by default, a primarily egocentric representation of space, defining a peri-personal and an extra-personal environment in relation to a body, and oriented along directions of perception and action (Berthoz, 1997; Legrand et al., 2007).

It manifests as a unified, supramodal and normally coherent experience, integrating and mapping different portions of space into a cohesive and continuous global workspace, as sampled by multiple, redundant sense modalities (notably audition and proprioception), and constantly filled with information from memory to complement the partial sensory sampling (e.g., I’m looking at a stage and listening to the music playing, but I know and feel that my friend is sitting next to me). Vision is central to the process, but it is not necessary to it (e.g., congenitally blind subjects also have a rich perception and imagination of objects and of others in space). The visual field is an informationally rich and sharp representation of a portion of the ambient space, which is integrated into the broader, more encompassing conscious experience of a global

spatial field with a unified frame of reference and a representation of the body (Berthoz, 1997; Koenderink and van Doorn, 2008; McCormack, 2006). Thus all sense modalities, vision, audition, proprioception, vestibular processing, olfaction, touch, but also interoception are expected to contribute to its ongoing construction in real-time, and to populate its content adaptively (Berthoz, 1997; McCormack, 2006).

As suggested above, lived space ought not be considered a direct representation of the intrinsic spatial structure of the world, and can be conceived as an adaptive simulation, “user interface”, projection or controlled “hallucination” that is actively constructed by the brain in real-time (Douglas, 2007; Friston, 2010; Koenderink and van Doorn, 2008). Experientially, outside of pathological conditions, it engenders an illusion of direct realism, (i.e., of being conscious of the world itself directly) (Metzinger, 2004), which is probably a necessary condition for effectively representing our real environment and acting within it. But dreams, imagination, hallucinations, out of body experiences (OBEs), sensory substitution, vertigo, and various neurological syndromes, provide compelling evidence that our lived, moment-by-moment representation and subjective experience of ambient physical space is the construction of a virtual world. In fact, the normal, robust, ongoing conscious experience of co-localization in space of one’s subjective point of view and of one’s own body is contingent upon a process of multi-sensory inference, and reflects a dynamical balance in sensorimotor calibration (Blanke, 2012; Blanke and Metzinger, 2009). Accordingly, when we use terms like ‘embodied’, this should not be taken to commit us to direct realist, externalist, or “extended-mind” theories of conscious perception (see, for example do Nascimento and Noë, 2012). In our view, consciousness is a computational process performed entirely within one’s brain; even though, clearly, it bears causal connections to processes that extend well beyond the brain and, by intentionality, it represents objects, events, and states of affairs that are certainly not contained within the brain.

1.2.2. An elusive standpoint and a space with a representation of the body at the center

In default egocentric spatial consciousness, the center of lived space seems to coincide with an implicit vantage point, somewhat near or intersecting with the perceived location of the head. Phenomenologically, the vantage point has an ill-defined and elusive localization in the space, even though it seems to play a central, if usually tacit role, like that of an origin, in the organization of the space: one can locate objects in space, but one cannot precisely locate the vantage point from which the representation seems to emerge as a simple location in 3-dimensional space (Williford et al., 2012). It seems to behave like a virtual, internal 3-dimensional camera, which is reminiscent of the long-debated notion of an “inner eye” and the basis of some classical aporias of the theory of consciousness, (e.g., the homuncular regress problem, the problem of the location of the subject, and issues surrounding the “Cartesian Theater”) (Butchvarov, 1979, 1998; Dennett, 1991; Williford et al., 2012).

Philosophers have debated over centuries about whether one can find a “subject” in consciousness (e.g., Hume, 1739; Wittgenstein, 1922; Sartre, 1937). Such a subject would, among other things, provide an origin and foundation for its lived space. The debate has oscillated between affirming and denying its presence (Williford, 2010; Williford et al., 2012). If it exists, it remains phenomenologically elusive and cannot be accessed like normal objects. However, the intuition is pervasive and difficult to eliminate, and we wish to investigate whether we can make sense of it with a mathematical model, bracketing ontological and epistemological concerns.

Lived space is also characterized by the pervasive presence of a 3-dimensional representation of the body, which monitors the

state and controls the action of the actual body. The existence of such a lived representation of the body is supported by evidence from perturbations of proprioception in various experimental and clinical conditions (Blanke and Metzinger, 2009), by the pervasive presence of phantom limbs in amputated patients (which the patients can control to a certain extent as they would control the missing limb (Giummarra and Moseley, 2011)), and by ordinary mental imagery.

In contrast to most contents of consciousness, the lived body is normally always present in the conscious field, though, of course, not normally the object of explicit attention or conceptual representation. Its integrity, as a proxy for the integrity of the actual body, is critical for an embodied autonomous agent. As a space imbued with value, it is an anchor point for our efforts at preserving autonomy and well-being. Protecting the actual body in physical space requires an ongoing monitoring and appraisal of one’s own bodily position and state in space in relation to past, present, and future situations. The lived body is thus a kind of inferential representation of the real body in physical space, reliant upon memory, anticipation, and ongoing multimodal sensorimotor information. It constitutes a sort of virtual “user interface” for the representation and control of the actual body, just like the rest of the global workspace. It can be thought of as a kind of “virtual body avatar”, richly yet selectively informed by multiple sources of information (Blanke and Metzinger, 2009). We control our limbs through anticipated proprioceptive representations of their position in space, and we avoid and approach things based on their anticipated potential personal relevance in connection to their impact on the body.

In normal perception, the lived body manifests as a 3-dimensional phenomenon in a default egocentric mode, that is, from a first-person perspective. The representation is informed by proprioception and interoception and by other sense modalities, like vision, audition, and the vestibular sense, but also olfaction; all of these act as spatial filters. The lived body and the real body are related: their referentials are “superimposed” via ad hoc calibration for optimal representation and motor control. This implicit superimposition is dramatically revealed by phantom limb phenomena.

In imagination, or in abnormal perceptual situations such as body-swap illusions or normal effects of mirror reflections (Blanke, 2012; Blanke and Metzinger, 2009), the body also manifests as a similar 3-dimensional phenomenon, under various egocentric and allocentric modes, from first-person or third-person perspectives. In spite of their frequent relative separation in lived space during imagination, experimental situations, or the confrontation with reflections, both the implicit vantage point and the representation of the body generally remain identified as one’s own; they preserve their identification as belonging to the same agent in consciousness.

Such lived representation of one’s own body is present in many forms of subjective experience involving an explicit and reflective form of self-consciousness. It is present in mind-wandering, day-dreaming, and fantasy (Smallwood and Schooler, 2006). It is constantly cast, with more or less emphasis, into the ongoing narratives and scenarios that we unfold as subjects through imagination in order to aid us in appraising, interpreting, and manipulating physical and social situations so as to maximize utility, in a normally highly biased way.

1.2.3. A mechanism for spatial intentionality and perspective taking

Lived space necessarily involves spatial intentionality (i.e., the ability to relate to objects as a situated observer and to aim at locations in space) and intentionality as an organizing function is central to consciousness (Freeman, 1999). There is a sense that one’s awareness of objects goes out from a subject toward an object (Husserl, 1913). This is in evidence in spatial selective attention and is described as an “attentional spotlight” in the global

workspace model (Baars, 2005) (which invokes the metaphors of a 3-dimensional theater with an implicit observer and a stage). In the space of consciousness, points are not without valence or relevance; they are the focus of “beams” of attention or the sources of such “beams”. This spatial intentionality is independent of overt actions or orientation (e.g., head orientation and saccades). It can operate covertly or overtly across perception, imagination, action planning and self-representation, and acts as a principle of unification among these processes. It is key to the control of behavior (Schöne, 1984). We should add that in speaking of ‘intentionality’ here we have in mind the general feature of “object-directedness” that is presupposed by all higher forms of representation. We are suggesting that this “directedness” or “aim” is a fundamentally geometrical notion that, of course, develops in conceptual thought and language well beyond its roots in perceptual consciousness.

The dynamics of intentionality and orientation in lived space are controlled by a complex ensemble of parameters, which relate to biological constraints and the ongoing appraisal of the cognitive and affective personal relevance of objects, spaces and situations (e.g., we pay attention to what we fear or desire and where it is to be found) (Sander et al., 2003). Affective processes, in connection to memory, have long been conceived as playing a key role in driving intentionality and consciousness (Husserl, 1913; Damasio, 1999; Rudrauf et al., 2003; Varela, 1999).

Perspective taking, the act of representing space and its contents from a particular standpoint, is central to the overall process of comporting oneself in lived space by representing physical and social relationships from different angles (Lamm et al., 2007). It implements the spatial expression of intentionality in attention. It is itself driven by the detection of personally relevant information in connection to memory. Perspective taking is the common mechanism by which the frames of reference of perception, imagination, action-programming, and self-representation are related. In perception, the frame of reference of lived space tends to default to an egocentric stance, with a first-person perspective. In imagination, we can adopt any allocentric vantage point: one can imagine something from somebody else’s point of view or an object from a first-person perspective, or oneself being represented from a second-person perspective, or oneself relating to another or an object from a third-person perspective. We also can imagine the points of view of non-human entities, from tiny animals to a “God”, and from any distance (e.g., “at arm’s length”, “with a ten-foot pole”, “from 30,000 feet”).

This implies that possible other perspectives are, in effect, built into lived space. Evolutionary pressure has been placed on animals to represent themselves from the perspective of another (Humphrey, 2003) e.g., the prey as seen by the predator and conversely, which has enabled animals to develop luring strategies based on anticipation, and to create and exploit perspectival illusions for mating success (Kelley and Endler, 2012). In our model, access to representations of space in consciousness always implies perspective taking. Creatures with no capabilities for third-person point-of-view perspective taking would have a very limited consciousness, as they would be unable to use anticipatory local and non-local perspective taking to build an integrative representation of their surroundings; thereby escaping their situated first-person perspective.

1.2.4. A workspace in perspective with a horizon

In both normal perception and imagination, the experience of spatiality generally follows specific rules of transformation in which the sense of direction (front, back, up, down, left, right) and of relations of incidence are preserved. The alignment of objects along a straight line is preserved in changes of point of view (e.g., when perceiving or imagining roads under different points of view). Likewise, when seeing or imagining an object moving, we

normally perceive the incidences of its parts as being preserved; otherwise the object would lose its unity. From an evolutionary standpoint, representing objects in a way that preserves relations of incidence is a critical adaptive feature in a world of inertia, collisions, and encounters. Relations of incidence are personally relevant for embodied systems.

In lived space, objects, points of view and their relations of incidence are represented with the shape of 3-dimensional projections in perspective, along certain directions, from a certain standpoint; parallel lines converge on an implicit or explicit horizon at infinity. It is through the lens of this inherently perspectival geometry that a world of transcendent 3-dimensional objects in Euclidean space existing in independence of any particular perspectives (whose sizes remain the same irrespective of distance) is consciously experienced. This is familiar enough: in one sense, people “look tiny” from the top of a tower, but one does not think that they have shrunk; railroad tracks “seem to” converge at the horizon, but one knows they don’t. In this space of appearances, profiles, adumbrations and perspectives, the same objects can be said to grow, shrink, vanish, occlude one another, be “visible-from-here”, “inaudible-from-there”, and so on. We hold that this perspectival mode of representation is true of both perception and imagination (for both concrete and abstract objects).

Perspectival properties are most obvious in vision (although the extent to which perspectival properties dominate visual awareness is debated (Koenderink et al., 2010)), in which the horizon corresponds to the set of vanishing points along lines of perspective, e.g., as parallel railroad tracks converge toward a vanishing point. But as noted, lived space is not reducible to the space of vision. It is a multisensory, integrated phenomenon, incorporating the visual field into a larger space and feeding on sources of spatial sensory sampling other than visual ones. It integrates both perception and imagination into a common space of inference and is strongly related to embodied cognition, in which prehension and locomotion (i.e., actions) are as much central organising factors as perception or imaginary representations.

One can thus expect that imagination also operates within a perspectival frame of reference, which can be mapped onto the visual field or more generally onto the perceived surrounding space. For instance, when one imagines a cube moving away anywhere around oneself, the imagined cube appears in imagination to become smaller and smaller as it moves farther away, and it eventually vanishes into a point as it approaches “infinity”. This geometrical behavior of the cube is associated with a sense of depth.

Since objects vanish near infinity in all possible directions of lived space, lived space, as a perspectival space, appears from inside as bounded by a horizon at infinity, even though the horizon may be set very far in the virtual space of consciousness and remain outside or at the fringes of the scope of attention. We hold that everything present in spatial consciousness is represented in perspective in relation to this movable horizon no matter its relative position in the space around us.

Beyond phenomenological considerations, it is important to emphasize that such a configuration is integral to the mechanism of perspective taking, as it frames all subjective representations of space along specific directions of aiming from specific vantage points. The perspectival structure of the lived space that results is part of the implementation as well as an expression of intentionality in consciousness.

2. Model

2.1. General rationale

We engage in cognition and act in ambient space in reference to a 3-dimensional Euclidean space (the ambient physical space may

or may not be a 3-dimensional Euclidean space, but assuming a modest realism, it approximates such a space at the scale relevant to most of our practical pursuits). Starting from the vital necessity for biological systems to distinguish between changes in the world resulting from their own motions versus from outside transformations, Poincaré showed the importance of accounting for the relativity of rigid motion between the body and the world in perception and action programming, which can be made on the basis of the group of displacements in Euclidean 3-dimensional space (Poincaré, 2015). This particular group appears to structure empirical facts, outside of subjective experience. The action of the group, supported by the Euclidean space, correlates with the behavior of physical space and of fields of matter and light under relative motion.

However, physical space does not coincide with lived space as experienced in consciousness. The phenomenological postulates previously reviewed have led us to conclude that the geometry that characterizes lived space is not Euclidean: it is fundamentally perspectival and represents objects in ambient space in a manner that is best explained by projective geometry with a direction of aim and a scope (e.g., “from here” or “over there”, or “from here to there”; see Greenberg, 1993 for a historical introduction to projective geometry). It follows that the 3-dimensional characteristic phenomenology of the lived space is framed by a 3-dimensional projective space and that the mechanism of perspective taking corresponds to the action of projective transformations. In fact, the group of projective transformations contains the Euclidean group, thus projective geometry embeds Poincaré’s transformations.

We are proposing that the generative models used by conscious creatures are necessarily equipped with a projective geometry. Note that the generative models in play cover all aspects of how sensory data are generated. This means that projective geometry constitutes a necessary but not sufficient component of these generative models. Having said this, we consider projective geometry to be quite fundamental for conscious processing in the following sense: we associate conscious processing with inference about the consequences of action. This sort of inference is necessarily about the future and therefore acquires a temporal thickness (Wilkinson, 2010). Crucially, it is the sort of inference that answers questions like “What would happen if I did that?” The fact that we do things in a metric space therefore requires counterfactual inference about the future to conform to a projective geometry. From the perspective of active inference, every (fictive) perspective or view we take on a scene represents a particular movement or possible action, from which we have to select the best. This selection constitutes our intention.

Active inference is a corollary of the free energy principle, which subsumes approximate Bayesian inference and model selection. Common examples of this in the neurosciences include predictive coding (Bastos et al., 2012; Rao and Ballard, 1999) and Bayes optimal decision making under Markov decision processes (Friston et al., 2016).

The PCM incorporates the free energy principle, which mandates that the best perspective is that which results in the smallest free energy or expected surprise. Because surprise is (in information theory) self-information, expected surprise is then entropy or uncertainty. This means that the best among all plausible or allowable perspectives is the perspective that resolves the greatest amount of uncertainty, thereby providing a formal basis for action selection. Perhaps the simplest example of this is deciding where to look next. The selection out of all possible saccadic eye movements can be easily modeled in terms of minimizing expected free energy or uncertainty in terms of salience; either in continuous state space models (Friston et al., 2012) or discrete state space models (Friston et al., 2016, 2015; Mirza et al., 2016).

Crucially, expected free energy can always be decomposed into epistemic and pragmatic parts (Friston et al., 2015). The uncertainty reducing component is known as Bayesian surprise, epistemic value, information gain or salience (Itti and Baldi, 2009). The pragmatic part corresponds to expected surprise under prior beliefs about outcomes; i.e., prior cost. This means that the active inference gracefully subsumes expected utility theory and related (reinforcement) formulations in behavioral psychology.

However, note that neither expected utility theory nor reinforcement learning theory is equipped with the epistemic or uncertainty resolving aspects that are a natural consequence of active inference, under the free energy principle. We will not develop the mathematical details behind this aspect of active inference. For details see Friston, 2010 and Friston et al., 2015. Our key point is that the machinery that is necessary to evaluate the expected free energy, under a plausible action or transformation, must necessarily conform to a projective geometry.

The integration of a global projective geometrical level of representation and control into a free energy minimization scheme for multisensory inference in relation to preferences must (we assume) have been shaped through a long process of evolutionary selection to support the generation of resilient autonomous agents. In this treatment, we have not considered the learning of spatial transformations. In principle, this is usually subsumed under active inference, where the parameters encoding allowable transformations within the generative model are learned so that the (long-term average) of free energy is minimized. This usually translates into Hebbian or associative plasticity; for both continuous and discrete state space models (Friston, 2008). There are supraordinate issues which we do not address here that concern the epigenetic specification and learning (i.e. Bayesian model selection) over evolutionary time (i.e., natural selection) (Campbell, 2016).

The explanatory work of projective geometry in the context of active inference will be illustrated using a simple example below (with simulations). At present, we just need to note that to select actions and navigate in a metric space, it is necessary to implement a projective geometry to evaluate the consequences of (fictive) narratives. As noted above, an important aspect of this selection rests upon the salience or epistemic affordance associated with each transformation entertained.

2.2. Core features of the PCM

This overall framework can be modelled by integrating in a hierarchical manner: (1) a projective geometry engine, which is embedded into (2) an active inference engine; the two being coupled in order to access and evaluate information and control action (Rudrauf, 2014) (Fig. 1). The projective geometry engine implements a specific kind of geometry, projective geometry, that: (i) is based on perspective taking, (ii) plays a central role in spatial inference, and (iii) is capable of integrating perception, imagination and action-programming into a global operational workspace with seamless relationships between frames of reference. This active inference engine integrates the geometrical model into the larger framework of free energy driven, active inference, which is used as a general model of enactive, embodied cognition (Friston, 2010; Rudrauf et al., 2003; Varela, 1979). As part of the overall process of active inference, intentionality, which includes the driving of perspective taking and the motivation of action, is driven by the minimization of free energy, which directly relates to conditional probabilities in memory. These probabilities can encode a variety of expectations, including cognitive and affective preferences, as a function of locations in space. Sensory evidence contributes to the updating of prior beliefs in memory as the system interacts with its surrounding.

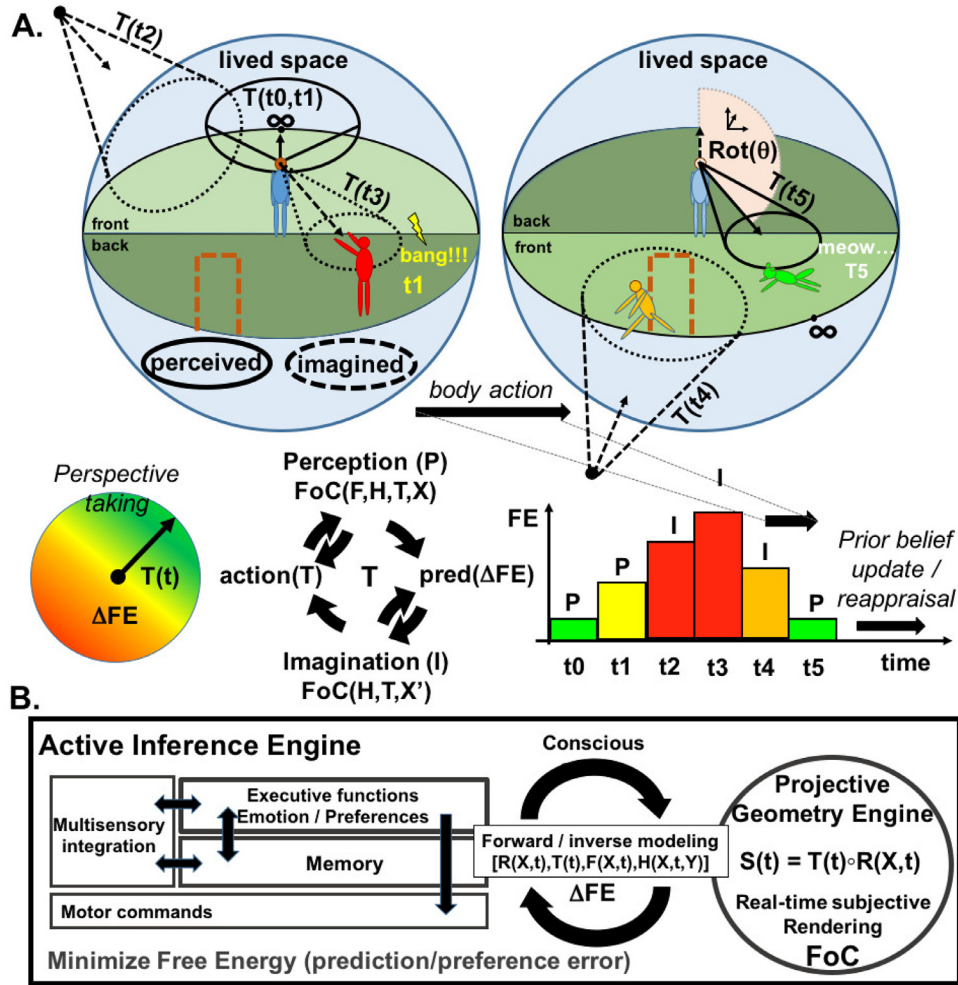


Fig. 1. Overall sketch of the model. A. Lived space S (blue spheres): subjective experience of space and its contents as perceived (continuous funnels on the figure) or imagined (dashed funnels). It is framed in a 3-dimensional projective representation in perspective of a world model $R(X, t)$ associated with a distribution of free energy (FE) in space, related to cognitive and affective personal relevance (colors on the manikins and door, and disk with color gradient). Perspective taking and the selection of corresponding projective transformations T are governed by the gradient of free energy across space δFE . The diagram with the arrows in circle represents the possible transitions of state: from perception to imagination and action, for the Field of Consciousness (FoC), which is partially driven by sensory evidence when it is focused on perception of the local environment, and entirely driven by prior beliefs in memory and simulation capabilities when it is focused on imagination of non-local spaces (e.g. when we imagine ourselves at home while we are at work, in the present, past or future), which cannot be based on direct sensory evidence. The FE defining the optimal perspective at a given instant is globally minimized, through cycles of perception, imagination and action, across time. The bar graph represents FE, as a function of time and of the different perspectives and processing modes (P: perception versus I: (non-local) imagination) adopted by the system (see text). B. Overall diagram of the active inference engine, embedding the projective geometry engine for subjective rendering, confronting prior beliefs with multisensory information and driving learning and action. The overall process minimizes free energy. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Projective geometry corresponds to a specific mode of access to information that is formatted by non-trivial projective transformations. It provides the principles by which a variety of representations or models of the spatial-temporal world, including the body, $R(X, t)$, which are stored in memory, can be accessed and mapped in the workspace from different points of view in an integrative manner. The process is based on the application of projective transformations $T(t)$, which implement perspective taking, such that the subjective experience of space at a given instant corresponds to:

$$S(X, t) = T(t) \circ R(X, t) \tag{1}$$

with \circ denoting the application, similar to an operation of composition, of the projective transformation T to the spatial world model R , yielding a space equipped with intentional directions in perspective, featuring observation beams and relating subjective perspectives of possible observers.

More specifically, $R(X, t)$ denotes the state at time t of a set X of internal variables that describe the environment, external world (for instance the objects in it, but also the fields on it, such as weather, temperature, light, and so on), and our own body. The state R depends on other variables Y , such as sensory inputs associated with sensory data D . And $S(X, t)$ denotes another state at time t of X that represents our subjective experience of the external world (objects in it and fields over it), and of our own body. The geometrical hypothesis tells that projective transformations and generalized transformations (for instance the perspective projections, as further described below), act as a set of transformations in the manifold of states of X , sending an unconscious state to a state that is accessible to consciousness.

Moreover, the geometrical hypothesis requires that the full group G of invertible projective transformations, such as changes of projective frames, is acting on the subjective transformations in the following sense: assume the state R is transformed in $S = T \circ R$ by an element T (that can be invertible or not), the replacement of T by the conjugated element $T_1 = g_1 \circ T = g_1 T g_1^{-1}$, where

g_1 is an invertible element of G , must satisfy the covariant rule: $g_1 \circ (g_2 \circ T) = (g_1 g_2) \circ T$.

This means that conscious states are not very different from unconscious ones: they result from taking a point of view or perspective using projective transformations. The choice of the projective operation $T(t)$ transforming $R(X, t)$ into $S(X, t)$ can therefore be assumed to rely on unconscious processing (it can be complex, as we can expect the system to combine several such transformations across representations of the self, other persons and objects).

A more precise symbolization, but also probably more cumbersome at this stage, should take into account two sets of variables, X for variables that are almost conscious, and Y for hidden variables that participate in the definition of X (e.g., related to sensory systems that are far from conscious access). Then the equation should be $S(X, t) = C(T(t), R(X, Y, t))$, satisfying, for any g in G , and state R of X and Y , the equation $C(g \circ T, g \circ R) = g \circ C(S, R)$. Importantly, the states of variables X and Y have to be understood as estimates of probabilities, which in some cases can be certainties, yielding fixed values.

Such a projective mode of access presents strong computational advantages and can seamlessly relate the spatial frames of perception, imagination and action programming. This model of space constitutes the spatial field of consciousness in our theory.

This spatial field of consciousness (FoC) is controlled by prior beliefs in memory and ongoing sensory evidence to serve active inference through the minimization of free energy. The process results in an algorithm of arbitration and learning that optimizes the projective representation of information (morphological or affective) and the orientation of intentionality (spatial attention and action programming) within the workspace, specifically: (1) adaptive spatial inference (physical and social) to build spatial knowledge in memory (about shapes, structures, behaviors); and (2) ‘personal relevance’ as mapped onto spatial knowledge in memory. The embodied agent minimizes free energy, as related to the discrepancy between beliefs and evidence, and thus optimizes its representation of spatial information and control of action.

Free energy (FE) is a mathematical way of expressing the surprise or improbability of some sensory impressions under a particular (generative) model of how those data were generated. In the PCM, free energy provides an upper bound on surprise or self-information such that:

$$FE > -\log P(D|R, T, \dots) \quad (2)$$

This (simple) formulation shows that FE is always greater than surprise. Therefore minimizing free energy (or selecting transformations T that minimize expected free energy) will always reduce surprise (or resolve uncertainty). Here, D are sensory data that are available to the active inference engine. In other words, we associate any possible transformation with a model that, in practice, would become a hypothesis about “what would happen if I did that”. As noted above, this enables the expected free energy to be associated with any transformation (or its products); where the products S may also be equipped with prior preferences (i.e., $\log P(S)$) of a particular outcome S that has low cost or surprise).

The relationship between free-energy driven active inference and projective geometry in the PCM can be more specifically formulated as follows. FE minimization is a general expression of Bayesian inference that incorporates constraints on computational (and statistical) efficiency: as with every Bayesian approach, it relies on a choice of variables to estimate and use internally, and on relations between them. The form of these variables and relations is crucial. A specific geometry furnishes a way to make precise and organize this form; in fact a specific geometry defines particular observations to be made and particular rules prescribing how the observations and priors have to be analyzed. Here we can refer to

Poincaré’s analysis of the most convenient geometry for managing the ambiguous relations of individuals with the external world. In the specific case of consciousness, the most convenient geometry is a projective one, notably because projective geometry is the natural geometry of arrays of directions from any point to all other points; it relates “me” observing and “me” being observed together, taking into account several frames of observation from places to places.

Moreover, free energy minimization offers a highly flexible, fully computable framework that can be associated with concrete algorithms for actually performing quantitative predictions and advanced simulations of agents implementing the PCM, including in guiding perception, action and decisions. FE offers a solution to the problem of implementation in addition to having a strong theoretical back story. The behavior of the model reproduces psychological descriptions invoking states of affairs of personal relevance (rewards) and utility. It allows us to understand and manipulate psychological dimensions computationally, based on a general principle of optimization of resilience, in order to test hypotheses about the role of cognitive and affective processes in behavior. It also allows us to make connections with brain function and more generally shows how geometry can be compatible with probabilistic algorithms for learning and decision.

2.3. Incorporation of models of appraisal and motivation

The free energy framework combined with projective geometry conveniently enables us to explicitly incorporate appraisal and motivation theory by setting prior beliefs in memory and rules of approach and avoidance (Damasio, 1999; Gray, 1990; Sander et al., 2003). We can thus specify the mechanics of embodied consciousness and characterize perception, imagination and action relative to affective and motivational models, which can be related to given adaptive goals of survival and well-being and compared based on simulation outcomes. We place emphasis on these affective aspects because they are essential psychological and ethological concepts that a unified computational framework proposed as a model of embodied consciousness should be able to incorporate.

We can interpret the free energy framework, under some specifications, as implementing an intuitive notion of “personal relevance”. The notion subsumes innate and acquired abstract and concrete “values” or preferences, yielding cognitive and affective saliences, as built into the prior beliefs about the world of a given system. In this context, personal relevance and motivational valence are scored by the expected free energy under a particular transformation (entertained by the projective geometry within the generative model). As noted above, the motivational valence or value is encoded in terms of prior beliefs about (counterfactual) outcomes S , while the personal relevance of those outcomes corresponds to the salience or epistemic value of reducing uncertainty about how these outcomes are obtained, i.e., epistemic affordance. Clearly, the salience or epistemic part of expected free energy is inherently personal, because it corresponds to the particular viewpoint or perspective implicit in each possible transformation. We operationalize personal relevance as the negative logarithm of a field of conditional probabilities that represent prior beliefs, namely, events with a given affective and motivational valence for the agent or “value” will occur across a given portion of space. This constitutes a belief system or hypotheses about the world stored in the agent’s generative model. Through active inference the agent confronts its beliefs with a given factual world, which corresponds to sensory evidence that can be experienced by the agent as a result of its sensory coupling with the world. Beyond appraisal and motivation, notions of motor commands and sensorimotor forward and inverse modeling can be naturally incorporated into the free energy framework and become integral to

the design and understanding of specific systems performing active inference. In active inference, motor commands become descending predictions of proprioceptive (or interoceptive) signals, which elicit motor (or autonomic) reflexes (Adams et al., 2013; Barrett and Simmons, 2015; Seth and Friston, 2016). By the same token, corollary discharge becomes the corresponding descending predictions from the generative model to exteroceptive modalities (e.g., what things would look like if I moved there) Mirza et al. (2016). The PCM can thus perform active inference within the real (e.g., through robotics) or simulated (e.g., through virtual reality) environments that define a space of interaction and a world model.

2.4. The Field of Consciousness in the PCM

The minimization of free energy drives perspective taking by setting the parameters of the projective transformation T based on a combination of cognitive and affective parameters related to the expectations of the system. This process plays a central role in intentionality and the motivation of action. For instance, typically, we want to move away from that scary thing over there we believe to be dangerous, we want to move closer to one who attracts us, and we stay put to continue what we are enjoying as long as other motives do not interfere. In the process we sample sensory evidence and revise our prior beliefs accordingly, and consciousness helps to transform implicit computations into explicit ones that are sensitive to observations.

In the PCM framework, under the minimization of free energy, the FoC is not only a spatial field, but also an integrated “force field”. The agent is governed by a dynamical system that drives the mechanics of its embodied consciousness through perspective taking across virtual and real spaces in order to maximize spatial knowledge and personal relevance and control action. We note that our model does not offer an explanation as to why the force field is phenomenally conscious; it limits itself to addressing why consciousness has the shape and dynamics of such force fields. The spatial gradient of free energy corresponds to the spatial gradient of preferences.

The overall FoC at each instant, including the agent’s intentionality (i.e., the direction of its aim), is driven by the dynamic interplay of ongoing multisensory information, memory and affective biases and serves to appraise and decide on the adequate course of action, based on accumulated, anticipated and observed information. By combining spatial and affective priors under the free energy principle, we can specify the principle of action that governs embodied intentionality and actual action in our model in a way, for instance, that minimizes the likelihood of unexpected and adverse events and maximizes positive outcomes for the embodied system. The framework allows us to encode space (including shapes) and affective evaluation together so that active inference relies on an “aesthetic” process, so to speak, which is consistent with the strong link between spatial and emotional memory in neuroscience (Varela, 1999).

When we associate transformations with an attentional beam, we are appealing to the notion of attentional salience and to the epistemic value entailed by minimizing expected free energy (Mirza et al., 2016). This should not be confused with the perceptual attention to ascending sensory streams; usually cast in terms of precision weighting in predictive coding formulations of active inference (Feldman and Friston, 2010). This means the attentional beam corresponds to the epistemic affordance that characterizes a particular transformation and ensuing perspective on the world.

2.5. Illustration of the PCM

Fig. 1 illustrates the general properties of the PCM from a phenomenological standpoint. The subjective experience of space and

its contents as perceived (continuous funnels on the figure) or imagined (dashed funnels), i.e., lived space S for short, are framed in a 3-dimensional projective space (blue spheres) in perspective, through a mechanism of perspective taking. Personal relevance (as associated with sensory qualities and feelings in conscious experience) is represented in the colors on the manikins. Perspective taking is governed by projective transformation T , which is applied to the current world-model R in memory, under the influence of executive functions, emotions, and ongoing multisensory integration. The minimization of free energy drives the dynamics of the system, including the sequence of perceived or imagined projective scenes (i.e., little “mind movies”) that plays out in its consciousness in an emotionally charged context. The overall process is implemented through the coupling of (1) a real-time projective geometry engine, performing subjective framing and rendering, and (2) an active inference engine, appraising perceived and imagined projected situations based on internal models and prior beliefs. The body is represented and experienced as a central reference in the space, and the visual field is only a part of lived space.

A hypothetical scenario can be used to illustrate the basic operation of our model (see Fig. 1). A loud explosive noise from behind captures the attention of a man governed by the model (blue manikin). He stands in a living room with an open door in the back (as an accessible piece of knowledge $R(X)$ stored in his memory)(Fig. 1-a left-tier)($t_0 - t_1$). Alarmed by this stressor (free energy is increased), he starts imagining (the projective geometry engine is solicited) that he might be under attack (he has a strongly biased prior belief that one can now be attacked anytime anywhere from behind, based on what he recently saw on TV and heard on the radio repeatedly). This anticipated attack manifests in his momentary consciousness as a mental projection $S(X, t_2) = T(t_2) \circ R(X, t_2)$ of a spatial representation $R(X, t_2)$ of the imagined situation, which is construed from a certain imaginary standpoint (embedded in $T(t_2)$). The man does not confuse this counter-factual imaginary projection with a perception (he is not hallucinating): we assume that as long as perception is tied to direct sensory evidence one knows what one seems to be perceiving and also knows what one does not seem to be perceiving, in particular when information is based on long-term memory access, so that the counterfactual scenario imagined is known not to be among the situations currently perceived. As part of a process of active inference, the projection nevertheless demands action in order to confront predictions with sensory evidence.

The man expects, when he will have turned around, to see another man (red manikin) pointing a loaded gun at his head, in front of the only door (t_3). He slowly starts to turn around (bodily action and $Rot(\theta)$), hindered by a sharply growing fear (free energy spikes very high)(Fig. 1-a right-tier). He imagines a possible more positive outcome, preparing accordingly for action: to run away through the door and escape the danger (free energy is reduced slightly by anticipating this scenario)(t_4). But instead of this dire outcome, as he now stands in front of the door (following body action)(t_5), he sees the door but no gunman in front. There is a strong smell of plastic smoke, the TV is on the ground broken. The cat (green figure) jumps out of nowhere and runs away, which he sees first then only hears, all of which he perceives as one unique scene in his global consciousness of space, $S(X, t_5) = T(t_5) \circ R(X, t_5)$. He now understands that the cat must have managed to make the TV fall to the floor, causing the loud explosive sound. Relieved (FE is lowered), he now smiles at his mistaken panic-attack, staring emptily at the door, updating his prior beliefs in memory (i.e., reappraising) towards the greater likelihood of one’s cat messing around than of being attacked at home.

2.6. 3-dimensional projective structure of the FoC

Here we further specify the geometry embedded in the PCM as a model of the subjective experience of embodied consciousness that defines the FoC.

2.6.1. The FoC as a 3-dimensional projective space

Technically, a projective n -space \mathbb{P}^n can be defined on fields of numbers \mathbb{K} , as the collection of all lines that go through the origin of the vector space $\tilde{V} = \mathbb{K}^{(n+1)}$. Thus by definition, $\mathbb{P}^n(\mathbb{K}) := (\mathbb{K}^{(n+1)} \setminus 0) / \sim$, with $\setminus 0$ expressing through a set subtraction symbol \setminus that the origin of the space is excluded from the projection. This is because if 0 were included, all the points would become equivalent. The symbol $/ \sim$, expresses the equivalence relation which relates points along the same vector lines. The vector space \tilde{V} has one more dimension than the projective space $\mathbb{P}(\tilde{V})$. In what follows we only consider the case of the field \mathbb{R} of real numbers, which is a model of continuity. Moreover we will mainly limit ourselves to $n = 3$.

In a projective space, the projective subspaces are defined as the sets of lines that belong to vector subspaces of $\mathbb{R}^{(n+1)}$. For instance, a projective line is the set of lines through 0 in a 2D plane, a projective plane is the set of lines through 0 in a 3D space, and so on. Hyperplanes correspond to subspaces of codimension one. When a hyperplane is selected, chosen with an equation that defines it in $\mathbb{R}^{(n+1)}$, its complementary subset in the projective space \mathbb{P}^n has the structure of an affine space \mathbb{E}^n (associated with the corresponding vector subspace of dimension n as a space of translations). Then parallel lines in \mathbb{E}^n converge toward vanishing points p^∞ at infinity, the set of which defines the horizon hyperplane.

2.6.2. The FoC under the action of projective transformations

Projective spaces and representations in projective spaces are ruled by the action of projective transformations from the Projective Linear Group $PGL(\tilde{V})$. These transformations govern global changes of perspective taking. Key parameters of the generative models that we consider thus become projective transformations T . One practical advantage is that linear algebra and matrix operations can be used to represent the whole process.

For a projective space in 3 dimensions, the group corresponds to the set of all 4×4 invertible transformation matrices M , which transform points in the ambient space (represented by homogeneous coordinates) to points in projective space through changes in perspective taking in a 4-dimensional space (M is a subset of the possible transformation matrices T , which include non-invertible transformations). The group includes all affine transformations (rotations, translations, shearing and scaling) and projective transformations that define direction of aiming and perspective in the space while preserving relations of incidence, without necessarily preserving the plane at infinity. The reciprocal switch between points of view is governed by simple projective parameters. All possible projections are thus related via a set of transformation matrices, which are invertible and can be estimated. In this context, T denotes such an invertible transformation g , which corresponds to an invertible matrix M_g in $PGL(\tilde{V})$.

It is also important for the model to include *generalized projective transformations* in the projective space \mathbb{P} . They correspond to non-invertible 4×4 matrices different from 0 , or non-invertible linear applications of \mathbb{R}^4 . A particular case is the projection from a point O in \mathbb{P} ; in this case the kernel of the associated linear map is the vectorial line corresponding to O . From a plane to another plane it is an ordinary *perspective*, the most usual for drawing, that projects the horizontal plane, including its horizon at infinity, to the vertical plane in front of the implicit point of view.

A fundamental result of projective geometry is that all projective transformations from a given plane to another plane, i.e.,

a transformation that preserves pencils of lines, sending lines to lines, are obtained by composition of usual perspective transformations centered at variable points (Artin, 2016).

2.7. Relations between FE, projective geometry and the optimisation of preferences in the PCM

The projective geometry engine serves as a unified level of representation and control for the integrative process of active inference: (1) about the causes of limited sensations and the uncertain consequences of action (Clark, 2015; Friston, 2010; Hohwy, 2013); and (2) about the gradient of personal relevance in space in relation to embodied cognition and multiple dimensions of appraisal (Damasio, 1999; Gray, 1990; Murray et al., 2015; Paradiso and Rudrauf, 2012; Rudrauf et al., 2003; Sander et al., 2003; Varela, 1999).

It provides the agent with a method of situated representation of space based on which the process of active inference can progressively build an integrative non-local representation R of the Euclidean world, which can be associated with indices of personal relevance, or likelihoods of affectively valenced events in memory based on the FE framework. R and its valuation are not directly accessed in consciousness as such, as they correspond to the entire framework of relational knowledge and beliefs potentially available to the system (or a subset of it). Conscious access is always a contextual, situated process which builds upon perceptual inference (Hohwy, 2014), using perspective taking to support generalizations and action programming.

There is a one-to-one mapping in the PCM between the elements of R in space and the valuations in memory of cognitive and affective personal relevance across levels of appraisal. The projective representation S acts as a spatial filter, which frames both the world model R and the associated mapping of preferences encoded in the prior beliefs in memory. This provides a principle of driving for perspective taking across perception, imagination and action that orients consciousness according to the expected spatial distribution of preferences. Through FE, the process is tuned to be sensitive to surprise, that is, the discrepancies between expectations in memory, e.g., as accessed through conscious imagination, and the actual state of affairs, as perceived. The balance between discrepancy and confirmation contributes to drive the parameter of free energy.

As a generative model is learned by revision of priors or “Bayesian updating” under the driving of free energy, minimizing free energy corresponds to optimally approximating the true posterior probability of the state of the effectively observed world, both in terms of content and personal relevance in our model. Free energy minimization thus operates as a biased operator of knowledge generation (biased by preferences) when associated with a mechanism for updating prior beliefs.

In this general framework, “active” inference (which implies actual or imaginary motion) is a method of optimization for the collection of observations (e.g., from sensory data D), aimed at updating previously estimated rules for hidden variables (e.g., the world model R) and parameters (e.g., valuations of locations and objects). It selects observations which are best predicted (Friston, 2005, 2010). In some sense, this is a reverse optimization algorithm, or an optimal conditioning for further likelihood maximization. Importantly, in this framework, violations of expectation can be resolved through two complementary strategies: (1) revision of prior beliefs (e.g., internal models of the world) or reappraisal (depending on learning rules), in order to ameliorate prediction by reshaping memory; or (2) adaptive actions aimed at confirming prior expectations, thus enacting biases of the system (e.g., which can lead sometimes, in relation to learning parameters, to ignoring part of the world because it does not match prior perspectives).

As a result, the system maintains a dynamical balance between, on the one hand, the use of fixed prior beliefs to resolve uncertain situations and, on the other hand, the updating of prior beliefs based on new evidence that violates expectations or preferences. The arbitration of this balance governs the choice between strategies of cognitive reappraisal or of active approach or avoidance. It can be expected to vary broadly depending on the design, experience, developmental emotional history, and the current situation of the system. Nevertheless, in nature, the situatedness, autonomy, embodiment, and limited processing power, memory and energy of embodied autonomous systems, as well as the relative stability of their environments, can be expected to make biases or prior beliefs carry a heavy weight in inference and choice and to impose a slow and reluctant process of updating in many contexts. In other contexts (like the resolution of everyday perceptual uncertainty) the updating is normally rapid and eager.

Free energy thus functions as a control parameter, which arbiters and motivates, like a tension to reduce the (costly) implementation of changes in the embodied agent's states, through cognitive reappraisal or adaptive action, as they appear necessary to reduce past, current or anticipated violations of expectation. In our model, free energy fuels and puts stress on the workload of the inference engine and drives learning. Minimizing free energy amounts to maximizing the predictive power of internal models in memory (including prior goals and preferences) in relation to sensory information. Because of its intensive dimension and driving role in active inference, free energy can be interpreted from a psychological standpoint, as contributing to attentional drives, emotions and the motivation of reappraisal or action.

The concept of free energy straightforwardly subsumes key features of emotion as a motivational force for reappraisal and action. It can be expected to relate to a feeling of anxiety when it increases as the embodied agent misses its objectives, motivating the agent to act or paralyzing it but forcing it to reappraise its predicament; and to feelings of relief as *FE* decreases when everything turns out to be fine, perhaps as initially expected or hoped in imagination (Cunningham et al., 2013; Damasio and Rudrauf, 2006; Joffily and Coricelli, 2013) (see Fig. 1 and accompanying scenario).

One key specification on the free energy framework in the PCM is that it drives intentionality (the “aiming at” function in the conscious field) and the corresponding perspective taking (the specific projective transformation T): gradients of free energy across space, as mapped in memory and evidenced through sensory sampling, directly drive the selection of T and of the corresponding possible direction of perception and action. Free energy minimization thus performs an optimization of intentionality in the conscious field, which behaves similarly to the needle of a compass aligning with a magnetic field. The algorithm selects the perspective (and corresponding projective transformation) that minimizes the total amount of free energy associated with objects under the scope of the current field of consciousness, acting like a magnifying device, tuning down everything that is not directly under its lens (see Clark, 2015 for a related discussion at a different level of integration). This results in the optimization of personal relevance according to specific preferences: e.g., we look towards the fridge when we are hungry, and listen to the people behind that are talking about us. For instance, a high likelihood of positive personal relevance (e.g., pleasantness and safety) in consciousness will be associated with a low level of free energy, whereas a low likelihood of positive personal relevance will be associated with a high level of free energy (or stress). These dynamics operate across perception, imagination and action. Free energy itself is not directly accessible to consciousness, and its minimization process relies on unconscious activities. However, part of the knowledge that evolves through the optimization process becomes conscious, and this part necessarily involves the embedding in a lived workspace.

2.8. Projective geometry as an optimal tool for embodied active inference

From the standpoint of natural selection, given the central role of spatial inference in our model, the phenomenology of consciousness is expected to present strong and specific computational advantages, especially for spatial inference itself. Projective geometry implements a framework that enables intentionality, perspective taking and action planning in the system in a manner that can be driven according to the globally optimal principles of free energy minimization. But projective geometry in and of itself is also optimal for a variety of problems pertaining to situated spatial inference. In other words it is the best geometrical model to perform spatial inference from situated perspectives (Hartley and Zisserman, 2003).

In the context of active inference, the inference engine must progressively update the representations R stored in memory by building a more and more comprehensive representation of the spatial and temporal world based on a limited number of situated observations and a set of preferences. In the process, the engine selects situated perspectives, (i.e., transformations T), which optimize the acquisition of new multisensory spatial information D based on available prior knowledge and under the driving of some contextual mapping of personal relevance. For example, I'm curious, I see a door in front of me, I don't know yet what is behind it; I can perhaps expect a small room that I represent in my mind from different perspectives compatible with my expectations, which will be confronted with the actual structure of the room soon by my opening the door and entering in order to update my world model R accordingly. Such a problem of spatial inference can be optimally solved using projective geometry as a computational method. Projective geometry is optimal for forward and inverse modeling in space because it enables the manipulation of a rich set of frames of reference.

Thus, by combining a given set of perspectives from perception using projective geometry, a system can progressively infer a broader and more precise representation of the 3-dimensional world in a manner that is independent of the particular point of view and that can be stored in memory for future retrieval. Such spatial inferences can be carried out based on epipolar geometry, which is part of projective geometry and can be formulated as a multiview process relying on the theory of abstract cameras (Hartley and Zisserman, 2003). It is computationally tractable, as it can be approximated with linear algebra. It is reliable, as the shape and distribution of objects in 3-dimensional space can always be inferred through projective transformations and a small set of additional priors. The approach is analogous to the methods used for the 3-dimensional mapping of objects, e.g., from sets of 2-dimensional pictures (Hartley and Zisserman, 2003), as used in architecture, virtual reality, engineering and industry (more generally, projective geometry is the basis of Computer Graphic Imaging (CGI) and computer vision).

Moreover, with projective geometry, the frames can link complex perceptions with complex predictions and actions in a seamless way in order to explore and build a representation of the environment. (Higher dimensional projective spaces might need to be considered to account for the multiple dimensions of consciousness (consistent with multimodal sensorimotor integration), but the essential characteristics of the phenomenology we focus on here, the 3-dimensional lived space, can be understood by considering 3-dimensional projective space).

Given estimated sensory data $D(t_1)$ and internal model R representing the state of objects in space, it is possible to estimate multiple possible perspectives on the current world-model within projective spaces and to estimate the projective transformations induced by displacements of the frame from one place to another in

the space. This estimation produces the expected values of the sensory data $D(t_2)$ after displacement. The system can then use a set of real displacements and compare the new sensory changes with the predicted observations expected from new points of views. The whole process relies on the general role of geometrical invariance in the preparation of action and the prediction of the effect of actions on perceptions (Bennequin, 2014). These mathematical properties allow the system to run a continuous optimization process that updates the hidden rules and the internal model based on spatial cognition.

2.9. From the PCM to an analytic definition of perception, imagination, and action

The question of the distinction between perception and imagination has been a classic topic of debate in philosophy, psychology, and phenomenology (Sartre, 2004; 2012; McGinn, 2009; Kind, 2016). A model of consciousness should be able to account for their differences and relations, for the unity of intentionality and lived space across these forms of consciousness, and for how they relate to action. The PCM suggests a way, in connection to active inference, of understanding this in a formal manner, which we use later on to account for certain psychological phenomena.

2.9.1. Perception

Perception can be associated with a specific subset of generalized projective transformations, the perception spectrum for short (these transformations do not belong to the proper projective linear group and correspond to linear projectors parallel to a line on a hyperplane equipped with a Cartesian frame, see the section on the workspace under the action of projective transformations). (1) The point of view is located close to or within the space occupied by the head of the body in an egocentric mode. In normal conditions, this means that the projection operates from an embodied first-person perspective with the real body as a physical referential and the virtual body being correspondingly centered. Such a setup provides optimal conditions for construing perceptual evidence as it registers the center of projection in space with respect to the locations and directions of the sampling of space that is performed by actual sensory interfaces (e.g., the eyes, ears and the body envelope with the skin) based on sensory data D , which incorporate intrinsic constraints of orientation and calibration related to their source organs. (2) The frame is oriented in relation to the body referential and main axes: the aiming or “look-at” projective vector tends to point forward by default, but this may be modulated by attention, and the up-vector is normally defined from inertial information (e.g., the direction of gravity sensed by the vestibular system). (3) The projection parameters (e.g., field of view, perspective) are set to a specific calibration range (e.g., encompassing the visual field).

At each cycle of perceptual inference, the system must identify the projective transformation T that will best fit sensory data D given its sensorimotor state. Forward models of sensory organs and processing are expected to serve as priors in order to reduce the search space for T . With such priors, perceptual access can assume a fixed subset of possible projective transformations T to explain sensory data D , as sampled by sensory organs and conditioned by internal priors R , and can quickly focus attention on personally relevant information. In other words, perceptual access, which must be as close as possible to real-time processing, operates based on a regularization principle in order to solve the inverse problem of representing the 3-dimensional spatial causes of D in the workspace and frame them according to personal relevance. Only transformations T^p from the perception spectrum will normally be selected for forward modeling. (We note that our notion of “sensory data” here is neutral with respect to various com-

peting theories of perception. We are not committing ourselves to classical sense-datum theories. See Williford, 2013 for some discussion. All we need here is the idea that in one way or another the sensory organs provide an independent source of input and correction for the continually updated world model we describe.)

2.9.2. Imagination

The imagination can be associated with the full group of projective transformations $PGL(P)$ and generalized projective transformation as well, the imagination group for short. It includes all possible compositions of perspectives from higher dimensional spaces, inducing changes of frames in the space, including the “God’s eye point of view” near infinity, which offers a view of almost the entire space, from any possible direction. It can be applied to conscious access in the context of a representation of a world R that can be realistic or a fantasy world.

In the PCM, the imagination serves as a space for inference, based on anticipations and simulations of the self vis-à-vis itself and the world. It is used: (1) to assess past, present and future physical and social scenarios according to personal relevance; (2) to learn from experience by updating internal models in memory accordingly; and (3) to program future actions based on the simulation of multiple perspective changes.

The imagination group can be seen as the perception spectrum of external consciousnesses or points of view on the space. It notably includes transformations of the observer itself, in particular of its location in and outlook on the space. It can always be used to infer how a scenario would look under perceptual constraints from an egocentric, first-person perspective. Transformation matrices can be computed and applied in order to compare perceptual situations (e.g., after motion) with imaginary situations (e.g., as anticipation of perception before motion).

2.9.3. Perception versus imagination: computational complementarity

2.9.3.1. Computational similarity: implications for inference. In this framework, imagination and perception appear fundamentally related through projective geometry. Mathematically, they correspond respectively to the full projective group and to the set of projections from points to planes. They share projective transformations so that they may seamlessly interact across the projective workspace via the operation of perspective taking. This relationship is essential in the process of active inference through which the agent accumulates spatial knowledge about the self and the world. Actual perceptions can be compared to anticipations of perceptions based on imagination in order to revise prior beliefs and further minimize free energy.

2.9.3.2. Computational simplicity: an attentional filter. The projective setup makes comparing perception and imagination computationally simple. Sensorimotor contingencies can be used for the calibration of a consistent metric in the projective space. This can be done based on a unit spherical geometry in the 4-dimensional real vector space (such a geometry is well adapted for vision, in which distances are measured with angles). Alternatively, it can be done by choosing a plane at infinity and a Euclidean distance in affine space (which is adapted to voluntary motion of the body).

Projective spaces then offer natural error metrics, which are based on an intuitive space of objects in perspective that can be directly compared, for instance, using measures of distances between shapes in perspective (e.g., the point-to-point difference between the expected smile and the blank face one receives). Because the space is in perspective, carrying out the comparison in projective spaces automatically down-weights the importance of information that is located farther away in relation to the attentional scope: the precision or sensitivity to spatial discrepancies or noise is larger for near objects than far objects. The scope of the projection thus acts

like a lens and an attentional filter, rendering it a powerful computational tool.

2.9.3.3. Differences: the role of imagination in expanding operational knowledge and energy saving. Imagination and perception are nevertheless clearly distinguished in the model. The respective weights of sensory data D versus data from episodic and semantic memory (as opposed to sensory buffers), in the current definition of the accessed R , as well as that of different prior beliefs about possible transformations T for perspective taking, form multiple levels of constraint for the distinction between perception and imagination. In perception, D drives the inference. In imagination, contrary to perception, the agent builds its representations of a local world more based on data from memory and internal simulations than on online sensory data D . Consistent with these distinctions, in sleep, the precision (i.e., weight) afforded to sensory data is attenuated through neuromodulatory mechanisms, permitting a fictive exploration of projective geometry that can take on a fantastical aspect as we dream, all the more as intentionality, perspective taking and actions are driven by a purely imaginary field of personal relevance (Hopkins, 2016).

Another important difference between perception and imagination is that imagination, contrary to perception, can be deployed in a manner that is largely free from physical and inertial constraints (e.g., in imagination one can jump to the Moon in a fraction of a second), including constraints on the calibration of the frame of reference itself. This is also true about social constraints that can be released in imagination. In contrast, perception must submit to the independent order of the physics of inertia in the world within which it is operating. The imagination allows the agent to explore multiple possible worlds and actions without having to visit them and act (i.e., with fewer consequences). It is a framework for the optimization of action plans, which will de facto minimize energy expenditure.

2.9.4. Action

Embodied autonomous agents need to act based upon their inferences. Programming or predicting action requires the extraction of parameters of motion. According to the model, conscious access carries out a projective mapping $S = T \circ R$ of a representation R of the local world (including the body) according to cognitive and affective preferences. This representation is itself a mapping of a locally Euclidean, ambient affine space \mathbb{E}^3 , corresponding to the physical space. The projective space \mathbb{P} and its affine subspace \mathbb{E}^3 are linked by geometry. Passing from one to the other only requires linear algebra and adequate calibrations of the respective frames of reference (e.g., between the distance I cover in the real world and the distance I seem to cover in projective space). Such calibration needs to be actively performed by dedicated systems (e.g., the cerebellum)(Fig. 3).

The projective linear group \mathbb{P} contains all the Euclidean transformations of the affine subspace \mathbb{E}^3 . With proper calibration, transformation matrices embed the main parameters of displacements in \mathbb{E}^3 : translations, rotations, and twists. The parameters of real motion can thus be estimated for action from imagination-based projections of the motion, with a corresponding series of projective transformations anticipating perception. The system can always predict based on its internal models how the environment will appear from a given position by choosing the corresponding transformation T at any point in time and by using it to extract the parameters of motion. The system will thus be capable of imagining optimal navigation paths based on its current representation of the local world R .

The classical concept of efference copy (Droulez and Berthoz, 1991) denotes an internal sensorimotor prediction accompanying an order for movement aimed at preparing the organism for a new

expected state and serving as a reference for control. The efference copy informs the central nervous system about the active contribution of the subject in the perceptual process. Part of the concept can be reformulated projectively. The projective workspace acts like a space of prediction. It uses a set of projective transformations T' to predict how R should appear after motion under a new standpoint. This prediction corresponds to a projective notion of efference copy in our model. After effective motion, such an efference copy embodied in imagined motion can be directly compared to perception-based projections in the projective workspace. Furthermore, transformation matrices accounting for the difference between actual and anticipated perception can be computed based on projective geometry and used in order to compensate for errors of prediction in action consequences, e.g., in relation to unmodeled perturbations in motion. Note that in the programming of action the minimization concerns the expected free energy, not the actual free energy. Alternatively, the predicted sensory data under a posterior belief about a transformation could be realized automatically through motor reflexes (or saccadic eye movements; see above). This is the tenet of active inference in which descending predictions or efference copies provide references or set points for (reflexive) action.

2.10. The imagination: an algorithm for escaping local minima of free energy in the FoC

The arbitration of the balance between remote imagination (in time or space) versus local perception is critical for adaptation. In the PCM the FoC shifts between the two processes. This capacity to switch between local and non-local processing functions as an optimization algorithm: when the local FoC does not yield a better minimization of free energy, i.e., increase preferences, non-local imagination can take over.

The stability of free energy across time can be used as a higher-order control parameter to motivate a switch from a local perception mode to a remote imagination mode during which the agent scans its memory (its current prior belief map) to identify occluded or distant locations at which personal relevance would be better optimized. The switch can be implemented based on the monitoring of the stability of free energy across time. Stable free energy throughout exploration of the local environment in the FoC suggests that no further optimization of the FoC can be expected locally. When the local FoC and its immediately reachable surroundings do not provide the agent with access to locations that increase preferences, the best strategy for the embodied agent is to search beyond its situated perspective by projecting itself away. When in a certain range of stability of FE , the system projects itself at remote locations from different perspectives, searching for the direction of maximal utility, i.e., minimal free energy, and programs action accordingly, based on its belief in memory. It is an empirical fact that imagination is often triggered by a lack of change in the environment, e.g., during low-level repetitive tasks, which induce mind-wandering (Smallwood and Schooler, 2006).

In the PCM, when free energy becomes too stable (e.g., as measured with first-derivative stability analysis), the system may be trapped in a local minimum with no better perspectives. For instance, when the embodied agent cannot find a location to optimize its situation, i.e., FE does not decrease, the system might wander around its immediate location aimlessly, which, interpreted psychologically, may translate into feelings of panic, hopelessness or boredom, depending on the level of baseline free energy. More generally, we expect the dynamics of free energy to relate to basic and social emotions such as surprise, fear, anxiety, relief, joy, sadness, shame and pride, in the driving of perception, imagination and action.

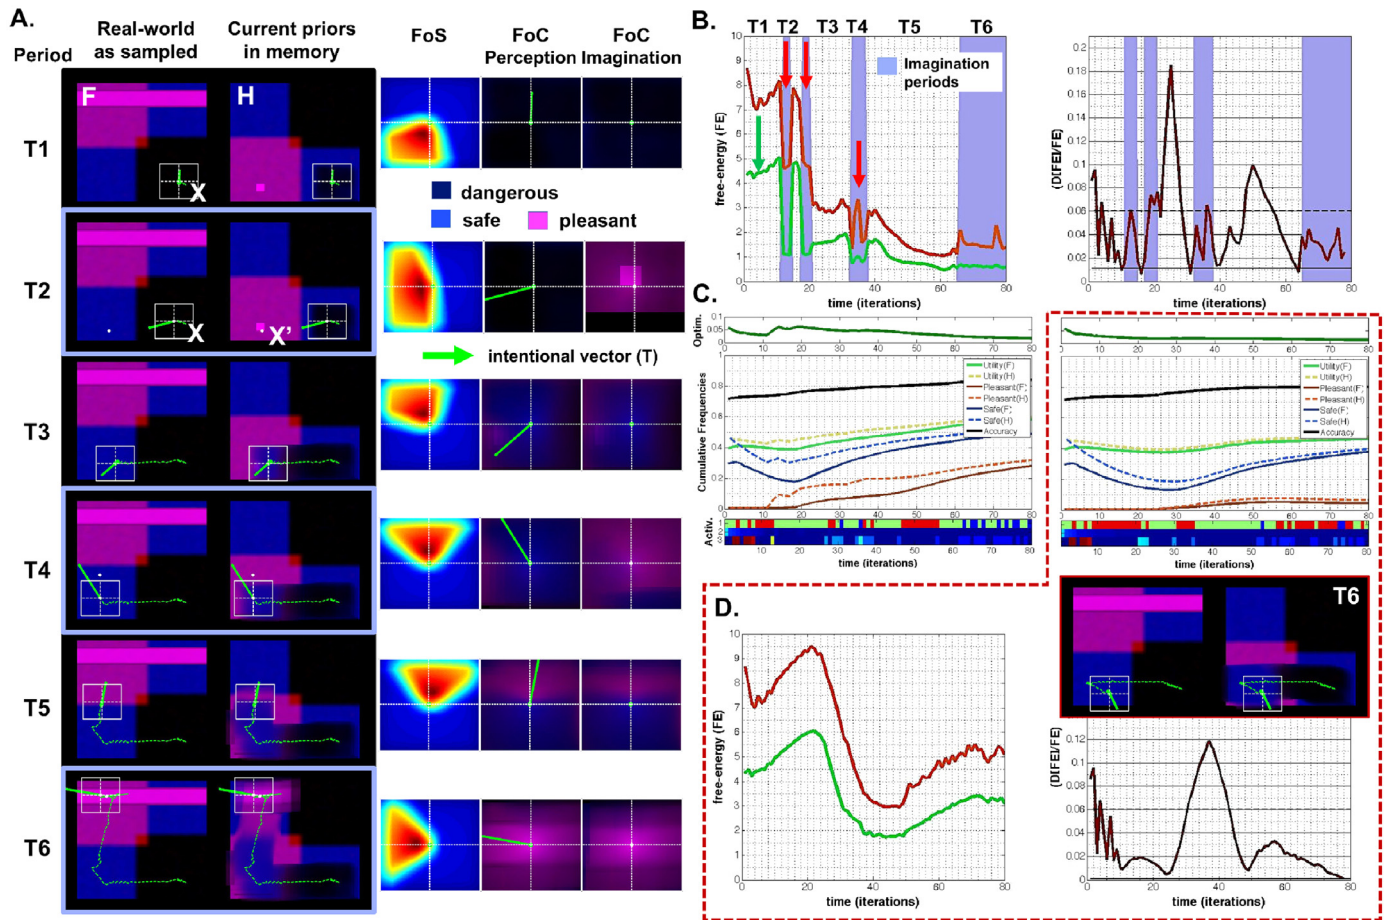


Fig. 2. Reduced-model simulation results. A. The results are organized as a table. Each row corresponds to a time period in chronological order from top to bottom. Processing modes alternate between Perception and Imagination (pale blue frame). Real-world personal relevance is represented by a matrix F and the current prior beliefs in memory by a matrix H . Personal relevance is color coded from dark blue for “dangerous” (unsafe and unpleasant), through lighter blue for safe, to pink for pleasant situations. The agent can be located at x within these maps. The field of consciousness (FoC) encompasses a subset of the world map centered at x . It is a linear combination of F as sampled by the field of sensation (FoS) and H . An intentional vector T (green arrow) points in the direction of maximum safety and pleasantness according to the minimization of free energy (FE). B. Left chart: average FE for the whole FoC in red, and the optimal perspective in green. Right: absolute value of a smoothed version of the first derivative of FE, $d(FE)/FE$, scaled by the current value of FE, as a function of time (number of iterations). C. Analytics characterizing internal and external parameters of the agent (see text). [Model parameters: n : 400 (world size, $n \times n$ matrix); $nfoC$: 50 (half span in pixel of the FoC); nit : 80 (lifespan: total number of iteration); $wEvP$: 0.9 (Perception mode: relative weight of local sensory evidence over prior beliefs in FoS); $wEvl$: 0.3 (Imagination mode: relative weight of local sensory evidence over prior beliefs in FoS for local environment tracking); $dEth_l$: 0.01 (Threshold: lower Value, for the FE derivative statistics controlling the switch between Perception and Imagination); $dEth_h$: 0.06 (Threshold: higher Value, for the FE derivative statistics. The Imagination mode is triggered when the dFE statistic cross the lower threshold and stops when it crosses the upper one.); $wlocFoC$: 0.4 (Imagination mode: relative weight of local overall FoC in the definition of the intentional vector); lr : 0.5 (‘Learning rate’, number of time perceptual evidence weights in the average $H(FoC, t + 1) = [lr * FoC(F, H, x, t) + FoC(H, x, t)] / (lr + 1)$); $speed$: 7 (Action: agent nominal speed in number of pixels in the matrix)]. D. Results of simulations for the agent with an impaired capacity for imagination (see text) [Model parameters: n : 400; $nfoC$: 50; nit : 80; $wEvP$: 0.9; $wEvl$: 0.3; $dEth_l$: 0.001; $dEth_h$: 0.06; $wlocFoC$: 0.4; lr : 0.5; $speed$: 7]. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

In line with these ideas, free energy derivatives (i.e., the changes in free energy as a function of time) have been proposed to be related to affective valence and to control through their dynamics the elicitation, expression and transition of various affective and motivational states (Cunningham et al., 2013; Joffily and Coricelli, 2013).

In the PCM, the imagination thus belongs to a class of optimization algorithms that incorporates mechanisms to mitigate the risk of solutions that are trapped in a local minimum of energy. Through active inference and the ensuing balance between exploration and exploitation, the system further maps the free energy landscape using actions and updating prior beliefs in order to further optimize its FoC . Free energy (and its derivatives) can be thought of as a state or phase space shaped by a landscape of conditional probabilities that describes where perception, imagination and action are most likely to rest or move to and from. The ongoing intentional state of aiming ensuing from FE min-

imization projects the agent across this space, making it move towards states or attractors further minimizing free energy, either through cycles of action and perception or through imagination. This renders the capacity for resilience of the system globally optimal.

2.11. The body model as the central reference of the space and its projective avatars

The PCM as presented so far generates a workspace embedding a mechanism of conscious access to spatial information in perspective based on perspective taking and driven by the minimization of free energy in relation to encoded personal relevance in space. In embodied systems, consciousness as a level of representation is also characterized by the pervasive presence of a 3-dimensional representation of the body (see postulates above), which is itself a space of high personal relevance.

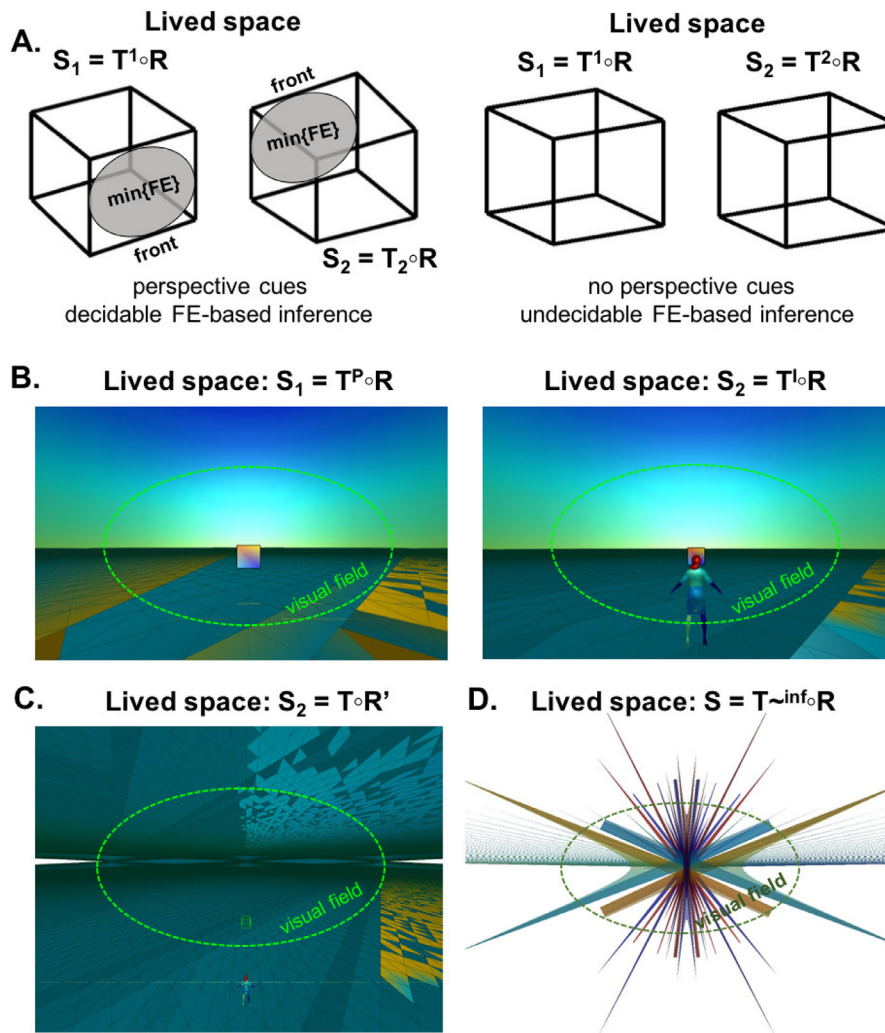


Fig. 3. Account of psychological phenomena. A. **projective solutions to the Necker cube.** Projection of two complementary views S_1 and S_2 of a generic wireframe cube R in a projective space (rendered on an image plane). **Left:** The point of view is located near the projected cube in ambient space: perspective information helps to disambiguate the inference about the two possible orientations of the cube, T_1 and T_2 (see gray ellipses indicating the front). As a result, free energy (FE) always possesses a minimum for a unique transformation T and perception can stabilize. **Right:** The point of view is placed far away but seen with a strong zoom. No perspective information remains to disambiguate the orientation of the cube. Free energy does not possess a unique minimum, each option is unsatisfying, the inference is globally undecidable but irrepressible, and the system can only oscillate between the two possible outcomes. B. **Projective solution to VR-mediated dissociations in phenomenal selfhood.** The visual field is indicated (dashed-green line) as a subset of lived space. **Left:** Projection from the normal perception spectrum: $S_1 = T^P \circ R$, the projective frame is calibrated in Perception-mode. **Right:** Projection from the pure imagination subgroup: $S_2 = T^I \circ R$. Because of the experimental conditions, a projective solution that is not normally used for perception will be chosen as the best explanation of sensory data. C. **The anti-space beyond the plane at infinity.** Rendering of the projective space with a transparent plane at infinity and no clipping, revealing the involution of the space mirroring the back of the ambient space (see text). D. **The “God’s eye” vantage point on the projective space.** Rendering on a 2-dimensional image-plane of the projective space and its mapping of the world model (used in (b) and (c)), from the “God’s eye” point of view at infinity. Almost all the space is visible but completely warped and the world model manifests as structures with complex symmetries. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2.11.1. The body as a projection in conscious access

Applying our model, we hold that the body at stake in subjective experience is of a projective nature: it is accessed in consciousness projectively, like other objects in space, in perception, imagination and action. This hypothesis implies that there is a body model (or schema) B that is embedded in the current representation R of the surrounding world, serving as a mobile referential in the space (Berthoz, 1997; Damasio, 1999; Penfield and Boldrey, 1937). It also implies that the body model and its dynamics are projectively mapped into the conscious workspace based on the application of projective transformations T to control perspective taking according to $S(B) = T \circ B$. (Technical remark: here, for clarity, we separate B from R , but we could introduce the part X_B of internal variable X that concerns the body, and keep R for denot-

ing the complete state (probabilistic law) of all variables, including X_B .)

In normal perception, the set of projective transformations used for bodily conscious access will correspond to the perception spectrum, yielding an egocentric, first-person perspective access to the body in space. In imagination and action, the full projective group can be used for perspective taking on a 3-dimensional representation of the body. In abnormal conditions, many dissociations and personally relevant violations of expectation can be envisioned between the real body and the projected-perceived and projected-imagined bodies.

The body model nevertheless plays the role of an essential reference in the space, irrespective of perspective taking: we always know where to find it in normal conditions, no matter how unfamiliar our surrounding environment may become. It is animated

and updated in real-time in consciousness based on multisensory information D .

2.11.2. The projective body, emotion and feeling

The body model and its internal structures define at each instant portions of space, which are imbued with personal relevance: it is the core space for the experience of pleasure and pain, and for the interoceptive, homeostatic monitoring of the health and well-being of the agent (Damasio, 1999).

As an integrated representation and a critical asset for the embodied system, the projective body serves utility and adaptation in both physical and social contexts. It is the object of constant affective appraisals (Moors and Scherer, 2013; Murray et al., 2015; 2015; Scherer, 2009) and is highly valued. As part of core consciousness, it is a space of experience, which, through interoceptive sensations, emotion and feeling notably, often reflects the consequences and discrepancies of actual and imagined actions (Damasio, 1999)

But we can further wonder how feelings themselves as subjective experiences could be conceived in this projective framework. We suggest that bodily feelings can be in part characterized by a spatial phenomenology. We experience our body as a sounding board (James, 1913), and bodily feelings are perceived as partly localized in the body compartment in relation to specific organs and systems (e.g., the heart, the integumentary system). Feelings are also experienced as representations of motor tendencies, with anticipated postural and behavioral changes, which also yield spatial representations.

We thus can conceive of bodily feelings as projective representations of the body in space, in part or as a whole, undergoing transformations, including representations of the internal state of the body, ultimately driven by the dynamics of free energy. More specifically, these representations can be driven by (1) ongoing real-time sensorimotor information, (2) memories, and (3) anticipatory imagination. The nature of the relationships between such mental processes and the gamut of actual physiological processes in the body *hic et nunc* remains an open question.

At a higher level of integration, the projective workspace can also be thought of as an internal environment for building, monitoring and maintaining a self-image in relation to social cognition, in which the body image as appearing under certain perspectives plays a central role and is constantly reappraised in terms of personal relevance (Rosenberg, 1965; Lacan and Fink, 1966, 2002; Kihlstrom and Cantor, 1984; Kappas, 1996; Manstead and Fischer, 2001; Paradiso and Rudrauf, 2012).

3. Results

In this section, we apply the PC Model to a variety of behaviors and psychological phenomena. We begin with the application of a reduced version of the model to simulate its basic behavior and follow with the analytical application of the full framework to account for a range of perceptual and cognitive phenomena, including some intractable problems in the philosophy of mind. All the simulations were prepared using in-house code implemented in Matlab (Mathworks) both for the reduced-model and for generating 3-dimensional projective spaces and 2-dimensional geometrical images for illustration of the broader analysis. We used linear algebra and Matlab's opengl rendering engine for patches based on wireframe meshes either created by us or freely available online (e.g., the human model).

3.1. Reduced-computational model

We derived and implemented a simplified version of our model, which includes an embodied agent embedded in a "flat land" (2-

dimensional plane), with perspective taking and action governed by free energy minimization. This computational model functions as a basic kernel onto which a full projective geometry engine operating on a 3-dimensional virtual world (filled with meshes of objects, avatars and landscapes labeled with personal relevance) will be added in the future. We apply this simplified model because it preserves essential features of the full model and offers a general demonstration of its behavior and utility.

It is a reduced model because it is not based on an optimization of 4^*4 matrices of projective transformations for a fully 3D world. It operates on a simplified model of perspective taking and projective geometry, adopted for the simulations, in a 2D world. The computation of free energy is also simplified to that of surprise per se. (The Matlab code implementing these simulations may be requested from the corresponding author).

3.1.1. Definition of the model

The world is modeled as a 2-dimensional finite plane ($n_x \times n_y \times k$ matrix; k standing for the number of dimensions of personal relevance or appraisal being encoded as conditional probabilities)(see Fig. 2a). This world could be a room or a country depending on the scale being assumed. There is an "objective" world map F , corresponding to actual expected facts, as sampled by the agent's sensory interfaces, and a "subjective" world-map, H corresponding to prior beliefs or hypotheses about facts stored in memory. Each 2-dimensional location on the maps, with coordinate x , is associated with two conditional probabilities related to personal relevance, encoding safety and pleasure as $P(\text{Safety}|H, F)$ and $P(\text{Pleasure}|H, F)$. These probabilities represent valenced situations or events that are more or less likely to happen at a given location and within a given portion of space (e.g., an object, a room) according to F or H : "safe" (S), and "pleasant" (P). The two probabilities are independent, as we consider that the same locations or portions of space can be the support of events with various degrees of S and P . When the probability of S is low, situations are likely to be dangerous, when that of P is, situations are likely to be unpleasant. (We note that F and H represent a concrete simplification and unpacking of what we denote generically with X above).

An agent is embodied in the world plane (green dot in Fig. 2a). Its field of consciousness (FoC) is defined as a local compartment of the plane (see box with dashed cross on the charts) with the agent at the center: this is the world that it can perceive and imagine just around it under a particular perspective. A subset of the FoC is updated based on sensory evidence D across a field of sensation (FoS) corresponding to the regions of the FoC informed by sensory evidence and dependent on the actual orientation of the agent in the world map, its sensory interfaces (e.g., eyes) being oriented in space. The complete FoC at each instant is computed as a linear combination (1) of perceptual evidence, as sampled by the "current field of sensation" (FoS) over F , and (2) of the spatial complement of the FoS in the FoC , which maps the current priors H in memory at the corresponding locations x . The FoC is thus a spatial combination of perceptual evidence and prior beliefs.

The agent constantly tries to optimize its field of consciousness with respect to personal relevance through corresponding appraisal layers. It continuously scans its FoC within an attentional beam around an "intentional vector" T (see Fig. 2a) pointing in the direction of its interests. The scope (field of view) of its attention defines a possible field of sensation (FoS) and direction of action. The attentional beam is down-weighted as a function of distance to emulate perspective effects (perspective division). It can vary in orientation and scope. The optimization of the intentional parameters of orientation and scope is based on the minimization of average free energy (FE) within the FoC across possible perspectives (Fig. 2b). Free energy is simply approximated in this reduced-model as the negative logarithm of the conditional probabilities

(or surprise) encoding personal relevance. Intentionality is optimized within a *FoC* centered at x by selecting the intentional vector and scope of the attentional beam, i.e., a perspective T , which minimizes $FE = -\log(P(\text{Relevance}(\text{FoC}(x))|H, F, T))$. T then corresponds to a locally optimal perspective (green vectors on the Figs.). As a result (Fig. 2b), the selected perspective has the lowest average FE (green lines in the figure) (compare with average FE in the overall *FoC*, red lines in Fig. 2b). The relative contribution of H and F in the computation depends on whether the agent operates in a 'Perception-mode' (high-load of sensory information) or in 'Imagination-mode' (high-load of prior information from memory) (see Fig. 2a-b). FE is related to motivational drives in the following way. First, in order to maximize the likelihood of pleasure, the agent's intentional vector will tend to point towards directions of anticipated maximal pleasure. This implements a mechanism of "approach" (Gray, 1990). Second, in order to minimize exposure to danger, the vector will tend to point towards directions of anticipated maximal safety. This implements a mechanism of "avoidance". The intentional vector is the sum of the two corresponding vectors.

Action is governed by the intentional vector, and position x is updated at each cycle by integrating a simple linear differential equation based on a unit intentional vector u and a coefficient setting the speed s : $x(t+1) = x(t) + su$. Priors are updated at each cycle based on the current *FoC*. The prior map H is locally averaged with the current *FoC* in order to reflect the impact of perceptual evidence and conscious experience on beliefs, with a learning rate r that controls the weight of the current *FoC* in updating past beliefs in H .

Following the principles defined above, when free energy becomes too stable (first-derivative stability analysis), for instance when the agent is not finding any location around its current location that clearly improves its situation, i.e., further minimizes free energy, the agent enters into a "pure Imagination mode" (a form of active mind-wandering), during which it scans its memory (its current prior map) in order to identify locations X' beyond its current *FoC* at which free energy could be more efficiently minimized. The switch between cycles of perception-action and imagination is defined by upper and lower thresholds bounding a smoothed version of the normalised derivative of FE: $d(FE)/FE$. The imagination mode is triggered when the lower threshold is crossed and turned off when the upper threshold is crossed. Imagination consists of the access to priors in memory via the generation of a *FoC* which is based on the prior map H versus fact map F . Nevertheless, local sensory evidence continues to contribute to the *FoC*, based on a set relative weight (even during imagination, consciousness normally keeps track of its local environment). The imaginary *FoC* map that minimizes FE is selected (model selection) as a desired destination, and will weigh in the definition of the intentional vector.

During the Imagination mode, imagination tends to drive action. The intentional vector is defined based on the most desirable location x' provided through a process of active "mind-wandering". The same equation that governs the Perception mode is integrated to actualize the motion, with a set speed parameter.

3.1.2. Simulation results

Fig. 2 illustrates the simulation results. At T_1 (Fig. 2a), the agent starts stuck in an environment which appears worse than expected as it updates its prior map and wanders around in circles without finding a more optimal outcome from the standpoint of its local *FoC*. Free energy is high (Fig. 2b) and stable. The green arrow on the left chart of Fig. 2b indicates the difference in average free energy content between the selected local optimal perspective T and the overall *FoC*. In other words, the agent has no "hope" in the local area it has been sampling through its *FoC*. At T_2 , the agent starts imagining itself at a location at which FE is minimal (red ar-

row) according to its prior map in memory and thus which maximizes the likelihood of pleasantness and safety. The agent proceeds to move towards this location. At T_3 , the agent is shown arriving at a location that does not exactly match its prior beliefs about personal relevance. According to its current *FoC*, the location is rather safe but not particularly pleasant. Free energy is lower than before but remains relatively high accordingly (see plateau between the two bursts of imagination indicated by red arrows in Fig. 2b). The situation is better but far from optimal. As the agent explores the local area, it updates its prior beliefs to better reflect reality (F) based on the sampling of the world-map via its current *FoS*. At T_4 , the agent finds itself stuck in the safe but anhedonic area. The lack of change in FE, irrespective of the perspective being taken, triggers again the agent's imagination mode. The optimal imaginary *FoC* is now found by imaginary projection in a location with a high probability of safety and pleasantness. FE relaxes and drops by anticipation (red arrow), offering a new "hope" for the agent's *FoC* optimization. The agent starts its journey towards this new destination. The imagined location happens to fall outside the actual location of the area of pleasantness on the world map (see F versus H maps in Fig. 2a). But it happens to be near it. When the agent arrives at the projected location (T_5), the actual area of maximum likelihood of pleasantness and safety begins to enter in its *FoC* through its *FoS*, and, in the process of minimizing free energy, the agent moves in this direction. At T_6 , the agent has now found a broadly optimal solution to the problem of minimization of FE, and remains stuck in it, enjoying its own *FoC*, without further optimization in perspective. FE is low and there is no local path to an even lower FE. After a while, the agent, which is "bored" (the first derivative becomes stable again FE), starts to imagine other options. But because it derives its non-local inferences entirely based on its prior map H during Imagination-mode and because that current prior map does not show any better areas outside the current area covered by its *FoC*, the agent goes back and forth locally, oscillating between Perception-mode and Imagination-mode, with its observations globally matching its expectations: it is globally satisfied. Its world-model in memory is now much closer to the actual world-map, but it still carries biases, which make the agent blind to potentially better areas and, as a result, unmotivated to further travel to remote locations, in the absence of any internal or external changes.

Fig. 2c shows a set of analytics from the agent. The middle chart displays (colored lines) cumulative average frequencies of different categories of perceived (continuous lines) versus imagined (dashed lines) situations in the *FoC* as a function of time, corresponding to the proportion of time spent experiencing: pleasantness (red lines), safety (blue lines), and overall utility (green lines), which were calculated as the average of the cumulative frequencies of pleasantness and safety. Overall, the agent optimizes utility: starting from a bad spot, it first experiences a decreasing sense of safety and as it imagines a better non-local outcome experiences by anticipation an increasing amount of safety and pleasantness. Imagination as a projective solution to the problem of local minima helps the agent to climb the gradient of utility and thus renders it more resilient. Because the agent learns by progressively updating its prior beliefs, its representation of the world in memory also progressively becomes more accurate as it explores the world (black line in the middle chart, which shows the correlation coefficient or predictive power between the prior map H and the world fact map F). The top chart shows the difference in cumulative frequency for overall utility between perception and imagination or the agent's level of "optimism". The agent starts with a phase of disillusionment, followed by a phase of optimism, progressively converging toward a balanced state of realism (the expectations of the agent better reflect the world). The bottom chart shows in color (from dark blue for low values, through green, to dark red for high values), the

cumulative frequency of three parameters related to covert and overt activity: (1) the angle of the attentional beam, (2) the quantity of rotation of the intentional vector, and (3) the amount of actual displacement. The agent's overall activity decreases as it optimizes its FoC climbing the gradient of personal relevance, and its attentional beam tends to become narrower. Fig. 2d shows the results of another simulation with identical initial conditions, the same agent, except for a lower value of the threshold controlling the triggering of the imagination mode as a function of the dynamics of the first derivative of FE. The agent has a dramatically impaired ability to rely upon non-local projective imagination, e.g., third-person perspective, and only experiences the world through a local FoC, e.g., first-person perspective. As a result, its exploration of the world is much more limited (see maps at T6). FE minimization follows a slower and smoother optimization. The level of FE remains high and sub-optimal at the end of the agent's journey, where it finds safety but not pleasantness, by mere proximity. The charts showing the agent's analytics (next to the corresponding chart for the first agent in Fig. 2) indicate bad performance in optimization of utility, with amelioration in experienced safety but barely any pleasantness experienced. The accuracy of the agent's representation of the world is also lower. The agent displays little optimism. The lack of imagination has rendered the agent less resilient.

These simulation results demonstrate the suitability and functioning of the general mathematical instantiation of the PC Model. We now concentrate on applications to more specific psychological phenomena, which we explain by reasoning based on the full model.

3.2. The puzzle of the missing origin

According to the phenomenological postulates, consciousness seems to embed a vantage point that is elusively connected to the center of lived space and plays a generative role in shaping lived space. The elusiveness of this vantage point is predicted by the PCM model. It follows from an integral, definitional property of a projective space. We saw that a definition of the n dimensional projective space \mathbb{P}^n can be obtained from a linear space \tilde{V} of dimension $n + 1$ by considering all the straight lines from the origin O . This can be understood as an embedding of the space \mathbb{P}^n in the space \mathbb{P}^{n+1} , with an additional general point of view in the supplementary dimension. The central point from which the projection is set up must be excluded from the projection to generate a projective space. By analogy let us consider the case of a projection on a 2-dimensional plane as in perspective painting or photography: to project 3-dimensional objects in Euclidean ambient space onto a canvas or film so that they appear in perspective, a center of projection or focal point has to be chosen outside of the plane of projection in ambient 3-dimensional space, otherwise all the objects would be projected onto the same point on the 2-dimensional plane that serves for representation. In other words, 2-dimensional images are really understood as 3-dimensional objects, with an elusive origin residing in a supplementary, implicit third dimension, which is necessary to set up a projection. This property generalizes to projective spaces of higher dimension such as 3-dimensional projective spaces like our lived space: to project 3-dimensional objects from a certain projective space (outside or in memory) in our projective space of 3 dimensions, a center of projection or focal point has to be chosen outside that space in a 4-dimensional space, which in practice requires the use of homogeneous coordinates in order to make the points in 3-dimensional ambient space correspond to 4-dimensional objects. Note the important theorem (Coxeter, 1961) asserting that the full projective group of a projective space of dimension n can be obtained by composing perspectives in a projective space of dimension $(n + 1)$.

Consequently, the embedding of our 3D space in the 4D space allows us to generate projective changes of frames as they are used in imagination, using perspectives from imagined points outside of the world.

In other words, the 3-dimensional lived space can be embedded in a 4-dimensional projective setup, with an elusive origin residing in a supplementary, implicit 4th (vector) dimension, allowing the projective group to act as changes of points of view.

Nevertheless, this extra-dimension is not mysterious. It is a parametric dimension, which controls perspective taking by configuring the direction, scope and perspective of the projection. In the numerical model, it is expressed in the use of 4×4 transformation matrices T (applied on homogeneous coordinates). The extra dimension with the excluded origin is essential in defining and generating the projection. It is fully computable but it does not have a direct geometrical (or subjective) manifestation in the 3-dimensional space. In 3-dimensional lived space, this extra dimension remains inaccessible. This means that according to our model, imagination as applied to the phenomenological 3-dimensional world, is generated from elements of a perception spectrum residing in a higher-dimension that cannot be directly grasped through a 3-dimensional representation of space.

3.3. Bistable illusions from ambiguous planar figures

The PCM can account for bistable visual illusions such as the ubiquitous Necker cube. Planar projection of objects in spaces or “images” tend to be decoded by the brain as representations of 3-dimensional objects from a given standpoint and interpreted as such in perception (Kim and Blake, 2005). Put in terms of our model, in such contexts, somehow perception always engages a situated form of imagination: it extrudes depth information as it renders images in consciousness, but it remains driven by sensory data, so it appears as a perception. The Necker cube illusion consists of the 3-dimensional illusion of a cube switching orientation, which is elicited by a visual confrontation with a 2-dimensional ambiguous figure (see Fig. 3-a). The ambiguous figure can be interpreted as the planar orthographic (affine Euclidean) projection I of the edges of a generic cube R (serving as model prototype in memory) on an image plane. The projection is ambiguous: two different orientations, R_1 and R_2 , of the cube in ambient space can have identical orthographic projections I_1 and I_2 on the image plane. In the experience, the projections on the retinas of such an ambiguous image plane provide sensory inputs D to the brain.

Based on the PCM, such an ambiguity would challenge any attempt by the active inference engine to solve the inverse problem of ecologically interpreting sensory inputs D as a likely generic cube R in 3-dimensional space with a given orientation.

Studies of the illusion have emphasized local rules for controlling the switch of the Necker cube (e.g., Hoffman's rule), and have considered the problem in relation to the “generic viewpoint assumption”, which implies that there is no a priori viewpoint position defined in such images but that a viewpoint is nevertheless assumed in interpreting it (Albert and Hoffman, 2000). In this framework, the bistability of the illusion induced by the ambiguity has received a Bayesian formulation (Albert and Hoffman, 2000). Our model complements this approach by specifying the projective geometrical underpinning of the phenomenon.

Solving the underlying inverse problem involves perspective taking: a projective transformation T that best fits the implied orientation must be selected to represent the cube in projective space. Unless strong priors are given about the object being of type cube R from the outset, the selection of R as a cube is also part of the process of inference. In our framework, this corresponds to maximizing the posterior probability $P(T, R|D) = \frac{P(D|T, R)P(T, R)}{P(D)}$ of the transformation T and of R , as a generic cube and internal model

in memory, given D from sensory information. If for an optimal selection of R the maximum of $P(T,R|D)$ was unique, i.e., there would be a unique T maximizing the quantity, then conscious access could be computed in a stable manner as (the prediction) $S = T \circ R$, which is the generic equation of the rendering of a cube in perspective with a genuine sense of depth and a given orientation in a projective space.

In the normal perception of a real cube in ambient space, extraction of color and shade information as well as perspective effects and binocular disparity related to depth would constrain the inverse problem. The two viewpoints would not have an equivalent posterior probability (i.e., there would be a unique optimal solution to the problem of selecting T). With a real wireframe cube presented at finite distance and projected on a 2-dimensional image plane, perspective and binocular disparity effects would remain to constrain the inference (e.g., the front side would appear bigger than the back side on retinal images, and there would be different angles of projection) (Fig. 3-a, left-tier). With the Necker cube's implicit orthographic projection on a 2-dimensional image plane, there are no perspective cues, and there is no binocular disparity (Fig. 3-a, right-tier). The transformations T_1 and T_2 that correspond to R_1 and R_2 respectively, i.e., the cube under the different possible viewpoints, have the same posterior probability and thus the same level of free energy (FE).

Moreover, in the group of projective transformations, the only transformations that are compatible with the absence of both perspective cues and binocular disparity entail an implicit point of view placed at a very far distance from the object, near projective infinity $T(\infty)$, with a small viewing angle so that the far distant object (the cube) appears big enough in the scope of the projection. For instance, observing very far objects with powerful binoculars results in their perspective cues vanishing and barely any binocular disparity. Contrariwise, observing objects from up close increases perspective cues and binocular disparity. In perspective painting, the implicit point of view (center of projection) is generally close to the main group of objects, and the wide angles used for close-ups in photography magnify perspective effects. Points of view placing objects near projective infinity (or equivalently points of view placed near projective infinity) correspond to a subspace of the possible configurations of a projective space, which is equivalent to a space of orthographic projection. Near projective infinity, the 3-dimensional projections of the two alternate orientations of the cube "look" identical and thus remain totally ambiguous. There is nothing in conscious access allowing the agent to further disambiguate the information and privilege one orientation over the other.

However, placing the point of view near an infinite distance from the object in order to explain the absence of perspective cues and depth disparity in D directly conflicts with ongoing information about the actual distance of the object in ambient space in the experiment, e.g., just in front on a computer screen.

By the end, (1) with no unique maximum in the posterior probability of competing transformations T , and (2) conflicting priors about relative distances, the inverse inference is maximally ambiguous and, outside of a metacognitively imposed voluntary control exerted on higher-order attention, the active inference engine can only oscillate between bistable percepts, with no help from conscious access to stabilize the representation.

3.4. Proprioceptive illusions and out-of-body experiences

Clinical and experimental situations can elicit experiential dissociations between the sense of the location of the subjective point of view and the perceived location of the body in ambient space, e.g., OBEs, syndromes of somatoparaphrenia, anarchic hands and phantom limbs, the rubber hand illusion, enfacement or full-body

illusions (Blanke, 2012; Blanke and Metzinger, 2009). In OBEs, subjects hallucinate their own body from a third-person perspective, e.g., as a phantom body above them, or travelling away (Stratton, 1899).

VR-based manipulations have facilitated reliable elicitation of full body swap illusions in the laboratory, which are analogous to OBEs (Blanke, 2012; Blanke and Metzinger, 2009; Lenggenhager et al., 2009). For instance, when a participant and a humanoid virtual avatar are being stroked simultaneously, an embodied illusion of spatial identification with the virtual avatar (for instance, projected in front of the actual participant) is experienced. The location of the body feels both near the point of view and at the avatar's location. Likewise, Ehrsson (Ehrsson, 2007) placed cameras behind the backs of participants and projected the image of the body of the participants in front using a VR device. This sufficed to induce a subjective impression in the participants of being literally seated behind their own backs and looking at themselves from this location (Hicheur et al., 2013; Lambrey et al., 2012; Thirioux et al., 2010).

Such dissociations between the subjective locations of the point of view and of the body in space are predicted and can be explained by our model. Base on the PCM, the body as experienced is accessed in consciousness in a projective way, like a multisensory-informed virtual body avatar. In normal perception, the projective body is rendered from an egocentric, first-person vantage point, as an object in space, which serves as a central reference to which the implicit point of view is attached. Only a subset of generalized, degenerate projective transformations T^p , the perception spectrum, will be assumed as priors for spatial inference and perspective taking in order to access the body representation projectively as $S(B) = T^p \circ B$.

In VR paradigms inducing the perceptual experience of sitting behind one's own back (Ehrsson, 2007) or similar phenomena of body swapping (Lenggenhager et al., 2009), multimodal sensory data D are conflicting (e.g., visual and somatosensory) in relation to the perspective normally assumed in perception. They are evidence for models of perspective taking, i.e., projective transformations T^l , that normally are only associated with pure imagination, e.g., in which the representation of the body is accessed from a third-person perspective. These transformations by no means require an actual displacement of the person in real space: they are a parameterization of the transformation matrices through which the representation of the body is accessed in consciousness through perspective taking (see Fig. 3-b).

In such VR paradigms, the process of spatial inference is challenged. Perceptual transformations T^p are expected, i.e., the prior $P(T^p)$ is high, because the process remains driven by sensory data D . But the conditional probability of the data $P(D|T, R)$ is maximal for a specific transformation T^l that belongs to the full group of invertible projective transformations, which is normally dedicated to imagination in our model. In such ecologically abnormal situations, the best bet for the agent is to use transformations T^l to explain sensory data D . The predictive power of the solution outweighs the priors of normal perceptual transformations. What would normally be only a possibility for the imagination thus manifests as perception, since the inference remains driven by sensory data D . The boundary between imagination-related and perception-related modes of conscious access is shifted in this highly unusual context.

3.5. Self-representation and pre-reflective self-consciousness

3.5.1. Self-representation as awareness of the workspace itself

A model of consciousness should be able to account for the central yet puzzling capacity of consciousness to somehow be aware of itself. This capacity for self-representation or self-acquaintance is a topic of intense debate in the philosophy of mind and cogni-

tive science (Kriegel, 2012; Kriegel and Williford, 2006; Williford, 2015; Williford et al., 2012; Zahavi, 2008).

Maintaining a certain neutrality about many contested issues here (e.g., the exact nature of the representation or acquaintance involved), we approach the problem starting from our definition of consciousness as a workspace with a characteristic spatial phenomenology based on projective geometry. In this context, consciousness represents world structures and itself in a spatial manner. Full self-representation for that model of consciousness could mean the ability to represent the entire space of representation itself, including the point of view.

When the point of view is placed near infinite distance from its target and the field of view is maximal, much of the space appears wrapped into the scope of the projection (see below). The entirety of the generative space cannot however be fully encompassed in this scope, since the origin of the space and the parametric extra dimension will always remain elusive.

Conversely, when the point of view is taken from somewhere near us, a kind of imaginary cyclopean eye, the world is roughly divided in two parts: the world at finite distance (itself divided in body space, prehension space and near locomotion space, Berthoz, 1997), and the infinitely far world, where we can only imagine things happening, but out of the scope of our own actions. Nevertheless, consciously, we can relate these things that are projected on the plane at infinity. This ability is consistent with the existence of natural projective metrics, in which all distances are finite, and in which the plane at infinity corresponds to the set of points at maximal distance from the center of vision.

3.5.2. The anti-space and pre-reflective self-awareness

Nevertheless, projective spaces embed further features of genuine self-representation that are remarkable. The mechanism of perspective taking defines a vantage point and a plane at infinity, and the complement of these elements (the represented world) in the generated space has the structure of a 3-dimensional affine space (on which a Euclidean metric can be defined). This geometry is non-trivial. Considering a large sphere, with a given fixed center (the standpoint) and a very large radius that tends to infinity (the possible horizon of consciousness), antipodal points on this sphere project to the same point at projective infinity in the projective space: the space is folded onto itself through the topological gluing of antipodal points. Thus, paradoxically, the back of the ambient or modeled space with respect to the vantage point (be it empty or informed with a representation of a world R from memory), is actually projected in front, but on the other side of the infinity plane, similar to a 3-dimensional virtual mirror image. This “antispaces” is usually clipped in digital imaging.

Such strange projective property of the involution of the space could not manifest in normal perception because the projective space is practically clipped from residual projections in the compartment beyond the horizon: in perception, the celestial sphere saturates the horizon with perceptual information, and world boundaries in the finite compartment of the space (e.g., walls) will tend to mask it. Moreover, one can expect that for computational efficiency, the parts of the world representation R mapped in memory that are behind the vantage point and outside the scope of the perspective do not generally need to be processed for normal conscious access. They can be ignored by attentional selection, prior to projective inference.

Nevertheless, such a symmetric mapping does exist and is always defined in projective spaces no matter which structures from R fill it, and it is mirrored in the very scope of the projection beyond the plane at infinity and could thus perhaps be explicitly observed if the horizon could be somewhat rendered less opaque. Thus, geometrically, in the projective space of our consciousness, in perception, imagination and action programming, one also al-

ways aims at one's own “back”, at an “anti-space”, when one aims forward at the horizon (see Fig. 3-c).

In our framework, this can be interpreted as one way in which consciousness is always tacitly aware of itself prior to any explicit reflection. Geometrically, the vanishing point I am aiming at, the horizon of my conscious space at positive infinity in front, is also the projection of the back of the antipodal point that is located behind me at negative infinity. My implicit vantage point is paradoxically located “between” the same unique ideal point at infinity, which seems at the same time to be two points along opposite directions. Although these statements may appear contradictory they are mathematically well defined in projective geometry, consistent and fully computable, and the behavior of the space in the finite compartment respects usual relations of orientation and incidence. Thus, according to the model, there is a projective machinery that is defined but remains normally implicit. The way it works as a projection does not imply that an explicit curvature will be apparent in such “elliptical” spaces.

In the variegated and multi-disciplinary literature on self-representation, the concept of “pre-reflective” self-consciousness has become an important subtopic of debate (Henrich, 1966; Zahavi, 1999; Frank, 2004; Williford, 2006; Kriegel, 2009; Williford et al., 2012). The idea that consciousness always involves a form of pre-reflective self-consciousness is the hypothesis that there is a genuine reflexive (as opposed to reflective) form of self-consciousness that (1) does not require a distinct, higher-order mapping or meta-representation S_2 (a representation of a representation) of subjective consciousness S_1 at the first-order, (2) is truly embedded in the space of awareness itself, as elusive as it may be, and (3) is a necessary basis for higher-order, reflective self-awareness. Interpreted variously, the notion has been criticized for being unanalyzable, ill-defined, or circular, whereas some, on the contrary, emphasize the importance of finding models that could incorporate such circularity without generating paradoxes (Williford, 2006; Williford et al., 2012).

We suggest that the peculiar topological properties of projective geometry could perhaps make sense of the hypothesis of a pre-reflective consciousness in a more formal way. This is at least one way in which a pre-reflective form of self-awareness can be built into the architecture of consciousness, and it is an intrinsic feature of our model.

3.5.3. Gaze behavior in mind wandering

Beyond its general role in shaping imagination under the driving of personal relevance, this peculiar geometry might also further illuminate the phenomenon of mind wandering. Mind wandering and daydreaming are directly linked to increased self-referential processing and self-consciousness (Smallwood and Schooler, 2006). When our mind wanders it tends to turn toward self-related inner concerns away from external information processing (Klinger, 1971) and to generate projective scenarios about the past, the present, and the future (Smallwood and Schooler, 2006). Attention is withdrawn from outside events, we are less sensitive to sensory data, and we may appear to others as phasing out or being distant. Our eyes seem to look through the vacuum to the external observer; that is, we lose binocular fixation on the proximal environment, and eye movements can become erratic (Reichle et al., 2010). In other words, we seem to down weight the processing of multisensory data D and weight up that of information from memory and internal simulations on the R being accessed, increasing, according to our model, imagination (about past, present and future events) over ongoing perception. Interpreted from a geometrical standpoint, the overt gaze behavior may suggest that we tend to enter into divergent parallel binocular states during self-referential processing. Interpreted as an index of a projective parameterization of a perspective on an imaginary space, divergent

parallel binocular states imply underlying projective transformations T that aim at or from infinity. Although this remains speculative and should be addressed experimentally in the future, our model would predict that, in such aiming, the subject may also be aiming at the anti-space, i.e., at her or his own “back”. Under such perspectives, a subjective point of view on the entire space might be approximated in a manner that is consistent with an activity of self-representation.

3.5.4. Mystical and psychedelic experiences

Various spiritual practices (e.g., meditation), the use of hallucinogenic drugs (e.g., mescaline), extreme life experiences (e.g., near death experiences), dreams, phenomena of entoptic vision induced by ocular pressure, and clinical conditions have been associated with reported subjective experiences of strange and fascinating spatial patterns in consciousness: radial patterns, spirals, funnels, cobwebs, honeycomb lattices and triangles, deep tunnels, which have then been a topic of artistic representation in many cultures (e.g., mandalas) (Klüver, 1926, 1966; Tyler, 1978; Fregnac, 2003; Siegel, 2005). Such patterns have been hypothesized to reflect the inner structure of spatial consciousness as a level of representation and a mode of access to information (Tyler, 1978, as cited in Fregnac, 2003). We can wonder whether our model can predict similar patterns, as would be expected from a valid model of consciousness, given their importance in human experience.

Projective spaces and transformations can lead to peculiar perspectives on the structure of the space, in particular when the implicit standpoint is placed at near infinite distance from the center of the world in the scope of the projection (which is a parameter to set in the corresponding transformation matrix). This is, so to speak, the “God’s eye” point of view. From this standpoint almost the entire space of representation manifests in the scope of the projection. As a result, symmetries and strange topological patterns may appear, which are reminiscent of the type of aforementioned patterns (see Fig. 3-d). We hypothesize that such psychological phenomena are a reflection of the application of “limit” projective transformations (e.g., placing the point of view near infinity), generally in a context of extreme conditions, in which the mind moves away from the normal constraints of perception and uses the imagination group widely to solve the inverse problem of representing a spatial world in a projective manner from a stream of sensory and mnemonic information. We may wonder whether such a projective setup, with its intrinsic dissociation between the subjective standpoint and the apparent location of the body in space, could be a natural ground for the emergence of beliefs in various metaphysical forms of dualism (Metzinger, 2005).

4. Discussion

We derived a mathematical model of embodied consciousness, the Projective Consciousness Model (PCM), based on projective geometry and active inference, which features a characteristic mode of representation of space or spatial phenomenology, resulting from a mechanism of conscious access based on perspective taking and which operates based on the minimization of free energy as related to a probabilistic field encoding cognitive and affective personal relevancies mapped across space. The model: (1) yields basic behaviors that are predicted by the theory and phenomenologically sound and accounts for a variety of cognitive phenomena that have not been linked previously in a unifying manner, (2) predicts core aspects of the characteristic phenomenology itself, and (3) renders explicit the link between phenomenological features and their functional and computational value as a specific mode of access to information and a support for the driving of action in the context of situated cognition.

The PCM directly connects consciousness, active inference and appraisal to resilience, in line with previous hypotheses (Kalisch et al., 2015; Rudrauf, 2014). Until now no model has been introduced that is both fully computationally tractable and integrates a comprehensive model of cognition and affective processing in embodied ecosystems, including a model of the first-person or experiential perspective. The development and implementation of mechanistic models of cognitive and emotional resilience are essential in order to precisely understand, in a predictive manner, the mechanisms mediating resilience, in the perspective of preventive and acute healthcare.

4.1. Neurocomputational hypotheses

The computational architecture implied by our model can be translated into general neurocomputational hypotheses, in accordance with the principles of neurophenomenology (Rudrauf et al., 2003; Varela, 1996). We have argued that such and integrative processes of conscious representation and control may rely on a form of virtualization in the brain, integrating information processing through layers of abstraction into a global, resilient, predictive world model (Rudrauf, 2014).

In this framework, we expect the relationships between the PCM algorithm and the functioning of neural networks, in particular at the microscale, to be complex and indirect. On a general note, the assumption that the brain uses a projective geometry as part of its generative model is neutral with regard to the particular neuronal encoding. For example, positions and metrics could be encoded as continuous variables in neuronal population activity. Conversely, there could be a discrete state space model of the sort suggested by classical receptive fields (with compact support). Irrespective of the particular encoding, our point is that there must be a formal connectivity among neuronal populations encoding allowable transformations that mediate belief updating among neuronal representations (see for instance recent hypotheses about neural computation referring to topological theorems that are also intrinsic to projective geometry (Tozzi and Peters, 2016)).

In terms of the functional architectures that the projective geometry model calls for, there is an abundance of empirical evidence implicating specific brain systems. On this basis, the model predicts that the brain should embed two main engines that are hierarchically coupled (Fig. 4): (1) an inference engine grouping higher levels of control concerned with homeostasis, emotion, memory, language and executive functions, or more generally personal relevance; and (2) a lower (sensorimotor) projective geometry engine, relying on multisensory integration and motor programming and concerned with representing the world and the body in the world and mapping gradients of free energy to drive intentionality. Active inference provides the overall hierarchical infrastructure and various message-passing schemes required for generating predictions and inferring the world (R) and adaptive transformations (T).

From an anatomical standpoint, we hypothesize that the inference engine relies on anterior cortical and subcortical systems, including limbic and non-limbic frontal and temporal association cortices, the amygdala and hippocampus. The projective geometry engine can be hypothesized to rely on posterior temporal-parietal-occipital, modal and multimodal sensory systems, in particular in parietal systems, encompassing exteroceptive, proprioceptive (including vestibular) and interoceptive processing, but also on frontal premotor regions. Recent evidence supports the existence of a partition of the cerebral cortex into two main anterior and posterior functional ensembles, in a manner that is compatible with our hypothesis of a posterior real-time projective geometry engine and of a slower anterior hierarchical active inference engine quantifying personal relevance (Mesmoudi et al., 2013).

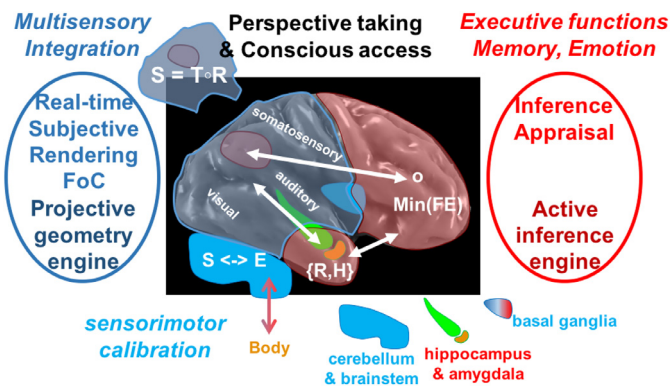


Fig. 4. Macroscale neurocomputational model. Lateral aspect of the right hemisphere of a brain model. The projective geometry engine is implemented in posterior cortical and subcortical systems (gray blue). The framing and rendering of the subjective experience of space through perspective taking is performed by those systems, which apply projective transformations such that $S = T \circ R$. The active inference engine is implemented in anterior limbic and non-limbic cortical and subcortical systems (red). It integrates memory (including world-representations R stored in mesiotemporal systems), affective processing of personal relevance (the amygdala and ventral prefrontal cortex), and executive components (the prefrontal cortex and basal ganglia). The cerebellum (bright blue), along with the basal ganglia, perform an ongoing sensorimotor calibration, relating the projective space to the Euclidean space of action \mathbb{E} . All the systems are coupled and interact through large-scale connectivity and networks (color arrows). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The tight relationship between spatial memory and affective processing in the brain, as mediated by the interactions between the hippocampus and amygdala, is in general support of a model of space in memory imbued with personal relevance: the strong interactions between the hippocampus and the amygdala and other limbic structures are thought to play a role in the labeling of memory of spatial locations with personal relevance (Gray, 1990; Sander et al., 2003). A large body of literature implicates the hippocampus in spatial representations and scene construction, ranging from place cells to sequential navigations to metric spaces (Kitanishi et al., 2016; O'Keefe et al., 1998; Zeidman et al., 2015). The parahippocampal system would then be expected (in our model) to play a key role in the representation R that is stored in memory and that is accessed in consciousness through projective transformations in the recall of scenes. Likewise, the strong interactions between the DMN and medial temporal systems in the context of autobiographical memory recall (Buckner et al., 2008) can be understood in the PCM framework as related to the process of addressing R and its labeling in terms of personal relevance in the model.

The construction and application of projective transformation T to both sensory data D and data from memory R are presumed to call on parietal processing. For example, the distinction between what and where pathways in the brain (Ungerleider, 1982; Ungerleider and Haxby, 1994) tells us immediately that a vast amount of the visual hierarchy (the dorsal stream, which includes much of the parietal cortex) is devoted to spatial transformations. This observation is endorsed by spatial neglect syndromes following damage to the parietal cortex (Andersen, 1997; Bartolomeo and Chokron, 2002). Furthermore, mental rotation tasks (Zacks, 2008) and scene construction tasks (Hassabis and Maguire, 2007) almost universally engage parietal cortex, including the precuneus and retrosplenial cortex. Work looking at saccadic eye control (perhaps the most obvious manifestation of spatial transformations) and sensorimotor transformations in motor planning implicates regions like the intraparietal sulcus (Vossel et al., 2015).

At a more global scale, we expect frontal-parietal-temporal interactions to be related to inference and perspective-taking via the

adaptive selection and application of projective transformations to spatial models in memory as a function of personal relevance. Conscious access across cycles of inference requires addressing information related to R in memory or to D in multisensory systems in an adaptive and well-scheduled way based on such transformations. Switching between these two perceptual and imaginary modes of access may be expected to be supported by medial and lateral fronto-parietal networks. There is an ongoing debate in the cognitive neuroscience of consciousness as to whether conscious access and experience require fronto-parietal interactions or only coherent activity in posterior cortices (Crick and Koch, 2003; Dehaene et al., 2006). We believe that this debate might be in part driven by a focus on different aspects of consciousness that may end up corresponding to the different hierarchical levels implicit in our model and its division into two main engines.

At a lower level of integration, the brainstem, thalamus, reticular nucleus of the thalamus, the precuneus and related regions can also be expected to play a role in the arbitration or switch between different projective modes (e.g., “Perception-mode” versus “Imagination-mode”), and in a form of arousal related to free energy. These brain systems can be expected to modulate, notably through mono-aminergic transmission, the precision or weight allocated to sensory data D in the process of projective inference and are key modulators of vigilance states. The cerebellum as well as the basal ganglia are expected to play a key role in the calibration of the different subspaces and transformations between them (e.g., between S the projective space of conscious access based on perspective-taking, the Euclidean representations of knowledge in memory R , which are independent of any point of observation, the associated field of personal relevance, and the compilation of predicted action into a space of effectors with biomechanical constraints, Fig. 4).

4.2. Real-world applications: implications for system design and autonomous computational agents

Beyond the scientific and theoretical value of a computational model of embodied consciousness, our approach has a strong potential for real-world impact. We argue that the emerging set of interrelated fields that include autonomous computational agents, robotics, and artificial intelligence (AI) as well as new developments in virtual reality (VR) applications can benefit significantly from the development of a mathematical model of embodied consciousness. The PC Model contributes by specifying what parts of consciousness (beyond intrinsic biological properties) could be transferred into an AI environment and the principles by which these components will function and become integrated. The model also provides specific predictions for how embodied AI agents will behave and an explanatory framework by which these predictions can be interpreted in human terms. This latter advantage is no small contribution: it has become increasingly clear that one of the most pressing problems with the application of advanced AI systems is their ultimate inscrutability (Yosinski et al., 2015).

Some of the principles at the center of our model are actually already heavily used, in bits and pieces, in industry: in filtering and image processing, in cybernetics and robotics, in cinema (CGI), in virtual and augmented reality. Our model, which can be interpreted as a formal theory of “enaction” and embodiment, though without the externalist implications normally associated with these approaches (see Rudrauf et al., 2003) integrates these principles into a coherent whole that offers the sketch of a design for embodied consciousness in artificial systems.

4.3. Real-world applications: consciousness and virtual reality

If the PCM is correct, consciousness accesses multisensory information D by framing it as a representation of space in a 3-dimensional projective space, supported by prior beliefs about the 3-dimensional structure of the space to be represented and about the forward models governing sensory processing. In this context, the framing of spatial consciousness by projective geometry may explain why it can make sense of 2-dimensional projections such as images and films from multiple standpoints and why Virtual Reality (VR) can be so effective at inducing an immersive feeling of realistic presence in a 3-dimensional world.

The PCM predicts that configurations of inputs provided to sensory interfaces (i.e., a projection of a geometrical model of a virtual world or of the real world on the sensory interfaces) will be interpreted and framed in consciousness as the experience of a world of 3-dimensional objects in perspective. In VR, the world represented, just as with normal perception, will be based on an inverse projection inferring the causes of the sensory inputs and implicating an operation of perspective taking. If the virtual world is simulated in a realistic manner, this simulated world is expected to appear to consciousness with a sense of direct realism, that is with a sense of natural evidence, similar to that of normal perception.

VR immerses participants in an artificial sensorimotor environment that takes cybernetic control of the interfacing of the conscious workspace with sensory inputs and motor outputs (Herbelin et al., 2002). Through metacognitive processing, we may well understand that cognition is immersed in a counter-factual virtual world, nevertheless the brain interprets the inputs and the consequences of the outputs on them, at a sensorimotor level, in an embodied and affective manner, as it would with the real physical world. Consciousness cannot help but access this information as it would for normal perception, since it is structured to be processed accordingly. The generation of the stimuli used for VR-devices (e.g., Head Mounted Display or HMD), is based on the application of projective geometry (and projective transformations) to 3-dimensional geometrical models of dynamical structures in Euclidean space in a manner that takes into account the real-time position and orientation of the HMD for perspective taking. Sounds are also presented through headphones according to geometrical and physical transformations to account for the localization of their sources in the world model. A persistent problem with VR experiences for decades has been the ubiquitous feelings of motion sickness and nausea. VR fails to be immersive and to induce a rich, natural sense of presence in another world when the sensorimotor contingencies, the responsiveness of the VR devices to motion (HMD, audio headset, controllers tracked in space), the sense of agency, and the accurate control of intentionality and action are limited by the technological setup. In other words, immersive presence is only achieved when the setup can flawlessly simulate the input-output relationships and parameters that condition the type of perspective taking embedded in the PCM.

Nevertheless, because its geometrical properties are so in tune with the generative mechanisms of consciousness we put forth, VR appears as an exceptionally promising tool for studying consciousness. The application of such VR paradigms in the context of paradigms of sensory conflict can reliably induce numerous states of bodily awareness in which phenomenal selfhood parameters can be manipulated at a sensorimotor level independently from: (1) the self-identification with a body as whole; (2) the self-location relative to the lived body; and (3) the framing in first-person versus third-person perspectives (Blanke, 2012; Blanke and Metzinger, 2009; Lenggenhager et al., 2009). We demonstrated above how the PCM can account for such illusions and phenomenal dissociations and how it can make sense of their paradoxical emergence in the context of non-ecological VR-mediated experimental manipula-

tions. VR could be further used to test hypotheses about real, normal and pathological human behavior based on direct predictions of the PCM as a function of its parameters. These types of manipulations in VR environments suggest experimental paradigms based on formal mathematical principles to quantitatively study processes such as optimal decision-making under stress as well as the structure and dynamics of consciousness.

4.4. Limitations

The model proposed here is only a starting point and a very primitive version of what we ultimately expect from a mathematical and computational model of consciousness in terms of complexity and formal rigor. However, we see the principles introduced herein as a basis for constraining more comprehensive models of consciousness and for exploring the relationships of the theory of consciousness to other scientific theories. Examples of the diverse applications and mathematical structures it implicates vary from topology to group theory, dynamical systems theory and beyond. But there are also connections to be made here with clinical and even pedagogical questions.

An essential question that remains open is that of the type of “field” and support the projective space is operating on in consciousness: the very fabric, so to speak, of the conscious space (e.g., a field of numbers?). If projective transformations operate in conscious access and provide us with the rules for the rendering of key spatial aspects of subjective experience in consciousness, the ontology or material support of the interface or medium for this rendering remains mysterious. We suggested that the whole workspace was somehow “virtual”, a sort of “user interface”, but this does not fully address the question. Images can be rendered and projected on a screen, but what constitutes the 3-dimensional “screen” of consciousness will remain unsettled here. From this standpoint, it is however important to mention the fundamental result established by Hilbert and Artin (Artin, 2016) according to which geometrical axioms of incidence that do not use any kind of numbers in their formulation (like Desargues’ and Pappus’ axioms) are sufficient to generate by themselves a commutative field of numbers. Numbers appear in an elegant manner through automorphisms of special automorphisms. Therefore, the construction of projective spaces from linear algebra does not lack generality. Moreover, projective geometry can contribute to the emergence of implicit computations, including comparison of ratios of length, like harmonic proportions (Coxeter, 1961).

It is certain that the structure of the conscious workspace is more complex than what this model captures. There is evidence and a rationale for suggesting that visual awareness, for instance, embeds a hierarchy of structuring spaces, including, at the top of the hierarchy, abstracted spaces representing the Euclidean properties of the world (Koenderink and van Doorn, 2008). However, projective geometry is a quite general geometry. When associated with a continuous field with metrical properties, it actually embeds affine and Euclidean geometry, and orthographic projections respecting Euclidean properties can be generated from projective geometry by assuming a vantage point near infinity.

The relationships between the PCM and the dynamical and intensive qualities of consciousness need more elaboration. There are dynamical components in our model, since it is based on active inference, as driven by free energy, which defines the principle and driving force for intentionality and action and implements a field of motivation based on the spatial mapping of personal relevance. The model features a complex interplay between cycles of perception, imagination, motor programming, and action, through projective transformations and complex mechanisms of inference related to personal relevance. These we expect must fundamentally constrain the pace of time-consciousness, as a moment-to-moment

construction of unified experiences involving “protection and retention” (Husserl, 1991). But the exact dynamics that link the inference engine and the projective geometry engine remains largely undefined and underdetermined: it will depend on the level of sophistication and specifications of future implementations of the basic kernel of the model and probably on the species of consciousness targeted. The status of intensive dimensions, e.g., connected to arousal, vigilance and emotions (Damasio and Rudrauf, 2006), as well as their differentiation into a range of qualitatively different emotional states, need to be worked out, starting from their connection to free energy (Cunningham et al., 2013; Joffily and Coricelli, 2013). Phenomenologically, we proposed that several aspects of bodily feelings are expressed as spatial transformations of the representation of the body in the workspace. To some extent, intensive dimensions, e.g., as with the feeling of more intense cardiac sensations when the heart rate increases (Khalsa et al., 2009), can be construed as spatial representations, e.g., as the heartbeat increasingly occupies our attention and interferes with our ability to carry on other mental tasks. In other words, we expect that some intensive dimensions of subjective experience will reduce to spatial quantities that are consciously accessed as such (e.g., the amplitude of a change in posture as a function of the intensity of an emotion) in connection to a field of personal relevance.

Nothing has been said here about many other key dimensions of experience: the gamut of sensory qualities, from colors to tones and smells. Even though, again, one might argue that essential aspects of the phenomenology of vision, smell and sound have a spatial manifestation, some key phenomenal features, like color, tone, and smell qualia, remain outside of the explanatory reach of the model as it is; however, they all certainly often carry personal relevance as sensory qualities.

Another significant missing item in our theory is the status and role of language in the workspace, how it interacts with the mechanism of conscious access and how symbolic representations (e.g., words) and projective imagination or mental imagery mutually influence each other. This is an important question, but it is beyond the scope of this contribution. It will be the topic of a future communication, as we have specific hypotheses about this issue (Rudrauf, 2014). We note that active inference as a generalization of Bayesian inference can be set up to operate based on propositional models and logical reasoning.

Finally, the PCM does not tell us how subjectivity, sensory qualities and representations with a phenomenal character relate, in their very phenomenal properties, to the computational and biophysical properties of the brain. Even though the model features a key relationship between computational efficacy and the mode of phenomenal access to spatial information in consciousness and can be used to derive neurocomputational hypotheses, it neither solves the Hard Problem (Chalmers, 1995) nor closes the Explanatory Gap (Levine, 1983) completely. On the other hand, it could well be that the construction of mathematical models that allow one to intelligibly link the phenomenology of conscious experience to experimentally confirmed neurocomputational implementations of the models are all that one could or should reasonably expect in this regard. In the absence of compelling arguments for dualism or panpsychism, the neurophenomenological method suggested here can be taken as a key to the development of a physicalist theory of consciousness (see Oizumi et al., 2014; Petitot et al., 1999; Rudrauf et al., 2003; Varela, 1996; 1999; Williford, 2006; 2015). And even in the presence of such arguments, since only few would deny that every difference in consciousness corresponds to a difference in brain processing, the method and the resulting models retain their full significance.

4.5. Future directions for the model

Future versions of the kernel of our reduced-model will integrate a more differentiated and sophisticated repertory of distributions of personal relevance as related to shapes/morphologies, sets of basic actions, physical and social models. They will also integrate a full 3-dimensional projective geometry engine coupled to a virtual world containing meshes of environments that will define a support for fields of conditional probabilities related to personal relevance. They will feature active inferences about new structures in space to enrich the world-model through multiview reconstruction processing. The world-map itself can be dynamic and easily updated to create events that force the agent to a change in perspective. Complex structures can be modeled with attached personal relevance maps. The kernel will approach and avoid these structures accordingly, e.g., avoiding walls, approaching doors (an implicit collision avoidance mechanics is embedded in the kernel; it is all a matter of setting up probabilities).

Already with the reduced-model, the instantiated agents can be rendered more or less intelligent and resilient, depending on the tuning of several parameters such as the rate and duration of pure imagination-mode, the prior updating rate, and the initial conditions in the situation of the agents, in terms of both factual and believed given information, along with the agents' developmental history. But the capacity of their consciousness will be enhanced and enriched by explicitly introducing into the model a layer of appraisal of self-relevance and self-concern based on a mechanism of perspective taking associated with the putative standpoint of other agents, i.e., a model of intersubjectivity based on Theory of Mind. This will require embedding and instantiating virtual simulations of other agents' consciousnesses in the algorithm of the agent itself based on the same principles used for the basic model.

We can envision using collectivities of such agents in interaction (they would be quite tractable computationally), which would have prior biases towards each other, e.g., some leading to mutual attraction, others to mutual repulsion, and others to stalking behaviors. Agents “liking” each other, could easily share their knowledge (prior maps) based on “trust”, and mutually update their prior maps through simple message-passing algorithms. In cases of disappointment, i.e., a mismatch between the information provided and actual sensory evidence, the probabilities of personal relevance attached to the agent source of information would naturally be negatively tuned down. Interesting social behaviors should emerge.

4.6. Conclusion and perspectives

The present Projective Consciousness Model contributes to the overarching goal of mathematizing phenomenology (Oizumi et al., 2014; Petitot et al., 1999). One important incentive for deriving generative models of consciousness is the possibility of using such models as mathematical tools for the development of a phenomenological psychology, based on sound formal and computational foundations. If our model contains some truth, a future formal psychological science may include the study (and the classification) of possible conscious states based on projective solutions from projective geometry in the context of processes of active inference driven by the minimization of free energy. Computational implementation of the PCM can be used to predict human behavior as a function of parameters across ranges of appraisal profiles, in normal and pathological contexts, with broad clinical perspectives. Another connected incentive for deriving generative models is to make connections with other fields of science that build generative models (physics, chemistry, biology, computer science), and to make connections with applications, including clinical and pedagogical models and applications. We believe that our model offers

a promising toolkit and potential formal foundations for psychology, neuroscience and artificial intelligence.

Acknowledgments

We thank for their help at various stages of the elaboration and writing of this work: Guillaume Marrelec, Andrea Samson, Sahib Khalsa, Jean Petitot, Will Cunningham, Didier Grandjean, David Sander, Kira Vinci, Fritz Lebowsky, Manfred Frank, Katharina Ohrnberger, Ben Meuleman, Manik Battacharjee, Orfeas Georgatos, Ryan Murray, Thomas Metzinger, Uriah Kriegel, Josh Weisberg... *In Memoriam*: Lucien, Charles, Charles, Jacques et Dominique Rudrauf, Kenneth Williford Sr, Tomis Kapitan, Francisco Varela.

References

- Adams, R.A., Shipp, S., Friston, K.J., 2013. Predictions not commands: active inference in the motor system. *Brain Struct. Funct.* 218 (3), 611–643.
- Albert, M.K., Hoffman, D.D., 2000. The generic-viewpoint assumption and illusory contours. *Perception* 29 (3), 303–312.
- Andersen, R.A., 1997. Multimodal integration for the representation of space in the posterior parietal cortex. *Philos. Trans. R. Soc. London B* 352 (1360), 1421–1428.
- Apps, M.A., Tsakiris, M., 2014. The free-energy self: a predictive coding account of self-recognition. *Neurosci. Biobehav. Rev.* 41, 85–97.
- Artin, E., 2016. *Geometric Algebra*. Courier Dover Publications.
- Baars, B.J., 2005. Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Prog. Brain Res.* 150, 45–53.
- Baars, B.J., Franklin, S., 2009. Consciousness is computational: the LIDA model of global workspace theory. *Int. J. Mach. Conscious.* 1 (01), 23–32.
- Barrett, L.F., Simmons, W.K., 2015. Interoceptive predictions in the brain. *Nat. Rev. Neurosci.* 16 (7), 419–429.
- Bartolomeo, P., Chokron, S., 2002. Orienting of attention in left unilateral neglect. *Neurosci. Biobehav. Rev.* 26 (2), 217–234.
- Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., Friston, K.J., 2012. Canonical microcircuits for predictive coding. *Neuron* 76 (4), 695–711.
- Bennequin, D., 2014. Remarks on invariance in the primary visual systems of mammals. In: *Neuromathematics of Vision*. Springer, pp. 243–333.
- Berthoz, A., 1997. *Sens du mouvement* (Le). Odile Jacob.
- Blanke, O., 2012. Multisensory brain mechanisms of bodily self-consciousness. *Nat. Rev. Neurosci.* 13 (8), 556–571.
- Blanke, O., Metzinger, T., 2009. Full-body illusions and minimal phenomenal selfhood. *Trends Cogn. Sci.* 13 (1), 7–13.
- Block, N., 1995. How many concepts of consciousness? *Behav. Brain Sci.* 18 (02), 272–287.
- Buckner, R.L., Andrews-Hanna, J.R., Schacter, D.L., 2008. The brain's default network. *Ann. N. Y. Acad. Sci.* 1124 (1), 1–38.
- Butchvarov, P., 1979. *Being Qua Being: A Theory of Identity, Existence, and Predication*. Indiana University Press.
- Butchvarov, P., 1998. *Skepticism about the external World*. Oxford University Press.
- Campbell, J.O., 2016. Universal Darwinism as a process of Bayesian inference. *Front. Syst. Neurosci.* 10.
- Chalmers, D.J., 1995. Facing up to the problem of consciousness. *J. Conscious. Stud.* 2 (3), 200–219.
- Clark, A., 2009. Spreading the joy? Why the machinery of consciousness is (probably) still in the head. *Mind* 118 (472), 963–993.
- Clark, A., 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav. Brain Sci.* 36 (03), 181–204.
- Clark, A., 2015. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.
- Conant, R.C., Ross Ashby, W., 1970. Every good regulator of a system must be a model of that system. *Int. J. Syst. Sci.* 1 (2), 89–97.
- Coxeter, H.S.M., 1961. *Introduction to geometry*.
- Crick, F., Koch, C., 2003. A framework for consciousness. *Nat. Neurosci.* 6 (2), 119–126.
- Cronin, T.W., 2005. The visual ecology of predator-prey interactions. *Ecol. Predator-Prey Interact.* 105–138.
- Cunningham, W.A., Dunfield, K.A., Stillman, P.E., 2013. Emotional states from affective dynamics. *Emot. Rev.* 5 (4), 344–355.
- Damasio, A., Rudrauf, D., 2006. *The biological basis of subjectivity: a hypothesis. Self-Representational Approaches to Consciousness*. Cambridge MA, MIT Press/Bradford Books.
- Damasio, A.R., 1999. *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. Houghton Mifflin Harcourt.
- Dennett, D., 1991. *Consciousness Explained*. Boston, MA: Little Brown.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., Sergent, C., 2006. Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends Cogn. Sci.* 10 (5), 204–211.
- Dehaene, S., Kerszberg, M., Changeux, J.-P., 1998. A neuronal model of a global workspace in effortful cognitive tasks. *Proc. Natl. Acad. Sci.* 95 (24), 14529–14534.
- Douglas, H., 2007. *I am a Strange Loop*. Basic Books.
- Droulez, J., Berthoz, A., 1991. The concept of dynamic memory in sensorimotor control. In: *Motor Control: Concepts and Issues*, pp. 137–161.
- Ehrsson, H.H., 2007. The experimental induction of out-of-body experiences. *Science* 317 (5841), 1048–1048.
- Elliot, A.J., 2006. The hierarchical model of approach-avoidance motivation. *Motiv. Emot.* 30 (2), 111–116.
- Feldman, H., Friston, K., 2010. Attention, uncertainty, and free-energy. *Front. Hum. Neurosci.* 4, 215.
- Frank, M., 2004. Fragments of a history of the theory of self-consciousness from Kant to Kierkegaard. *Crit. Horiz.* 5 (1), 53–136.
- Freeman, W.J., 1999. Consciousness, intentionality and causality. *J. Conscious. Stud.* 6 (11–12), 143–172.
- Fregnac, Y., 2003. Neurogeometry and entoptic visions of the functional architecture of the brain. *J. Physiol.-Paris* 97 (2), 87–92.
- Friston, K., 2005. A theory of cortical responses. *Philos. Trans. R. Soc. London B* 360 (1456), 815–836.
- Friston, K., 2008. Hierarchical models in the brain. *PLoS Comput. Biol.* 4 (11), e1000211.
- Friston, K., 2010. The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* 11 (2), 127–138.
- Friston, K., Adams, R., Perrinet, L., Breakspear, M., 2012. Perceptions as hypotheses: saccades as experiments. *Front. Psychol.* 3, 151.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G., 2016. Active inference: a process theory. *Neural Comput.* 29 (1), 1–49.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., FitzGerald, T., Pezzulo, G., 2015. Active inference and epistemic value. *Cogn. Neurosci.* 6 (4), 187–214.
- Giummarra, M.J., Moseley, G.L., 2011. Phantom limb pain and bodily awareness: current concepts and future directions. *Curr. Opin. Anesthesiol.* 24 (5), 524–531.
- Gray, J.A., 1990. Brain systems that mediate both emotion and cognition. *Cogn. Emot.* 4 (3), 269–288.
- Greenberg, M.J., 1993. *Euclidean and Non-Euclidean Geometries: Development and History*. Macmillan.
- Gurwitsch, A., 1964. *The Field of Consciousness*. Pittsburgh, PA: Duquesne University Press.
- Hartley, R., Zisserman, A., 2003. *Multiple View Geometry in Computer Vision*. Cambridge university press.
- Hassabis, D., Maguire, E.A., 2007. Deconstructing episodic memory with construction. *Trends Cogn. Sci.* 11 (7), 299–306.
- Henrich, D., 1966. Fichte's original insight. *Contemp. Ger. Philos.* 1, 15–53.
- Herbelin, B., Vexo, F., Thalmann, D., 2002. Sense of presence in virtual reality exposures therapy. In: *Proceedings of the 1st International Workshop on Virtual Reality Rehabilitation*, Lausanne, Switzerland. Citeseer.
- Hicheur, H., Kadone, H., Grezes, J., Berthoz, A., 2013. Perception of emotional gaits using avatar animation of real and artificially synthesized gaits. In: *Affective Computing and Intelligent Interaction (ACII), 2013 Humaine Association Conference on*. IEEE, pp. 460–466.
- Hohwy, J., 2013. *The Predictive Mind*. Oxford University Press.
- Hohwy, J., 2014. Elusive phenomenology, counterfactual awareness, and presence without mastery. *Cogn. Neurosci.* 5 (2), 127–128.
- Hopkins, J., 2016. Free energy and virtual reality in neuroscience and psychoanalysis: a complexity theory of dreaming and mental disorder. *Front. Psychol.* 7.
- Hume, D., 1739. *A Treatise of Human Nature*. Oxford University Press.
- Humphrey, N., 2003. *The Inner Eye: Social Intelligence in Evolution*. Oxford University Press.
- Husserl, E., 1913. *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy: First Book: General Introduction to a Pure Phenomenology*, 1st, Springer.
- Husserl, E., 1991. *On the Phenomenology of the Consciousness of Internal Time (1893-1917)*. Springer.
- Itti, L., Baldi, P., 2009. Bayesian surprise attracts human attention. *Vision Res.* 49 (10), 1295–1306.
- James, W., 1913. *The Principles of Psychology*. Read Books Ltd.
- Joffily, M., Coricelli, G., 2013. Emotional valence and the free-energy principle. *PLoS Comput. Biol.* 9 (6), e1003094.
- Kalisch, R., Müller, M.B., Tüscher, O., 2015. A conceptual framework for the neurobiological study of resilience. *Behav. Brain Sci.* 38, e92.
- Kant, I., 1999. *Critique of Pure Reason*. Cambridge University Press.
- Kappas, A., 1996. The sociality of appraisals: impact of social situations on the evaluation of emotion antecedent events and physiological and expressive reactions. In: *ISRE 1996: Proceedings of the IXth conference of the International Society for Research on Emotions*, pp. 116–120.
- Kelley, L.A., Endler, J.A., 2012. Illusions promote mating success in great bowerbirds. *Science* 335 (6066), 335–338.
- Khalsa, S.S., Rudrauf, D., Feinstein, J.S., Tranel, D., 2009. The pathways of interoceptive awareness. *Nat. Neurosci.* 12 (12), 1494–1496.
- Kihlstrom, J.F., Cantor, N., 1984. Mental representations of the self. *Adv. Exp. Soc. Psychol.* 17, 1–47.
- Kim, C.-Y., Blake, R., 2005. Psychophysical magic: rendering the visible 'invisible'. *Trends Cogn. Sci.* 9 (8), 381–388.
- Kind, A., 2016. *The Routledge Handbook of Philosophy of Imagination*. Routledge.
- Kitanishi, T., Ito, H.T., Hayashi, Y., Shinohara, Y., Mizuseki, K., Hikida, T., 2016. Network mechanisms of hippocampal laterality, place coding, and goal-directed navigation. *J. Physiol. Sci.* 1–12.
- Klinger, E., 1971. *Structure and Functions of Fantasy*. Wiley-Interscience.
- Klüver, H., 1926. Mescal visions and eidetic vision. *Am. J. Psychol.* 37 (4), 502–515.

- Klüver, H., 1966. *Mescal, and Mechanisms of Hallucinations*. [Chicago]: University of Chicago Press.
- Knill, D.C., Pouget, A., 2004. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27 (12), 712–719.
- Koenderink, J., van Doorn, A., 2008. The structure of visual spaces. *J. Math. Imaging Vision* 31 (2–3), 171–187.
- Koenderink, J., van Doorn, A., de Ridder, H., Oomes, S., 2010. Visual rays are parallel. *Perception* 39 (9), 1163–1171.
- Kriegel, U., 2009. *Subjective Consciousness: A Self-Representational Theory*. Oxford University Press.
- Kriegel, U., 2012. Précis of subjective consciousness: a self-representational theory. *Philos. Stud.* 159 (3), 443–445.
- Kriegel, U., Williford, K., 2006. *Self-Representational Approaches to Consciousness*. MIT Press.
- Lacan, J., Fink, B., 1966, 2002. *Écrits: A Selection*. WW Norton & Company.
- Lambrey, S., Doeller, C., Berthoz, A., Burgess, N., 2012. Imagining being somewhere else: neural basis of changing perspective in space. *Cereb. Cortex* 22 (1), 166–174.
- Lamm, C., Batson, C.D., Decety, J., 2007. The neural substrate of human empathy: effects of perspective-taking and cognitive appraisal. *J. Cogn. Neurosci.* 19 (1), 42–58.
- Legrand, D., Brozzoli, C., Rossetti, Y., Farne, A., 2007. Close to me: multisensory space representations for action and pre-reflexive consciousness of oneself-in-the-world. *Conscious. Cogn.* 16 (3), 687–699.
- Lenggenhager, B., Mouthon, M., Blanke, O., 2009. Spatial aspects of bodily self-consciousness. *Conscious. Cogn.* 18 (1), 110–117.
- Levine, J., 1983. Materialism and qualia: the explanatory gap. *Pac. Philos. Q.* 64 (4), 354–361.
- Limanowski, J., Blankenburg, F., 2013. Minimal self-models and the free energy principle.
- Manstead, A.S., Fischer, A.H., 2001. Social appraisal: the social world as object of and influence on appraisal processes. In: *Appraisal Processes in Emotion: Theory, Methods, Research*, pp. 221–232.
- McCormack, G.L., 2006. *Fusion and Binocularity*. In: Benjamin, W.J. (Ed.), *Borish's Clinical Refraction*. WB Saunders Company, Philadelphia.
- McGinn, C., 2009. *Mindsight: Image, Dream, Meaning*. Harvard University Press.
- Merleau-Ponty, M., 1945. *Phenomenology of Perception*. New York: Routledge.
- Mesmoudi, S., Perlbarg, V., Rudrauf, D., Messe, A., Pinsard, B., Hasboun, D., Cioli, C., Marrelec, G., Toro, R., Benali, H., et al., 2013. Resting state networks' corticotopy: the dual intertwined rings architecture. *PLoS ONE* 8 (7), e67444.
- Metzinger, T., 2004. *Being No One: The Self-Model Theory of Subjectivity*. MIT Press.
- Metzinger, T., 2005. Out-of-body experiences as the origin of the concept of 'a-soul'. *Mind Matter* 3 (1), 57–84.
- Mirza, M.B., Adams, R.A., Mathys, C.D., Friston, K.J., 2016. Scene construction, visual foraging, and active inference. *Front. Comput. Neurosci.* 10.
- Moors, A., Scherer, K.R., 2013. The role of appraisal in emotion. In: *Handbook of Cognition and Emotion*, pp. 135–155.
- Murray, R.J., Gerrans, P., Brosch, T., Sander, D., 2015. When at rest: "event-free" active inference may give rise to implicit self-models of coping potential. *Behav. Brain Sci.* 38 (e114), 41–42.
- do Nascimento, L.M., Noë, A., 2012. Out of our heads: why you are not your brain, and other lessons from the biology of consciousness. *Principia* 16 (3), 495.
- Oizumi, M., Albantakis, L., Tononi, G., 2014. From the phenomenology to the mechanisms of consciousness: integrated information theory 3.0. *PLoS Comput. Biol.* 10 (5), e1003588.
- O'Keefe, J., Burgess, N., Donnett, J.G., Jeffery, K.J., Maguire, E.A., 1998. Place cells, navigational accuracy, and the human hippocampus. *Philos. Trans. R. Soc. B* 353 (1373), 1333–1340.
- Paradiso, S., Rudrauf, D., 2012. Struggle for life, struggle for love and recognition: the neglected self in social cognitive neuroscience. *Dialogues Clin. Neurosci.* 14 (1), 65–75.
- Peirce, C.S., 1955. *The principles of phenomenology*. Philosophical Writings of Peirce. Dover Publications, New York.
- Penfield, W., Boldrey, E., 1937. Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain*.
- Petitot, J., Varela, F.J., Pachoud, B., Roy, J.-M., (Eds.), 1999. *Naturalizing Phenomenology: Issues in Contemporary Phenomenology and Cognitive Science*. Stanford University Press.
- Philippi, C.L., Feinseth, J.S., Khalsa, S.S., Damasio, A., Tranel, D., Landini, G., Williford, K., Rudrauf, D., 2012. Preserved self-awareness following extensive bilateral brain damage to the insula, anterior cingulate, and medial prefrontal cortices. *PLoS ONE* 7 (8), e38413.
- Poincaré, H., 2015. *La science et l'hypothèse*. Booklassic.
- Rao, R.P., Ballard, D.H., 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2 (1), 79–87.
- Reichle, R.A.E.S.J.W., Erik, D., 2010. Eye movements during mindless reading. *Psychol. Sci.* 21 (9), 1300–1310.
- Rosenberg, M., 1965. *Society and the adolescent self-image*.
- Rudrauf, D., 2014. Structure-function relationships behind the phenomenon of cognitive resilience in neurology: insights for neuroscience and medicine. *Adv. Neurosci.* 2014.
- Rudrauf, D., David, O., Lachaux, J.-P., Kovach, C.K., Martinerie, J., Renault, B., Damasio, A., 2008. Rapid interactions between the ventral visual stream and emotion-related structures rely on a two-pathway architecture. *J. Neurosci.* 28 (11), 2793–2803.
- Rudrauf, D., Lutz, A., Cosmelli, D., Lachaux, J.-P., Le Van Quyen, M., 2003. From autopoiesis to neurophenomenology: Francisco Varela's exploration of the biophysics of being. *Biol. Res.* 36 (1), 27–65.
- Sander, D., Grafman, J., Zalla, T., 2003. The human amygdala: an evolved system for relevance detection. *Rev. Neurosci.* 14 (4), 303–316.
- Sartre, J.-P., 1937. *The Transcendence of the Ego: A Sketch for a Phenomenological Description*. Routledge.
- Sartre, J.-P., 2004. *The Imaginary: A Phenomenological Psychology of the Imagination*. Routledge Press.
- Sartre, J.-P., 2012. *The Imagination*. London: Routledge. (Translation Williford, K., Rudrauf, D.).
- Scherer, K.R., 2009. The dynamic architecture of emotion: Evidence for the component process model. *Cogn. Emot.* 23 (7), 1307–1351.
- Schöne, H., 1984. *Spatial Orientation: The Spatial Control of Behavior in Animals and Man*. Princeton University Press.
- Seth, A.K., 2014. The cybernetic Bayesian brain. *Open Mind. Open MIND*. Frankfurt am Main: MIND Group.
- Seth, A.K., 2015. Presence, objecthood, and the phenomenology of predictive perception. *Cogn. Neurosci.* 6 (2–3), 111–117.
- Seth, A.K., Friston, K.J., 2016. Active interoceptive inference and the emotional brain. *Phil. Trans. R. Soc. B* 371 (1708), 20160007.
- Siegel, R.K., 2005. *Intoxication: The Universal Drive for Mind-Altering Substances*. Inner Traditions/Bear & Co.
- Smallwood, J., Schooler, J.W., 2006. The restless mind. *Psychol. Bull.* 132 (6), 946.
- Stratton, G., 1899. The spatial harmony of touch and sight. *Mind* 8 (4), 492–505.
- Thiroux, B., Mercier, M.R., Jorland, G., Berthoz, A., Blanke, O., 2010. Mental imagery of self-location during spontaneous and active self-other interactions: an electrical neuroimaging study. *J. Neurosci.* 30 (21), 7202–7214.
- Tononi, G., 2004. An information integration theory of consciousness. *BMC Neurosci.* 5 (1), 42.
- Tozzi, A., Peters, J.F., 2016. Towards a fourth spatial dimension of brain activity. *Cogn. Neurodyn.* 10 (3), 189–199.
- Tyler, C.W., 1978. Some new entoptic phenomena. *Vision Res.* 18 (12), 1633–1639.
- Ungerleider, L.G., 1982. Two cortical visual systems. *Anal. Visual Behav.* 549–586.
- Ungerleider, L.G., Haxby, J.V., 1994. 'what' and 'where' in the human brain. *Curr. Opin. Neurobiol.* 4 (2), 157–165.
- Varela, F., 1979. *Principles of Biological Autonomy*. New York: Elsevier/North-Holland.
- Varela, F.J., 1996. Neurophenomenology: a methodological remedy for the hard problem. *J. Conscious. Stud.* 3 (4), 330–349.
- Varela, F.J., 1999. The specious present: a neurophenomenology of time consciousness. *Naturalizing Phenomenol.* 266–314.
- Vossel, S., Mathys, C., Stephan, K.E., Friston, K.J., 2015. Cortical coupling reflects Bayesian belief updating in the deployment of spatial attention. *J. Neurosci.* 35 (33), 11532–11542.
- Wilkerson, W., 2010. Time and ambiguity: reassessing Merleau-Ponty on Sartrean freedom. *J. Hist. Philos.* 48 (2), 207–234.
- Williford, K., 2010. Pre-reflective self-consciousness and the autobiographical ego. In: Webber, J. (Ed.), *Reading Sartre: On Phenomenology and Existentialism*. Routledge.
- Williford, K., 2006. The self-representational structure of consciousness. In: Kriegel, U., Williford, K. (Eds.), *Self-Representational Approaches to Consciousness*. The MIT Press.
- Williford, K., 2013. Husserl's hyletic data and phenomenal consciousness. *Phenomenol. Cogn. Sci.* 12 (3), 501–519.
- Williford, K., 2015. Representationalisms, subjective character, and self-acquaintance. *Open MIND. Open MIND*. Frankfurt am Main: MIND Group.
- Williford, K., Rudrauf, D., Landini, G., 2012. The paradoxes of subjectivity and the projective structure of consciousness. *Consciousness and Subjectivity*.
- Wittgenstein, L., 1922. *Tractatus Logico-Philosophicus*. London: Routledge.
- Yosinski, J., Clune, J., Nguyen, A., Fuchs, T., Lipson, H., 2015. *Understanding neural networks through deep visualization*. arXiv preprint arXiv:1506.06579.
- Yuille, A., Kersten, D., 2006. Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* 10 (7), 301–308.
- Zacks, J.M., 2008. Neuroimaging studies of mental rotation: a meta-analysis and review. *J. Cogn. Neurosci.* 20 (1), 1–19.
- Zahavi, D., 1999. *Self-Awareness and Alterity: A Phenomenological Investigation*. Northwestern University Press.
- Zahavi, D., 2008. *Subjectivity and Selfhood: Investigating the First-Person Perspective*. MIT press.
- Zeidman, P., Lutti, A., Maguire, E.A., 2015. Investigating the functions of subregions within anterior hippocampus. *Cortex* 73, 240–256.