Information Flow, Representation, and Awareness

Ernest W. Kent, *ewk@uchicago.edu.*

Currently Retired. Formerly: Chief Scientist, Intelligent Systems Division, The National Institute of Standards and Technology. Formerly: Director of Research in Information Technology for N. America, Phillips, N.V. Formerly: Associate Professor of Physiological Psychology and Psychopharmacology, The University of Illinois at Chicago.

Abstract

I begin with a concept of representation based on the flow of the inherent information contained in physical objects and argue that it is apropos for both physical objects and phenomenal character. The implications of this are explored in the context of awareness viewed as a mental representation of the world, in which "bottom-up", non-conceptual representation interacts with "top-down", intentional[1] states to produce a hierarchical representation from which both, and their mixed states, are simultaneously accessible, and which is actively maintained by Bayesian control theory principles. This view is then examined in the context of current neurobiological investigations of sensory and perceptual processes. I suggest that this approach, although not fundamentally at odds with many concepts in philosophy, may offer a more coherent way to view mind-dependent aspects of perception than many traditional philosophical positions, which do not align easily with the concepts of empirical neurobiology.

Overview

It has been remarked, (Hatfield 2011) that in the philosophy of perception, the majority of positions currently in favor espouse some form of direct realism, and that indirect or representative realism is largely abandoned. I'm not aware of any actual census of philosophers on the subject, but it may be so. Among empirical scientists on the other hand (again without a formal census) it seems that not only are the majority unaware of this state of affairs in philosophy, but most when informed of it receive the news with either mirth or incredulity. Clearly there is some need for a meeting of minds. Additionally, there are a number of concepts in philosophy regarding perception which are often presented as opposing views, whereas the theoretical models embraced by the majority of neurobiologists, control theorists, and some cognitive scientists e.g., (Raftopoulos and Muller 2006) often may be seen as containing ideas at least similar to both positions, but simply embodied in different aspects of the model's operation. Examination of such instances may provide avenues for progress.

I will attempt to assemble a view of representation in perceptual experience that emerges from neurobiology, control theory, and information theory. All of these fields have something of interest to contribute to the discussion, frequently in ways that overlap only partially or not at all with hotly-debated positions in philosophy, and almost never

1 To be understood in this paper in the sense of being about things, properties and states of affairs.

break apart along precisely the same lines. The exception is that such a perspective is at its core representational rather than relational.

I begin (Sections 1-4) by developing a definition of representation derived from the information necessary to describe an object, and the requirements for its causal transmission into perceptual awareness. On this basis I conclude that awareness must itself be considered a representation, and further, that there can be no final intermediate representation preceding awareness. I then (Sections 5-6) consider in some detail two potential problems: representation under scale transformations, and representation of space and time. From this perspective, I present and discuss (Sections 7-9) some of the major ideas current in the neurobiological community, and in particular the ways in which they are informed by control theory. Finally (Sections 10-11), I consider what implications may be drawn from this material and a few of the possible mappings between disciplines.

1 Varieties of Information

The common, modern usage of "information" as "knowledge communicated concerning a particular topic" corresponds fairly closely to our technical concept of "semantic information." The literature today includes many other specialized terms such as Shannon information, natural and non-natural information, structural information, free and bound information, intrinsic information, algorithmic information, physical information, quantum information, and others. Usage of these terms often has significant overlaps and not all authors use them in the same way. Confusion is likely if one is not very clear about the kinds of information one is discussing, so let me spend a moment on that.

We require a conception of information appropriate to discussing both the information content of mind-independent physical things and of mind-dependent representations of them. Because I will need to deal with how information about the world enters mental representations, it is important to understand in what sense I mean information that is in some way "part of the world". We do not usually think of natural objects in and of themselves as embodying "meaning" in the sense of semantic information as a property. After all, they existed long before anyone was around with the cognitive capability to experience meaning. The mathematical information theory of Shannon, while clearly applying to events in the physical world, is explicitly devoid of semantic content. [2]

Somewhere between semantic meaning and Shannon's theory of information we feel the need for a property of physical objects or events that instantiates what can be transmitted by causal processes that enter conscious awareness, and to which meaning may become attached. A piece of rock may have a differentiated structure that has been around for billions of years before any cognitive agent existed, yet the inter-related details of that structure may have meaning for a geologist. Some may prefer to call this "data" or perhaps "information content". Collier attempts to capture this in a formal theoretical framework which he calls, "Intrinsic Information."(Collier 1990) He says:

[2] Notwithstanding, mathematical information theory is employed by many projects seeking to provide a naturalized account of cognition. See Adams, (Adams 2003) for a very readable review.

"… if we assume even a modest realism there must be some property of objects that allows us to have information about them. This property must be causally based and causally communicable to us, as well as being commensurate with information in the vulgar sense."

He develops a formal definition of this property which has many concepts in common with Shannon, and asserts that, like Shannon's information, "… what is transmitted is information, irrespective of whether or not there is a cognitive receiver." Because standard information theory is based on probabilities and applies to ensembles, Collier appeals to Kolmogorov's "Algorithmic Information". This is essentially a way to measure information content of an object in terms of the minimum procedure that can reproduce it. It has a number of advantages for Collier's program including especially that it can apply to individual cases. Further, probability theory can be based on algorithmic information theory, providing a link to Shannon. (See Grunwald and Vitanyi (Grunwald and Vitanyi 2008)for an exposition.)

Collier uses the terms "Bottom up" and "Top Down" information to connote intrinsic information and cognitive, semantic information respectively, and this is very much in the spirit of the control-theoretic ideas I will present below in regard to the internal model of the environment maintained in the brain and its role in the interaction of these two streams of information. I don't know whether or not Collier's formalism can carry the freight of explaining the connection of intrinsic to semantic information, but that is not important to my purpose. Rather I draw on the concept of something akin to his intrinsic information as a necessary addition to semantic or intentional-cognitive information on the one hand and meaning-free information in the sense of the mathematical theory of communication on the other. The kind of information I am intending is that inherent in the differential structure of a physical object's properties. This usage is also very similar to that in Chapter 8 of Chalmers (Chalmers 1997). It is the kind of information that distinguishes one part or region of a thing from another and which is embodied in its essence and required for its description.

Thus things have real properties that have measures associated with them, and these measures have values on one or more abstract dimensions. The measures are thus represented by vector-valued functions on a state space. This is a descriptive formalism for values that are realized in the properties of the real physical events they describe. This information space formalism can provide a helpful way to organize our thinking about both physically-realized and phenomenally-realized information, and I will make use of it in both cases. I may add that events in the world have structure in time as well as space and so the information spaces they realize must include temporal dimensions (even if the relevant temporal information is "No change".) I believe a physically-realized information space captures the essence of what I find appealing in collier's work without committing to his particular mathematical formalism. I will employ the term *inherent information* to refer to it.[3]

Specifically, I will use "*inherent information*" to refer to information that physical or mental spatio-temporal events realize in their properties, and also to the representation of

3 For no better reason than that it seems appropriate and I haven't (yet) found anyone else who has co-opted the term.

that information in data structures called information spaces, containing vector-valued measures of these properties. I will also have need to refer to cognitive, semantic information in the ordinary, intentional sense, and its interaction with inherent information. I will try to be explicit as to which I am referring when it is not obvious from context. I will also use the terms "bottom-up" and "top-down" where appropriate to distinguish the two, as these will be very descriptive when considering the control-theoretic setting.

2 Inherent Information Flow

Chalmers makes the point that these physically-realized values can only be *information* insofar as they can be processed (as is also true of Shannon information), and this entails information flow. In its simplest form, information flow in an information theoretical context is the transfer of information from a variable x to a variable y in a given process. (Wikipedia contributors 2021) It is most often considered in the context of secure information-routing in computer systems, but here I am retaining the simpler and broader definition of transfer of information from one variable to another, and particularly the case of transmission from one medium to another.

That is, information inherent in the structure of the object is transmitted by becoming inherent, through causal connection, in the structure of successive events, which events are also realizations of their own information spaces. The substrate of these realizations is irrelevant provided that they have the capability of realizing a sufficient information space. For example, a particular spot on a rock may have a color, by which is meant the property of differential reflectivity with regard to wavelength. Causal, physical interaction with this spot will affect the properties of the wavefront of electromagnetic radiation reflected off it in such a way that there exists a mapping from the information space of reflectivities over the surface of the rock to the information space of wavelengths contained in the wavefront. The material of the rock is not transmitted, nor is its property of reflectivity transmitted. What is transmitted is a deterministic change in the structure of the wavefront, and hence in the information space realized by the wavefront, such that there exists a lawful mapping from the information space of the rock to the information space of the wavefront. The wavefront may in turn encounter something else and causally alter that thing's information space. In particular it may encounter the retina of a cognitive observer, and a causal chain of events may be initiated in the observer's nervous system that ultimately changes the semantic information content of the observer.

It should be clear that at each step, at least up to the point of interaction with semantic information, there is a lawful relation between the information content of the preceding stage and changes to the information content of the succeeding stage. Causal chains of interaction between the realizing events which alter their properties effect lawful mappings from the inherent information space of the initial event to those of subsequent events in the chain, thereby realizing what I will refer to as "causal flow of inherent information". To say that there is a lawful mapping does not mean that the value or dimension of a property measure is transmitted. Indeed, information flow between different sorts of media would usually make this impossible. Rather, the state of the subsequent measure space depends upon the state of the first in such a way that it carries information about the state of the first because there is some function entailed by the causal action between the realizers, and perhaps only empirically-defined, which has as its domain the information space of the antecedent realizer and as its range the information space of the subsequent realizer. This sort of flow of inherent information requires continuity spatially and temporally in an unbroken causal chain because the *mechanism* of the flow is between the real events and not between the abstract information spaces of which they are realizers.[4]

3 Inherent Information as Representation

Confusion can arise from discussions of how mental or physical representations "look" or "appear." These terms are theory-laden and can generate questions as to what is doing the looking, leading to regression arguments.(R. French 2018) Further, they carry an implicit assumption that an external world possessing a "look" to be like is even a coherent notion, but there is no such "God's Eye" view to be had. The way electromagnetic fields and collections of subatomic particles in a four-dimensional spacetime would appear *if they could* be experienced "as they are, out there" might be just as we experience them through the senses, but there is no reason to think so, nor to know what it would even mean to experience them other than through the senses. To say that one sees the world as it is "out there" also carries the implication that one somehow knows what it would be like to experience the negative of that statement; that is, to be aware of any intervening representation, and somehow find it recognizable as *not* appearing as things are "out there." There is of course no way of knowing that the look of the external world you experience right now is not *exactly* your way of experiencing, for example, a pattern of neural firing. Representations may or may not "look like' the object in this naive sense, but even if they do, their status as representations must depend on something else.

There are many deep controversies in the philosophical literature about the nature of representation, particularly mental representation, along the lines of representationalism and phenomenalism, especially with regard to issues such as reductive and non-reductive views of phenomenal content, arguments about transparency, the nature of qualia and many others.(Pitt 2020) The account of representation I will offer here is not motivated by these concerns, does not arise out of them, and is not intended to address them, although some may view it as coinciding with one or another such position in whole or in part, and I will comment on this in relevant places. In general, in the view from empirical neurobiology, the functioning of the brain does not divide up very well along the lines of traditional philosophical distinctions, and the concept of representation I am putting on offer here is intended as a concept appropriate for a coherent view of information processing across physical events, neural processing, and phenomenal experience, from a control-theoretic viewpoint that suggests that several of the traditional positions of philosophy are

4 It is sometimes asserted that information can be transmitted otherwise. For example, if it is known that A causes both B and C, and C is observed, information about B is obtained. Or, one may learn something by the non-receipt of a phone call. In these cases however the effective information path seems still to be by continuous information flow. Thus, one has to have the knowledge that A causes B and that knowledge must have been obtained (unless we are to have an infinite regress) by some direct path. Hence the path from observing C to knowledge about B has taken a perhaps lengthy and tortuous but nonetheless causal route through the cognition and memory of the observer. A similar argument applies to the phone call. Prior information must have arrived concerning its expectation, and caused the observer's cognition concerning the silent phone. In any case, such instances are unlikely to be relevant to the flow of inherent information from an object to the observer's sensory apparatus.

best related to particular parts or stages of the process. In particular, while the concept of "high-level" and "low-level" categories in phenomenal content has been discussed – see (Bayne 2009) for a review – it has not often been considered in the context of hierarchical models of representation with interacting levels of description such as are common in control theory and neurobiology, and which will be the focus of the present discussion. In this case many common points of view from the philosophical literature may seem entirely appropriate at one point in the representation and their opposite at another.

On the view being offered, the essential property of being a representation of something by inherent information content is defined very loosely as the encoding of the description of something using another something as the encoding medium. The point of departure will be to consider such representations as a continuous causal chain of events commencing with the distal object and ending in phenomenal experience. There will be little dispute that such a continuous chain of events exists, but in what sense do they all qualify as representations? Depending on your views on physicalism vs. dualism, you may ask in what sense some of these events could even have comparable properties.

To refine this idea, I will define a representation of inherent information as anything which realizes in its properties the contents of an information space for which there exists a veridical map (in the mathematical sense) between some subset of the information space of the object being represented and the information space of its representation, provided this mapping is instantiated by a causal chain of inherent information flow.

The causal chain requirement prevents accidental correspondences counting as representations, although in principle it could permit something like a painting if you're up to the task of specifying the causal chain. It also eliminates cases of two items which correspond to one another only because they are both representations of another thing – for example, two castings from the same mold correspond to one another because their information spaces are both mappings of the information space of the mold, and are causally linked to it. But there is no causal chain between them, hence, under this definition, one is not a representation of the other.

The need to allow a subset of the information space of the object represented is not only due to the possibility of situations such as partial views, but also because all real channels are noisy and have losses. Thus, an object is likely to have reflectivity in wavelengths to which the eye is insensitive, and mapping from that part of the object's information space is lost to the information space of neural firing rates in the retina. As a consequence being a representation in this sense is not all-or-none; representations may be partial and information may be lost or distorted in transfer from one stage of representation to the next. Information originating in the distal object may become diluted by additional, unrelated information embodied in succeeding events until it perhaps makes only a minor contribution, or eventually none at all.

Since the mappings may be mathematically composited, this process may continue through many such stages of representation, resulting in a single end-to-end map. In the previous example, the property of an object of differentially reflecting wavelengths of light is encoded in a change in the mix of wavelengths of light in the wave front reflected from it. This in turn may then causally encode the value in geometric patterns of illuminance in an image formed on the retina by the lens and then in turn differential firing rates of different populations of cones in your retina and thence in firing patterns of ganglion cells in the retina, neurons of the lateral geniculate, area V1 of the visual cortex, and so on, and at every stage a composition of maps yields a map from the information space of the initial object represented to that of each subsequent stage. Clearly each representation in this continuum of representations may, but need not, “look like” the thing represented in the ordinary sense of the term, and in fact makes no assumptions at all about what a representation “looks like.”

This manner of defining representation provides a nice way to discuss representation in perception because mind-independent objects in the external world and mind-dependent perceptions can be represented in the same manner, so that it is applicable to at least some kinds of mental representation. In particular, low-level phenomenal properties have inherent information in the form of differential structure of their properties (e.g., change from red to blue in a given direction), and the topological and temporal relations holding amongst them. This structure must realize some information space in the same manner as do physical objects. If there is a causal (if nomic) dependence and if a veridical map can be composited, then there is no bar to regarding the mental object as a representation of the distal, physical object. Certainly difference in the underlying, realizing substances can not even enter into the question, since it does not do so anywhere else in the causal flow of information.

It is widely held that something is a representation only if it represents some portion of the world in a truth-evaluable way. (Chalmers 2004; Pitt 2020; Bailey 2007; Hutto 2009) The veridicality of a mapping between the information space realized by a potential representation and that realized by the object it putatively represents is a proposition which is open to evaluation. For the present definition of representation in terms of inherent information this is clearly possible in the case of representations of one physical entity by another – just map the measurements of their physical properties. In the case of phenomenal representation, information enters our experience initially when realized in phenomenal properties of sensory experiences – qualia[5]. Some aspects of this phenomenal experience, for example the topological relationships among its discriminably differing parts, seem easily capable of evaluation with respect to the distal object comprising the other end of the causal chain. Other aspects, for example our experience of the color of some patch on the object, seem more problematic. Thus, we might say that we could evaluate the veridicality of color experience against surface reflectivity, and while this is true in the sense that we could establish consistent covariance, there is probably no coherent sense in which one could say that the phenomenal experience of red is a “true” experience of a long-wavelength reflectivity. This, however is because the question is incorrectly posed. It asks about the truth of a comparison between incompatible media. Rather, we must assess the truth value of the mapping between the measures contained in the two information spaces. The mapping from an information space to the medium that realizes it is always veridical by definition, but the mappings that exist between two different information spaces and hence

5 The term “Qualia” has a number of interpretations, some controversial. Here I am using it in the broadest sense simply to mean the phenomenal character of some portion of our low-level sensory experience.

the respective media that realize them have no necessary relationship, and so can be evaluated.

We could, for example, employ standard psychophysical techniques to obtain measures of hue, saturation and luminance from the phenomenal realization of our experience and demonstrate a (perhaps nomic) causal chain mapping this subset of the quale's information space to that of photometric measures of spectral reflectivities contained in a subspace of the information space of the rock. Psychophysical studies rely on our subjective access to phenomenal character — for instance, if we want to find out which retinal stimulations are linked to which particular colors. But the knowledge gained is still empirical and third-personal in nature, unlike our first-personal knowledge about the qualitative and categorical features of colors.(Dorsch 2011) If this mapping proved veridical, we might, perhaps, *then* wish to say that there was a sense in which the phenomenal red experience was a 'true' representation of the rock, however in this sense it would only be idiosyncratically true for this observer; it would also be true in this sense for another observer with an inverted spectrum.

These relationships are diagrammed in Figure 1. As this diagram makes clear, an important aspect of the concept of representation presented here is that phenomenal properties with measures in the information space of our experience do not need to be the same properties as those with measures in the information space of the represented object. It is only necessary that there exist a mapping between the two information spaces that can be demonstrated to result from a process of causal information flow. Thus it presents no problem to speak of *phenomenal red* as representing *physical red* while holding that they share nothing of an ontological nature. There is no conflict in believing that our experience represents the world while believing that the world is not like "what it is like" to have that experience.

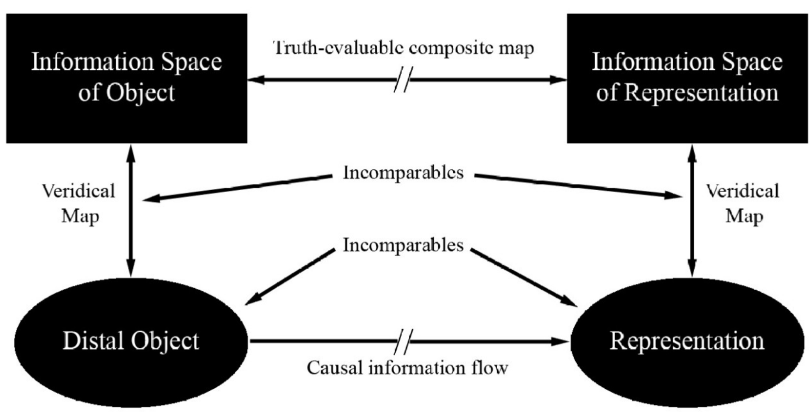

Figure 1: Figure 1: Inherent information is transmitted via the lower pathway. Measures of properties map each vehicle of information to an information space. Maps between information spaces allow evaluation of the representations if the causal pathway exists.

One might yet ask if it is possible for the proposition expressed by such a representation to be false. The answer is not straightforward, and depends on what you are asking. In the simplest sense, such a representation is never fully true. All real channels are noisy and lossy, so that no such representation is ever a perfect one. Thus, 'false' in this sense is not all or none, but comes in degrees. Perhaps a term such as "veridicality" would be preferable to emphasize this. In a different sense, it might be false because the causal chain is incorrect – for example it could have a perfect mapping to the information space of the putatively represented object but actually derive causally from a different object as in the case of the two castings. In this sense it would be the object of a mistaken identity by the perceiver. De Sá Pereira's distinction between 'accuracy' and 'truth' as conditions of satisfaction is useful here.(de Sá Pereira 2016) Yet another interpretation of true and false will emerge below when the hierarchical nature of representation by an internal model of the world is presented. It will be seen that a best-fit hypothesis about a high-level identification of an object, which might be considered a token of a type, may become the current representation based on lower-level representations of varying degrees of "fit'. In such a case a probability measure of the "truth" of the representation may be most appropriate.

The concept of representation under discussion is specifically intended to address both physical and mental representations of low-level, bottom-up, inherent information entering into control-theoretic representations of perceptual processing grounded in neurobiology. As such it is not, necessarily, an appropriate schema for representation of abstract, semantic content *per se*. In following sections we will be discussing the evidence that top-down information of a semantic character enters into and modifies perceptual representation of incoming, bottom-up information. To the extent that top-down semantic content enters into perceptual awareness and interacts with bottom-up information flow to modulate the phenomenal character of perception, this definition of representation is a useful one for dealing with the result.

4 The Information Content of Awareness

Whatever you think the ontology of our awareness of perceptual experience may be, it has phenomenal properties, qualia, which are differentiated from one another, and their differential structures have topological and temporal relations to one another, and those fundamental facts realize the inherent information content of our perceptual awareness; its representational properties. Every fact about awareness considered as a realization of an information space can be stated in terms of spatio-temporal distribution of phenomenal properties (in phenomenal space and time) without reference to how those properties "look or feel" just as it can be stated for a physical object in terms of reflectivities or frequency content without mentioning other facts about those physical objects. The associated measures in this case are those publicly accessed by the standard procedures of psychophysical judgments. To the extent that this information space contains a veridical mapping of the information space of an external event linked to it by causal processes, anything, including our phenomenal perceptual awareness, is by the present definition a representation of the external world.

Of course properties of conscious experience *do* have a look and feel as ‘qualia’. Conscious awareness ‘looks like’ something and is ‘what it is like’ to experience that particular information space, and that phenomenal character is a different, private, sort of fact about awareness. The linkage between the two is that the private phenomenal “look and feel” properties of our conscious awareness instantiate the same information space as is realized by our public judgments about their abstract psychophysical relationships. This is actually just the assumption behind psycho-physics in the first place; measures of phenomenal experience can be related to measures of properties of the stimulus and vice versa. The fact that our interest, unlike that of psycho-physics, may be in the private experience of this realization of the information space does not change this. Chalmers (Chalmers 1997) suggests that since physical properties and phenomenal experiences are realizations of the same information state this “double aspect” makes it tempting, although not necessary, to equate them. He says:

> “Wherever there is a phenomenal state, it realizes an information state, an informational state that is also realized in the cognitive system of the brain. Conversely, for at least some physically realized information spaces, whenever an information state in that space is realized physically, it is also realized phenomenally.” (p. 286)

It certainly seems correct to me that the two are realizations of the same information space, yet I would be cautious about equating them ontologically for that reason. The information state is represented by a vector in information space, and it could well be that it projects onto orthogonal dimensions of phenomenal character and a physical (for example, neural) property.

However, this is only to say that the explanatory gap remains. This “double aspect” nomically describes the flow of inherent information from the external to the mind-dependent, and links the two as representations (in the current sense) of the same mind-independent thing, which is the only issue with which the present argument is concerned. It is not tantamount to elucidating the *mechanism* by which this occurs, which is Chalmers’ hard problem of consciousness. It remains to explain how it comes to be that there exists some lawful mapping in the mathematical sense between the information content of conscious awareness and the information content of the distal objects; we still must either expand the ontology of physics to include private experience or expand the ontology of reality to include mind. Defining awareness as a representation in this sense therefore says nothing about the ontology of awareness, only about its information content.

Thus I am not concerned if you care to think of our awareness of the distal object as a mind-dependent thing that is “like” the distal object of perception by virtue of the information it embodies, or a neural firing pattern that instantiates that information without being “like” it, or the process of becoming aware of it, or the state of being aware of it. Similarly some may prefer to refer to it as an “appearance” or some similar term. All of these concepts and labels, and I believe anything else that anyone might propose to constitute veridical perceptual awareness, have in common the property that they realize an information space causally determined at least in part by an external, mind-independent object and are thus encompassed by the definition of representation discussed here.

The mind-independent object represented in visual awareness is the facing surface of the three-dimensional world, and usually in particular that portion of it that is presently defined as the distal object of interest by image-segmentation and attentional processes. Trivially, the distal object is not only an embodiment of its own information content, but also therefore an exact representation of itself, and its visible surface is a partial representation of itself which participates in the causal chain resulting eventually in the pattern of neural impulses in the observer's brain and so on to awareness itself as the final representation.

An important, immediate consequence of the definition is that since representation involves a *continuous* process of information flow, an infinite number of causally-connected representations must exist between the distal object and our awareness of it. As Venn remarked, we should replace the chain of causation with the “rope of causation.” (Venn 1866). Hence it makes no sense to ask what is the final representation before awareness for the same reason that there is no last real number before an integer. Physically, causal flow of information is a process continuous in space and time. If mind-dependent phenomenal awareness is indeed a representation as defined here, then what can be said about continuity of the path by which the information it represents reaches it? Depending on how you think about mind, there may be different answers:

1. If you are a physical materialist then this process is spatio-temporally continuous from the distal object through various physical realizations connecting it to our sense organs, and thence through the nervous system to whatever neural activity is the substrate of phenomenal experience.

2. If you are a property dualist who holds that mind somehow supervenes on neural activity of the brain, much the same statement holds, for phenomenal experience then depends upon a continuous spatio-temporal process of information flow to those neural events upon which awareness supervenes.

3. If you are a substance dualist who accepts that information is transferred from physical to purely mental processes, and that like Descartes you hold that mind does not have spatial extent, then spatial continuity within mind is vacuously true since all mental processes have the same separation (none), while temporal continuity still seems a requirement for any concept of causal information processes within mind, and is also likely true.

4. If you are a substance dualist who holds that mind has spatial extent, then the situation is unclear since we have exactly no information on how continuous causal processes might propagate in a spatially extended mental substance. Depending on what laws might govern interactions in an extended mental substance, saltatory propagation of effect might be possible between temporally continuous but spatially-separated mental events. Or not. Again, however, it is difficult to imagine causal flow without at least temporal continuity.

Possibly excepting the last case on grounds of ignorance, it appears that there can be no such thing as a final representation prior to awareness since a spatially or temporally continuous process can have no penultimate element before any given event. From this it follows that if awareness is a representation there can be no 'tertium quid' serving as the object of awareness. Thus if, for example, in the causal flow of information from distal object to awareness there were a mind-dependent object such as a 'sense datum' prior to phenomenal awareness there must yet be some further process carrying its inherent information into awareness. Hence awareness itself must be the final representation of the inherent information it presents. Awareness is then revealed as both the observer and the observed, in which the distinction between them vanishes, and along with it the problem of the homunculus and regress. The question of what is performing the act of observing the object of awareness is answered without need of any further iteration. The only final representation in the flow of information between the object and awareness is the inherent information content of awareness itself.[6] [7]

5 The Problem of Representing Scale

Consider a black and white checkerboard. When it is at some distance from the observer, the phenomenal experience will be that of a uniformly gray surface. Similar examples are easily available for color, connectivity, shape, etc. Clearly neither grayness nor 50% reflectivity are inherent properties of the distal object in this case, so how is this acceptable as a mapping of the information space of awareness to that of the object? It would seem that the issue here is scale, but can we legitimately include a scale dimension in the information space of the object? Worse, the scale of presentation interacts with the angle of regard since the acuity of the retina is not constant across its surface. It seems we must have dimensions of effective scale incorporating both radial distance from the observer and angle of regard, which are hardly inherent properties of the object. Rather they seem to be relations between the object and the observer[8]. This is an entirely different case from something like a "bent stick illusion" in which there is a perfectly good, continuous, causal mapping from the information space of the stick before the light carrying the information passes through the water to that carried by the light following the event. That mapping then composes with subsequent maps in the causal chain to form the overall map from object to awareness.

Instead, the distortions of scale are problems of information loss, not different in principle from those mentioned in regard to a lack of cones responding to ultraviolet or infra-red. Thus, at the appropriate scale, several black and white squares fall on the same retinal receptive field producing an average firing rate halfway between either; a saw-toothed edge in a similar manner evokes the same response as a straight line and shape information is lost; a void is averaged out and connectivity information is lost.

---

6 On this account, those who hold that perception is a two-place relation based on accepting the causal chain of events but denying that we are aware of the vehicle of representation would seem to have the problem of explaining how we are not aware of our awareness.
7 It is interesting that Aristotle invokes something very like continuous information flow (*De Anima* III 12), and also (for other reasons) holds that the senses are self-aware (*ibid*, III 2.)
8 Although some advance positions that do attach spatial relation to the observer to the representation of qualia. (Tye 2017)

Thus the phenomenal experience is no more *incorrect* here than in its failure to appreciate the ultraviolet, and the map is at least in this sense partially veridical. Also, there are inherent properties of the object, instantiated in its information space, that partially determine what the information content of awareness will be at any effective scale. In particular, the actual relative sizes and topological and geometric arrangements of the object's properties will determine how its representations transform under scale changes. But this still leaves unresolved the issue of how and where the effective scale changes, which are relations between the observer and the object, are themselves represented.

Up to this point, for convenience in discussing other issues, I have spoken of objects of our awareness as if their mereological natures were already part of awareness. They are not. At the early stages of perception the visual (or other sensory) field is presented as a whole. The necessary properties on which scene segmentation can be based such as color, shape, and so on are present, but scene segmentation based on them is a later stage of perceptual processing. Mathematically there exists a single, high-dimensional, information space containing all the measurement vectors realized by the entire scene. The problem of perceptual scene segmentation – that is, its decomposition into the sort of objects we have been considering - is precisely that of finding the optimal set of information sub-spaces into which to decompose the information space of the scene.[9]

The sensory field for which measures are encoded by this total information space contains, in addition to everything else in the scene, our phenomenal awareness of our own bodies. As Smythies has pointed out, what we are aware of is not our bodies but a perceptual representation of our bodies in the same form as any other component of the phenomenal scene; what he refers to as the "body image."(J. R. Smythies 1994) The information space of our awareness of the scene must then contain this body-image information. Other scene information available to the initial stages of perceptual processing includes the multitude of cues to depth from which the observer-centered depth map of the world is constructed, and hence the spatial relations of locations in the scene to one another and to the observer. Thus the correct way to think of the information space of a segmented object in our perception (as opposed to the information space of a physical object in the world) is not just that subspace of the information space of the scene containing the measures of the object's structurally-determined inherent properties, but that subspace which also includes the time-varying vectors structurally inherent in the scene that encode the object's distance, orientation and angle of regard with respect to the body-image of the observer.

Since the observer and the distal object are both structural parts of the scene, these additional vectors encode measures of structurally-inherent properties of the scene, and there is no basis for not including them in the information subspace describing any segmented subset of the scene, including any object of awareness. Any special status for an object-observer relation in this context is seen to be only an apparent one, and not different in nature from relations between internal areas of the object. Thus, in the early stages of perceptual representation, potential mereological subsets of the sensory

9 Where 'optimal' is defined relative to the needs and abilities of the organism.

field are treated as segmented regions and are not represented as semantic types in a relation with the observer, but as particulars individuated by their location in an observer-centric space.

In the information space of the sensory field of awareness, the phenomenal property of luminance of the checkerboard will change as we approach it, and this change will co-vary in particular ways with changes in the information-space vectors defined on measures of its phenomenal observer-centric location. Thus the gray phenomenal property occurs in a different part of the information space than the checkered phenomenal property. The two regions of the information space are separated along the dimensions of vectors encoding the effective scale property. At the other end of the causal chain of information flow, the information space of the distal object, as already discussed, contains information about the physical structure of the object which determines how its measures will transform under changes of effective scale. Thus the veridicality of the mapping may be established for any combination of phenomenal property and phenomenal location. A quick way to think about it is that finding the object's current point in the information space of awareness maps a retinal receptive field onto a region of the physical object and takes an average (and of course this is always part of the mapping process at all scales.) This scale-space approach is widely used in computer-based image-understanding systems. See (ter Haar Romeny et al. 1991) for an overview of the general principles.

When an object moves away from you it does not appear to shrink, yet there is some sense in which we have a perception that something is getting smaller, or at least changing. If this something is our phenomenal experience of the scale property of the relation then it might be thought of as reflecting a mode of presentation in the Fregean sense. (Zalta 2001)

There is one loose end here, which is how information about things such as the size of a receptive field or the acuity effects of astigmatism could come to be associated with a point in the information space of awareness, since these, unlike the body image, are not (directly) part of our phenomenal experience, and have variability from one observer to another in both. This point is best considered among the issues discussed in the following section.

6 The Problem of Representing Space and Time

Space and time enter importantly into our phenomenal awareness of the world; they provide the framework relative to which our other phenomenal experiences occur. Yet it seems curious that a part of our primitive and immediate phenomenal awareness could be composed of something that cannot be sensed, at least in any direct way. Still, most people report that with their eyes closed they are not aware of nothing, but rather of a blank, black field that seems to have extent. The same is not true of patients with cortical blindness after loss of area 17, who report awareness of nothing (despite often exhibiting "blind sight" phenomena.) Such observations suggest that there might be some apparatus for immediate phenomenal awareness of spatial extent. It is often assumed that a representation of the geometry and metric of the world we experience, and which we need in order to act in it, can be obtained from retino-topic mapping of the visual cortex together with cues to depth such as stereopsis. It is more complicated than this.

There are two aspects to the phenomenal awareness of space which were distinguished by Gibson as the visual field and the visual world. (Gibson 1950; Goldstein 1981). The visual field is the “raw experience” of the scene of which a painter must be aware when he paints a railroad track as two converging lines moving up on his canvass. The canvass is then a good metaphor for the visual field and clearly both have a topology, geometry, and metric. The visual world must also have a geometry and metric in order for us to get around accurately in the physical world, but that would appear not to be the geometry and metric of the raw visual field. The geometry of the image on the retina is given to us as projective by the physics of the eye, but that of both the phenomenal visual field and the phenomenal visual world has been a matter of long debate, and has included Euclidean, spherical, hyperbolic, and variable. See (R. E. French 1987) for a review of this material. It is only my aim here to consider how this information comes to be instantiated in the information-space of awareness, whatever may be the particular answer to the question of the precise nature of its metric.

The first question is, “Given that the image on the retina has a topology, geometry, and metric, independent of the scene contents, how does that information come to be encoded in the resulting neural activity?” It is well-known that the retina has a hierarchical organization of receptive fields such that excitation within a receptive field contributes to the activity of particular neurons. It is often assumed that the basis of representation of the spatial relationship among these receptive fields is the retinotopic mapping of these fields onto cortical area V1, but this is not so.[10] Only an external observer and not the brain itself has any knowledge of this mapping. Similarly the brain has no clock with which to learn the temporal relations of activity in different neural pathways. How then do the spatial and temporal relationships of points of stimulation on the retina become encoded in neural function?

In a paper concerning the information in the optic nerve Koenderink (J. Koenderink and van Doorn 1982) discusses the ability of a system without a clock to reconstruct the temporal order of events from information in the optic nerve simply on temporal overlap of “signals.” In a later series of papers (Jan J Koenderink 1984; J. Koenderink 1984; J. J. Koenderink and van Doorn 1987; J. J. Koenderink 1987; J. J. Toet A. ;. Blom, J. ;. Koenderink 1987; J. Toet A. ;. Blom, J. ;. Koenderink 1987; 1988) he elaborates a similar argument for the encoding of space in the functional activity of the visual system, independent of any retinotopic mapping, and capable of discovering and encoding such things as the geometry and dimensionality of space from covariance of activity in visual receptive fields (which stand as analogous to the “signals” of the earlier paper.) The approach is easily extended to a space-time.

10 In fact, this sort of mapping occurs in many places in the nervous system where spatial information is not involved. It is likely just a matter of conservation of resources in arranging interconnections of neurons that have a lot of traffic with each other.

I will briefly summarize his argument:

1. The geometric structure of the visual perceptual world cannot be obtained from retinotopic mapping on the neural apparatus of the brain since the neural processes involved have no knowledge of it and it is only available to an external observer. The responses of neurons to stimulation are all the information available to the system. These responses would not be changed if the neurons were spatially relocated but still retained their functional connections. Thus, the geometric order of the visual field must be encoded in a functional order of neural activity.

2. An adequate mathematical basis for encoding the geometry of the visual field in neural activity exists in the hierarchical structure of overlapping receptive fields and the possibility of computing covariance relations among neural elements. Thus, correlation between two neural elements on stimulation of the retina may be taken as a measure of overlap between their receptive fields. From topology we know that such sets of covers of sets can be used to develop geometric properties of spaces composed of the elements of such sets (*cf.* for example, Lebesgue covering dimension, Brouwer dimension.)

3. Using idealized but otherwise accurate examples of neural connectivity in the visual system, Koenderink demonstrates the mathematical possibility of building up a pre-geometry (an ordered lattice structure) that correctly characterizes the most important facts about the geometry of the space of the visual image, including its dimension. Since this geometry is encoded in neural activity it is then accessible to perceptual processes.

This description fails to rise to a full geometric description in that it does not provide a metric. This is a feature, not a bug – a point to which I shall return below. The real visual system is less regular than the idealized models Koenderink uses for mathematical exposition, but its relevant features are functionally the same. What he proposes or something very similar to it is clearly necessary as a model of how the geometry of the sensory world projected on the geometric and hierarchical structure of the receptive fields could be computed and encoded in the nervous system and hence to enter phenomenal awareness. Importantly, it also demonstrates what *cannot* be computed from it.

Let me now bring time into this picture. Since a neuron is basically a leaky integrator that sums activity over time until it fires and then resets itself, it can be regarded as having a receptive field that extends in time as well as space. That is, just as input anywhere within its spatial receptive field will contribute towards firing, so will any activity occurring within the integration interval. Thus in the simplest case a neuron can be considered to possess a series of temporal receptive fields sequentially extended in time, represented by successive integration periods separated by firing events. Its successive integration periods do not overlap and hence its own temporal receptive fields cannot overlap themselves in time. The temporal receptive fields of other neurons receiving input from the same and other spatial receptive fields can of course temporally overlap them, similarly to the case with spatial receptive fields. Thus, the visual receptive field structure consists of sets of 3D (that is, 2 space + time) covers of elements of a visual space extended in time. Therefore a basis exists for encoding time as another dimension within the context of constructing a spatialised perceptual world in a manner analogous to the treatment of the two-dimensional visual field demonstrated by Koenderink.

The availability of an easily-computable measure of spatial and temporal organization of the retinal receptive field structure that is encoded in neural function explains how the system can represent a statement such as “This blue quale is adjacent to that blue quale” without recourse to knowledge of cortical physical structure that is unavailable to it. This kind of information is so foundational in our experience of the visual field that it takes some thought even to realize that it is information that must be obtained and represented. The fact that the covariance must be obtained by acting in the world to generate variance on the retina is of particular interest since it entails acquisition of an idiosyncratic map of adjacency and acuity in the retina of an individual during maturation and aging, a requirement since individual variation would preclude acquiring it through evolution. This gives us part of the answer to the problem introduced above concerning representation of the angle of regard with respect to a segmented area of the sensory field, which is necessary for correctly mapping effects of acuity and astigmatism.

But that it is still a good ways from a full spatial construction of the 4D perceptual world we know, especially without a metric. How then do we obtain a metric, and how do we understand that metric in relation to the individual proportions of our limbs so that we may accurately act in space? Berkeley (Berkeley 1709) was an early proponent of the idea that the relation of the two-dimensional image on the retina to distance is an association between sight and touch, and that this must be learned. Gibson (Gibson 1950) enunciated the motor theory of perception according to which the co-variance of visual and kinesthetic sensations upon actions of the body discover the spatial relations between the two. This implies that the perceptual understanding of the metric of visual space must be represented in terms of the metric of proprioceptive and vestibular space; that is, represented in terms of joint angles, limb-lengths, and body orientation.

Thus when the infant encounters a tactile sensation together with a certain proprioceptive and kinesthetic experience, and simultaneously with these the sensation of a particular visual stimulus within the geometry of the image, it is common to suppose that hand-eye coordination is learned. This can be restated in part as supposing that covariance matrices of the sort Koenderink discusses are being acquired and that what is being learned is the X,Y,Z,T pre-geometry of the phenomenal space of the full perceptual world, so-encoded, and including the pre-geometry of the body-image extended into that phenomenal space. However, a meter stick and a system of units grounded in physical space is still lacking, and they must have a representation within a fully-metrized phenomenal space. They also must be represented as something other than sensory phenomena, since the metric of sensory perception is itself to be given by reference to this meter stick. To be more exact, we need a set of metric values that can be placed, subject to certain conditions, into correspondence with pairs of points in the topology of phenomenal space, and this set of values must have a representation within the phenomenal rather than physical world. Nonetheless, to be useful in the control of behavior, this should be subject to the constraint that distances so-assigned to point-pairs of phenomenal space shall stand in the same relation to each other as distances between the corresponding points of physical space under some continuous map. How can this be achieved?

Whenever we make a motion, it is the result of an outflow of hierarchically-organized motion-planning activity in the neurons of the motor system, encompassing many structures in the nervous system.[11] Thus a high-dimensional vector of "intended magnitude and direction" exists for each motion, that is realized by this efferent neural activity for the effort planned. The resulting physical action in space will subsequently produce some result in the afferent, sensory neural apparatus, and ample opportunity exists for the computation of covariance of efferent activity with the resulting activity in the proprioceptive, kinesthetic, tactile and visual systems. The "truth" of the proposition that the motion was correctly planned is established when tactile or visual confirmation is obtained that the target point was reached. This confirmation enables the nervous system to learn the correct correlation between efferent motor system correlates of physical motion and the results off these motions in sensory space. Clearly this correlation will differ for every individual, and for a given individual over their lifespan, as it depends on physiological variables such as limb-length.

As a result, vector differences between successive states of the efferent representation at successive states of an action provide a set of differential values represented in functional neural activity that can act as a metric for the point-pairs of the pre-geometry of physical space involved in each state of the action, including the body-image itself, which has been derived as above from the simultaneous order of activities of the afferent sensory system. Thus any two events in sensory space-time such as "my fingertip is touching the left corner of the page" and "my fingertip is touching the right corner of the page" have representations derived from visual, tactile, kinesthetic and proprioceptive activity, and there will exist a vector difference between the two states of the efferent activity corresponding to each that may be assigned to that phenomenal point pair as a metric. Since the relation between the point pair and the metric value will have been mediated through the actual geometry and physics of the world in which the action occurred, a neural representation of their geometry in the physical world exists, and this provides an observer-dependent metric for a full geometry of the external world. (This is simplistic; the body has mechanical redundancy and there are many ways to move one's arm that put the fingertip at the same point in space and there are different conditions of resistance to motion, etc., but none of it is in principle incapable of computation over a series of actions.) The phenomenal and neural representations of physical space need not have the same metric, nor need it be the actual metric of physical space; they can be rubber-sheet distortions of one another and we

11 The hierarchical elaboration of motion in the nervous system was first understood in the 1870s by John Hughlings Jackson; a non-technical modern overview of the operation of these structures can be downloaded here: https://www.ndsu.edu/faculty/pavek/Psych486_686/chapterpdfs1stedKolb/kolb_10.pdf.

would not know it, but their correlation matrices induce a map between them expressing the necessary functional relationship to ensure accurate action in the world. This is not the metric we use to describe Euclid but it is the metric we use to catch a ball and to metrize phenomenal space.

In this view then, the representation of phenomenal space is comprised of a topology supporting an ordered lattice provided by sensations, and a metric defined on it by interacting with physical space to generate those sensations. In essence, our limbs become our measuring rods. What could be more natural for a brain that has the function of controlling the action of the body in space?

7 Robotic Interlude

Control Theory offers many concepts of significance for understanding perceptual processing in the brain. Some of these also can be put into analogy with the terminology of philosophical discourse. The idea of a hierarchical internal model of the world, in which many levels of description, both intentional and non-intentional, are simultaneously available, is of particular importance.

While information theory is well-known to most, control theory has less of a following. It is, however, essential in neurobiology to understanding the functions of the nervous system and by extension, their implications for philosophical issues. It is, approximately, the mathematical treatment of the problem of applying information to the control of goal-directed systems. Trivial examples are thermostats and automobile suspensions. It deals in concepts such as feedback, feedforward and stability under perturbation. The brain *is* a control system; that's what it does - what it evolved to do. Higher cognitive faculties in humans are a minor, late addition, and even those are employed the majority of the time to further refine biological, goal-directed behavior.[12] Describing the operation of the nervous system in control-theoretic terms is a problem of hideous complexity. We can understand the general outlines of a lot of it, but understand only bits and pieces of it in detail.

In particular, inside the nervous system, information flow is dominated by feedback, and information flow from post-perceptual neural processes flowing back into perceptual processes forms the most important determinant of the state of phenomenal experience. Here the sequential flow of information from one representation of inherent information to the next, while yet remaining a causal chain, becomes complex to the point that an "intervening representation" becomes difficult even to define. The naive view of neurobiological vision, that it consists of successive stages of analyses of the retinal image which finally become the data which enter consciousness, is simply not tenable.

Therefore, before tackling the evidence for mechanisms of neurobiological representation in perception, I would like to review the operation of a simple, artificial system that will clearly display some general concepts and permit comparison with some

12 The ability to divert these resources to abstract problems may well be a design flaw that survival of the fittest will ultimately correct.

philosophical terminology. This is the design of a vision system for a robot that was constructed in the 1980s with the objective of mimicking some important aspects of biological image processing. (Kent and Albus 1984) It in no way captures the complexity of neural processing, which does not come apart into neat boxes and simple features, but it will help us to get a start on it.[13]

Please examine Figure 2, which consists of three columns. The column on the left with boxes “SP” is the sensory processing column. The center column with boxes “WM” is the system’s internal model of the external world. The right column with boxes “AP” is the action planning system. The arrows indicate directions of information flow in the system. This is a hierarchical control system, and each horizontal level is a control system in its own right for dealing with information at some level of description. The arrows for level three have been labeled, and each of the other levels would have similar labels. At each level, the WM box contains the system’s current belief about the state of the world at that level of description. It receives feedforward input from the AP box at that level concerning the current state of action planning, and employs that to predict what the sensory state should become as the result of that action if its current assumptions or “beliefs” are correct. That prediction is passed to the SP box, which compares the prediction with accepted data arriving from below. If the prediction and the data match within acceptable limits, a statistical best-fit is made and the prediction, adjusted to the best fit, is accepted and passed up to the next level of SP as well as back to the model, which updates its belief about the current state of the world. If the mismatch between observation and prediction exceeds some level, the model must strive to find a different belief about the state of the world that matches it acceptably, and a “perceptual reorganization” may occur. Asynchronously with this process, the boxes of the AP column are receiving commanded states from above at their level of description and resolving them into finer-grained commanded states for the level below, along the ideas noted in footnote 11. To do this it needs to know the state of the external world, and it obtains this from the “best guess” embodied in the beliefs (assumptions) of the world model at its level.

What is omitted for clarity is that at each level, multiple, parallel lower levels receiving input from a single higher level are instantiated to deal with subdivision of the problem in narrower scope but greater detail. There are many advantages to the use of world models in control systems, and most control systems of any complexity require them. (Francis and Wonham 1976) The complexity of hierarchical control systems is rarely required in small-scale engineering applications, but they are of great theoretical interest in neurobiology, and are employed in robotics and artificial intelligence as well as plant-wide control of manufacturing.(“Hierarchical Control System - an Overview | ScienceDirect Topics” n.d.; “Hierarchical Control System” 2020; Raisch and Moor 2005)

A clearer idea of the interaction between data inflow and current model may be had by examining figure 3. This illustrates how detected edge-points maintain the internal model of a rigid rectangle as the state of the system evolves. At each of the two levels

13 Viewed as an instantiation of a functional theory of perception, if such a concept is applicable to machines, it at least has the advantage of having been built and demonstrated to actually work.

shown, there is an input from WM representing the current "belief", and a collection of lower-level descriptions to be matched to it. There will always be noise and error so that the fit is never perfect. Kalman filtering is used to weight the model vs. the data in making a best estimate so that the model will adjust smoothly unless there is a difference that exceeds the expected effects of noise. The prediction may be based on *a priori* knowledge of properties, such as the fact the shape is rigid, in addition to simply working with last accepted locations.[14]

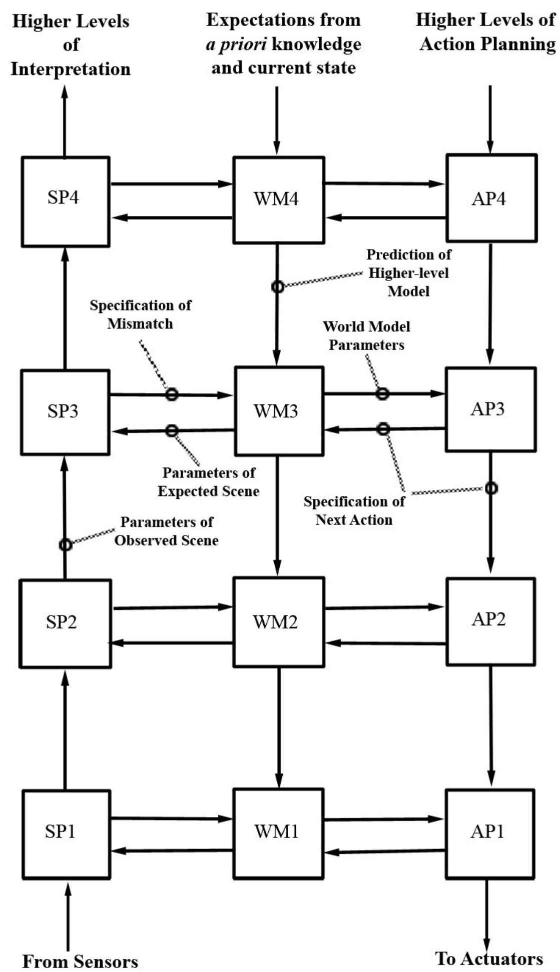

Figure 2: Simplified scheme of functional information flow, with path descriptions shown for one hierarchical level. SP: Levels of sensory processing. WM: Levels of world model description. AP: Levels of action planning.

Can we build any bridges between this kind of system operation and philosophical terminology? I think to some degree we can, although not a perfect fit. First, at the lowest levels of the system we are getting some kind of smoothed and filtered representation of actual physical events impinging on the sensors. This can be compared with "qualia realism," the view that claims that conscious character is an intrinsic nonrepresentational property of visual experience. On the other hand, from the highest levels there is descending information flow in the form of a representation that compares favorably with an intentional "type" derived from narrow content. Of interest then is what happens in the middle. At each level the type suggested from the higher level is particularized as a token by interaction with ascending analyses from lower levels, so that we get, for example, a located and oriented token of the rigid-rectangle type. This in turn becomes a token of a type for subsequent comparison and prediction. Thus from top to bottom the flow is to representation of increasingly particularized broad content, while the bottom-up flow is from particular properties to increasingly abstracted intentional representations.

In sections 9 and 10, I will argue that our phenomenal awareness either is, or is of, a controlled world model maintained by the brain in a manner similar in principle to

14 This is not to be taken as indicating that the brain actually engages in this simplistic, sequential "grandmother cell" approach to image understanding (although it worked for the robot), but only to illustrate the concepts of servoing model-derived expected input to observations to correct the model.

this simple example. Further, the evidence suggests that our awareness seems to be of all of the levels simultaneously, as if they were a transparent stack, so that when there is no stable higher-level organization available, we will still see the "qualia realism" level. Thus the answer to a number of philosophical disputes of the form "Is the nature of perception X or Y?" may be "yes".

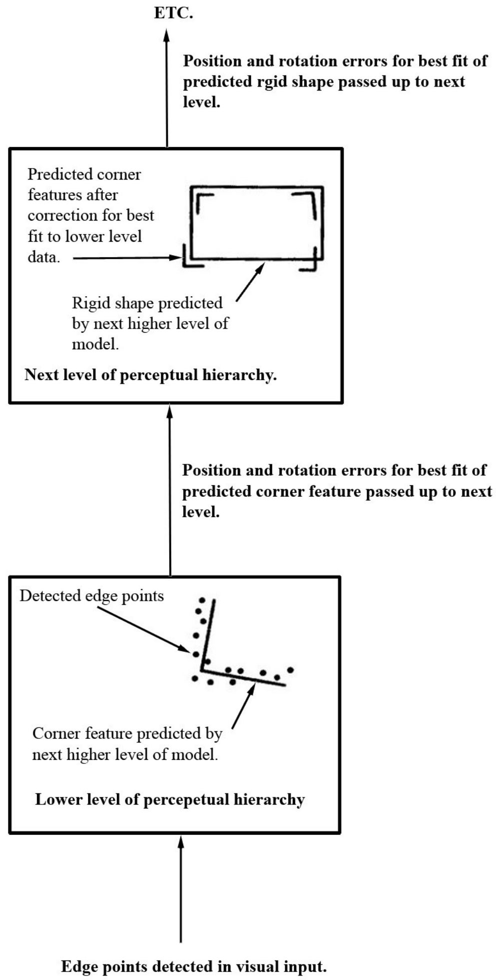

Figure 3: Conceptual example of error-fitting the prediction from the model to input from a lower level at two hierarchical levels.

8 Phenomenal Character and Representational Properties as Awareness of a Controlled Internal Model

The current majority view of neurobiology is that our phenomenal awareness is best thought of as either being, or being the mental correlate of, an integrated, neural, predictive model of our sensory-motor environment that is constrained at the bottom level by non-intentional input from sensors, but subject to reorganization and influence by a host of intentional influences at every hierarchical level of description.

This picture has frequently been referred to as our perceptual experience being a "controlled hallucination" (Nir and Tononi 2010; Parkkonen et al. 2008). By this is meant that the brain continually constructs predictive hypotheses in the form of models of what we will experience next, and these perceptual models are continually corrected to conform to the actual, current sensory input. It is then this predictive "hallucination" which we experience rather than the raw sensory data. None of this is particularly new, and the basic idea goes back to Helmholtz (von Helmholtz 1896). It has been very precisely developed in mathematical control theory as described in the preceding section, and is the mechanism at the heart of all complex control systems. Indeed, it can be proven that such an approach is necessary for optimal control, and the brain's evolutionary history is as an optimal control system for the organism. In these terms, stripped of a bit of attention-grabbing nomenclature, the "controlled hallucination" is perceptual awareness of a neural model of the external environment that is developed by the brain and "servoed" to the sensory input by various statistical procedures in which the error between the estimated sensory input from the model and the actual sensory input is used to continuously correct the model. Control of action at every moment is then guided by the current state of the model rather than by raw sensory input. As such it qualifies as a representation of inherent information under the present definition, inasmuch as it is causally linked to the environment it represents through the senses (in addition to many other inputs), and it realizes an information space that maps veridically to that of the external scene. (In fact, the whole point of this vastly complex apparatus is to maintain the best possible fit between the model and the scene in the face of noisy and incomplete data.)

The evidence supports that it is not our sensory input of which we are perceptually aware, but rather this predictive hypothesis, this model of *impending* sensory state, generated like dreams and hallucinations. If this is true then little of the raw sensory input ever reaches consciousness in any direct manner. Instead we are indeed experiencing something like a predictive hallucination about the world which is continually updated via error signals. In particular it is asserted that the appearance of which we are aware has its origin not in the immediate distal world but in what amounts to our imagination, informed by beliefs, knowledge, memory, intended actions, and the current state of the model itself, however much that imagination is constrained by current sensory input.

From the considerations of section 4, it would then be the case that:

1. This internal model is, or is the neural correlate of, the final, phenomenal, representation of the world.
2. This internal model is, or is the neural correlate of, our phenomenal awareness.

I would like to lay out the evidence for this position in some detail.

9 The Internal Neural Model of the World

That the brain maintains a multi-modal, internal model of the external world, synthesizing information concerning all the senses, including proprioception and kinesthesis (Clark 2013; Land 2014; Lee 2016; H. Barlow 1987; H. B. Barlow 1994; Ito 2008; Kawato 1999; Picton and Stuss 1994; Friston 2012; Petro and Muckli 2016; Jeannerod 1994; Lange and Haefner 2016; Jerath, Crawford, and Barnes 2015), and producing a unified model of the body within an organism-centered model of the world is widely accepted. The model is continuously adjusted on the basis of intended motion in order to predict impending sensory states, and by using knowledge about the world continuously to track and update the evolving state of affairs. The resulting predictive model is then used, rather than raw sensory input, as the guide to control of action (Brown, Friston, and Bestmann 2011; Cerminara, Apps, and Marple-Horvat 2009; Hikosaka et al. 2002; Imamizu et al. 2000; Jeannerod 1994; Kawato 1999; Kwon and Knill 2013; Lalazar and Vaadia 2008; Seidler, Noll, and Thiers 2004; Shadmehr and Holcomb 1997; Sülzenbrück and Heuer 2011; Synofzik, Lindner, and Thier 2008; Thoroughman and Shadmehr 2000; Wang, Arteaga, and He 2013; Daniel M Wolpert, Ghahramani, and Flanagan 2001; D. M. Wolpert and Kawato 1998; Daniel M Wolpert, Miall, and

Kawato 1998). The sensory and connative functions are thus highly integrated. We don't really have a sensory system so much as a sensory-motor system.

From a control-theory standpoint, this servoed model approach is a necessary evolutionary response to the fact that sensory input is partial, incomplete, sporadic, ambiguous, unstable and very noisy. The method is common to all advanced control systems, both artificial and biological. Our later, higher intellectual functions have been incorporated into an existing underlying machinery of sensory-motor behavioral control in order to improve its performance, particularly with respect to longer time-frames of action. Control theory describes and constrains the behavior of all feedback-dominated systems, and the brain is the stellar example of a feedback-dominated system, rather than a linear system that sequentially processes information into ever more refined forms (Layton, Mingolla, and Yazdanbakhsh 2014; Lee et al. 1998; Scholte et al. 2008; Seidler, Noll, and Thiers 2004; Somers et al. 1999; Williams et al. 2008; Wokke et al. 2013; Kerkoerle et al. 2014). In such systems the functional characteristics of the system are controlled by the properties of the feedback loops. Control theory is about obtaining optimal performance, and it can be proven that internal models are necessary for optimal control (Francis and Wonham 1976). The nervous system uses Baysian statistics, with the current state of the model as an estimator for the next state of the world, to deal with noisy and missing input information (Fiser et al. 2010; Knill and Pouget 2004; Lee 2002; Lee and Mumford 2003; Rao 1999). It draws on multiple sources of knowledge to implement this, ranging from phenomenal beliefs[15] about the current state of the world already embodied in the immediately preceding state of the model itself, to knowledge of the properties of objects, to abstractions such as models of naive Newtonian mechanics (Cheron et al. 2014; McIntyre et al. 2001; Zago and Lacquaniti 2005). This knowledge is not necessarily verbal, and presumably is largely shared by non-verbal animals. While there is evidence that language-like representations do penetrate to affect lower levels of the sensory systems, the evidence is disputed and not yet clear. Thus I am speaking here primarily about the kind of epistemic knowledge that informs our actions; that is, naive physical understanding of objects in our perceptual world and of how objects in the world, including our bodies interact, and how those interactions are contingent on the form of our behaviors.

The process by which this proceeds is generative, Baysian, statistical inference (Fleming 2014; Friston 2005; Hinton 2007; Lange and Haefner 2016), in which bottom-up connections are driving, and top-down connections are both driving and modulatory. The internal model of the world is hierarchically adapted to form the statistically best fit to the raw data at multiple levels of description. It is much more complex than this suggests, and the brain, being evolved rather than designed, is not constrained to do similar jobs in identical ways at every point, but this is the basic plan. The fundamental process is predictive coding, and is thought to be the essential means by which the sensory-motor system acts (Edwards et al. 2017; Fiser et al. 2010; Friston 2010; Hohwy, Roepstorff, and Friston 2008; Huang and Rao 2011; Jehee and Ballard 2009; Knill and Pouget 2004; Kok and de Lange 2015; Lee and Mumford 2003; O'Callaghan

15 Which may be conceptual beliefs or non-conceptual beliefs about the world available to any organism, depending on the level of representation.

et al. 2016; Petro and Muckli 2016; Pickering and Clark 2014; Rao 1999; Young 2000). The process continuously keeps the model in synchrony with the distal world, and the predictions of the model are continuously adjusted to account for current behavioral acts[16]. If at any level of the hierarchy the error-signal strays too far from expectation, a new hypothesis may be formed and a perceptual reorganization occurs (Wang, Arteaga, and He 2013).

None of this requires any necessary relation to consciousness, and it would be just as necessary, and work just as well, for a zombie. It was evolved long before man walked the earth, and one of the structures principally involved in its operation, the cerebellum, is phylogenetically very old. Although cerebellar function evolved early for the control of action based on sensation, its associations with all parts of the cerebral cortex have continued to develop to include those involved with cognition and emotion (Schmahmann 2010), including individual-specific functional networks, and it plays a role in adaptive plasticity. In the human, the cerebellum contains the majority of the brain's neurons, and it has evolved most rapidly in the great apes and humans (Barton and Venditti 2014) underscoring its important role for higher functions in the operation and utilization of the model. It is increasingly understood to play an important role in sensory-motor integration and feedback/feedforward control of the internal world model at all levels of perceptual processing. (Daniel M Wolpert, Miall, and Kawato 1998; Synofzik, Lindner, and Thier 2008; Cerminara, Apps, and Marple-Horvat 2009; Ito 2008; Cullen and Brooks 2015; Imamizu et al. 2000; “The Underestimated Cerebellum Gains New Respect From Brain Scientists” n.d.; Hogan 2004; Marek et al. 2018; “Cerebellum” 2020)

10 The Internal Model is Both the Content *and* the Medium of Our Perceptual Awareness

Empirical evidence supports the claim that the information state realized in this sensory-motor model of the world, and not that of the immediate sensory input, is the information content of our conscious, perceptual awareness (Cichy, Heinzle, and Haynes 2012; Albers et al. 2013; Chong, Familiar, and Shim 2016; Naselaris et al. 2015; Knauff et al. 2000; Harrison and Tong 2009; O’Herron and von der Heydt 2009; Muckli, n.d.; Stokes et al. 2009). Given the very close, albeit nomic, causal connection between brain states and mental states, this is hardly surprising. This is not to say that there is a model of the environment somewhere in the brain that “looks like” the external world. This is not a sense datum. The model is both hierarchically organized and distributed. Different aspects of it, and at different levels of description, are the information content of the multitude of sensory processing areas – over 30 in the visual system alone. Moreover there appears to be no one central location where all the parts are drawn together; rather, the unity of perceptual awareness is a hegemony of richly-

16 To gain a good, intuitive feel for how this process appears in real-time operation, I recommend you take the time to work through the visual experiment in the later section “Cognitive Penetration” below, which provides a convincing demonstration of the model dynamically adjusting the perceived geometry of objects under body motion when conscious assumptions about orientation are put in conflict with knowledge about rigid objects.

interacting perceptual domains, each dealing with some aspect of the whole (S. Zeki and Bartels 1999).

It appears that neural activity embodying this distributed, hierarchical model not only entails perception but is entailed by perception. In the forward, or bottom-up direction, the dependence of perception on this representation is clear. Stimulation of these areas is well-known to produce perceptual experiences. Lesions in specific portions of the sensory systems eliminate perception of corresponding specific perceptual attributes without eliminating others, as in the various clinical agnosias (reviewed in (Semir Zeki 1993)) Further, the ability to perceive elements structurally and functionally antecedent to the agnosic deficit (such as accurate perception of individual component features while lacking face recognition) is retained, demonstrating that no simple, sequential analysis leading to some later perception is being interrupted.

In the opposite, top-down direction, activity in the neural apparatus subserving this representation is modulated, and can be initiated by, perceptual events that do not arise directly from sensory input. The experimental evidence supports the conclusion that all the things we might consider non-sensory sources of, or influences upon, perception, while originating in different underlying neural events, are ultimately expressed in our perceptual awareness via the same neural pathways that mediate the entry of information from the sense organs into awareness (Muckli, n.d.; Cichy, Heinzle, and Haynes 2012; Stokes et al. 2009; Albers et al. 2013; Chong, Familiar, and Shim 2016; Naselaris et al. 2015; Knauff et al. 2000; Harrison and Tong 2009; O'Herron and von der Heydt 2009). For example, induced changes to perception such as volitional figure-ground reversals or ambiguous figure changes are repeatably associated with activity in these same neural circuits in the absence of changes to sensory input. Subjects asked to make volitional shifts in figure-ground reversal images show characteristic activity in the frontal cortex which is followed by activity in the visual cortex coincident with report of the shift, even though the stimulus reaching the visual cortex from the sensory inputs is unchanged (Wang, Arteaga, and He 2013). Thus, it has been found that non-sensory events which alter perception, such as volitional acts, ambiguity resolution, and illusions (Kornmeier and Bach 2012; Wokke et al. 2013; Kornmeier and Bach 2004; von der Heydt 2015; Seghier et al. 2000; Layton, Mingolla, and Yazdanbakhsh 2014) alter neural activity in these sensory areas, and so do internally-generated kinds of perceptual imagery such as hallucinations (Howard et al. 1997; ffytche et al. 1998; David et al. 1996), dreams (Dresler et al. 2011; Horikawa et al. 2013; Igawa et al. 2001; Nir and Tononi 2010; Revonuso 2006), and imagination (Kosslyn et al. 1993; Kosslyn and Thompson 2003; Aleman et al. 2001). Finally, "filling in" phenomena, which are a form of hallucination involving the perception of synthesized features where none exist or are different, are a routine part of ordinary visual processing (De Weerd 2006; De Weerd and et al 1998; Komatsu 2006; Luiz Pessoa and De Weerd 2003; L. Pessoa, Thompson, and Noë 1998; Rees and weil 2009; Spillmann and Weerd 2003). In general the visual system at every level passes on only boundary information about homogeneous regions of color, luminance, texture and so on that are implicit in their boundaries. Our perceptual awareness of these regions is constructed by 'filling in.'

The case is made therefore that neural activity in the sensory apparatus underlies perceived imagery whatever its sensory or non-sensory source, and is either identical with, or the neural correlate of, our perceptual awareness. As such, the perceptuo-motor model is driven by both bottom-up sensory inputs and top-down inputs from non-sensory sources. With this in hand we can ask what are the sources of the information that reaches conscious awareness. The crucial point is that the model, and hence perceptual awareness, is not just a representation of immediate sensory input from the external world but rather a reconstruction of the nervous system's best estimate of the state of the external world, derived from many sources such as object knowledge, past experience, expectation, perceptual learning, servo-estimation, filling-in procedures from boundary conditions, ambiguity resolution and the like (Gilbert and Li 2013; Kandel, et al, 2014; Cavanagh 2011; Naselaris et al. 2015; Williams et al. 2008; Harrison and Tong 2009; Chong, Familiar, and Shim 2016; O'Herron and von der Heydt 2009; Muckli, n.d.; Stokes et al. 2009; Kornmeier and Bach 2012; Wokke et al. 2013; Kornmeier and Bach 2004; von der Heydt 2015; Seghier et al. 2000; O'Callaghan et al. 2016; Lee and Nguyen 2001; Parkkonen et al. 2008; Somers et al. 1999; Cheron et al. 2014; Smith and Muckli 2010; Lee 2002; Hochstein and Ahissar 2002). In fact, Lamme and Roelfsema have proposed, based on backward masking and transcranial magnetic stimulation studies, together with measured latencies of activity in different classes of interconnecton, that the fast feedforward activity through the various hierarchical cortical regions can be identified primarily with pre-attentive and unconscious visual processing, while the slower lateral and recurrent feedback activity responsible for most cognitive and contextual, attention-driven processes of perception are to be identified primarily with conscious visual experience (Lamme and Roelfsema 2000) Such a division would correspond to the feedforward pathway being identified with the "SP" column of Figure 2, while the collateral and recurrent feedback processes would correspond functionally to the "WM" column of the figure. Obviously in this case conscious visual experience is to be identified with the functioning and maintenance of the internal model rather than with direct sensory input.[17]

Thus it appears that phenomenal experience is largely fabricated by the observer in real-time, with reliance on sensory input primarily for adjustment of the internal model to synchronize it to a statistical best fit to the input stream. The production of our visual experience is inferential, not analytic (Young 2000). The weight of neurobiological evidence supports the conclusion that the majority of the information provided to our awareness (John Smythies claims 95%) is provided not by the distal object but by these other, internal sources, only synchronized to the external world by the current state of sensory input. A different way of saying this is that that of which we are perceptually aware is not a representation of the world but a representation of our *justified beliefs* about the world. Where the model contains the current state of our belief and justification in this sense is measured by the statistical error of the fit between the model's prediction and sensory input. The model is being continually adapted and adjusted to minimize this error.

17 None of this would seem to support a Direct Realist position.

At various hierarchical levels of the model, post-perceptual stages of analysis as well as cognition and long-term memory, all provide input and feedback to help instantiate the model from current knowledge or belief about the world. This instantiation, continuously corrected through synchronization to current input, then becomes not only the basis for our immediate sensory awareness, but also for longer-term epistemic understanding of the world which provides the basis for the model and is learned over time, beginning in infancy, as we explore the world. It is justified and reinforced when sensory-motor hypotheses, instantiated as neural models of our environment, are confirmed or disconfirmed. For example, when we reach out to touch something and receive tactile confirmation, the model is confirmed; if not the system learns how to improve the modeling process (Golub, Yu, and Chase 2015).

Conceptualizing what I have described here as a mechanism for perception risks imposing our own conceptual categories on a brain that was evolved rather than designed. I have already discussed the central place of the model in motor control, and there is evidence that it also participates in many other mental functions. For example, there is psychophysical and neurological evidence that it is recruited as a simulation engine to extract intentions from the motions of other animate beings and to predict their future actions by simulating the observed action and estimating the actor's intentions based on a representation of one's own intentions. In particular, neurophysiological evidence supports the existence of a matching system between perception and action, which is recruited during imitation (Blakemore and Decety 2001). It should not be surprising that a general system for modeling and simulation acts similarly in service to many other functions as well as those we categorize under our artificial concept of “perception.”

11 Wait, that’s not what I meant!

All of this is of course empirical science which never has the luxury of certainty. This is just our best theoretical interpretation of the current state of the empirical, experimental evidence. It has only gotten better-supported with time and new techniques, but that could change tomorrow. Also, there are of course alternative interpretations, conflicting evidence and theoretical differences in the extensive neurobiological literature. But, for the sake of argument, let us suppose that this picture is, at least in outline, correct. Further, although the real biology has no such uniformity of function nor clear-cut demarcations of levels as portrayed in figures 2 and 3, for the sake of manageable discussion, and as a beginning rather than an end, let us suppose that at least the major functional concepts are captured by something similar to the almost comically-oversimplified picture in the figures. Specifically:

1. Predictive representation - The content of perceptual awareness consists of a continuously evolving model of the world that is the nervous system's best estimate of the current state of the external world.

2. Peceptuo-motor integration – The model is the motor system's guide to the anticipated state of the external world, and the perceptual system's guide to the effects of intended action on the sensory state.
3. Hierarchical organization – The model exists at many scales of description, in both space and time, with many levels of granularity. The descriptions at every level are simultaneously available to awareness.
4. Bayesian estimation - The model represents a belief about the state of the world that is justified by continuous testing of its predictions against raw sensory input, and constantly adjusted to maintain a statistical best fit to the input.
5. Top-down driving and modulation - In addition to perceptual belief instantiated in the current state of the model itself, it employs non-phenomenal inputs such as knowledge about known properties of objects and naive physics, and even volition, to constrain the time-evolution of its synthesized model of the world driven by both bottom-up sensory inputs and top-down inputs from non-sensory sources.
6. Final common path – The perceptual mechanism realized by the model is a final common path for phenomenal experience regardless of its origin, such as dreams, hallucinations, and imagination, in addition to sensory input.

What then can we see in these ideas that enables us to build bridges to philosophical terminology, even if the fit is not perfect? The above-mentioned features of the predictive model hypothesis suggest that its acceptance would imply certain positions concerning intentionality, representationalism, cognitive penetration, and direct realism among others. I examine some of these here.

Representation

Philosophers tend to put great stress on falsifiability and conditions of satisfaction in discussing representation; biologists and control theorists not so much, although many mathematical techniques in control theory carry an implicit assumption of incorrect data. By almost anyone's definition the internal model as presented here is a representation. It instantiates both physically (in neural activity) and phenomenally (in conscious awareness), information that represents to us that the world is a certain way that is verifiable. Beyond this it is less clear what sort of representation it is. The hierarchical structure of the model and the multiple paths of information flow and interaction admit of various interpretations at different levels, so it is far from clear that any single notion of representation applies throughout. Its bottom-up function in ordinary perception fits the definition used earlier in the present paper, derived from causal information flow and mapping of information spaces. It is driven by causal processes originating in the distal object, and it certainly realizes, when successful, a valid mapping of its information space. Indeed, that is the whole point of its hypothesized operation. On the other hand, taking this view alone treats other inputs such as semantic, top-down information, simply as perturbations and noise like any distortion in the causal chain, and this clearly ignores important new aspects of the representation.

If instead we look at the situation with dreams or hallucinations in the absence of sensory input, the suggestion from neurobiology is that these employ the same mechanism for representation of their phenomenal character. In this case however, in place of the “distal object” there would be the internal, neural processes responsible for generating or recalling from memory the content of the dreams or hallucinations. That such internal neural processes produce hallucinations has been well-understood since the pioneering work of Penfield. (Penfield and Rasmussen 1950) In this case, both causal information flow and veridical mapping could plausibly be assumed between these processes and the mechanisms of phenomenal representation. The phenomenal character presented by the model would then fit the same definition of representation as for veridical perception, but “stood on its head”, with top-down input from internal processes as the input in place of bottom-up input from a distal object. If you accept that ordinary phenomenal character in veridical perception is your phenomenal experience of neural firing in the visual cortex, then it is certainly no great leap to suppose that hallucination or dreaming is your phenomenal experience of the neural activity generating those events, presented to you via top-down inputs to the same internal modeling apparatus. A composition of these two views of the model could be considered as two, *interacting*, instances of the original definition of a representation, each with its own input. This of course fits nicely with the characterization of perception as a “controlled hallucination.”

Since our access to the model seems to be of all hierarchical levels simultaneously, it is as if we were looking down through a transparent stack of representations at different levels of description and possibly of different sorts. Consider this image.

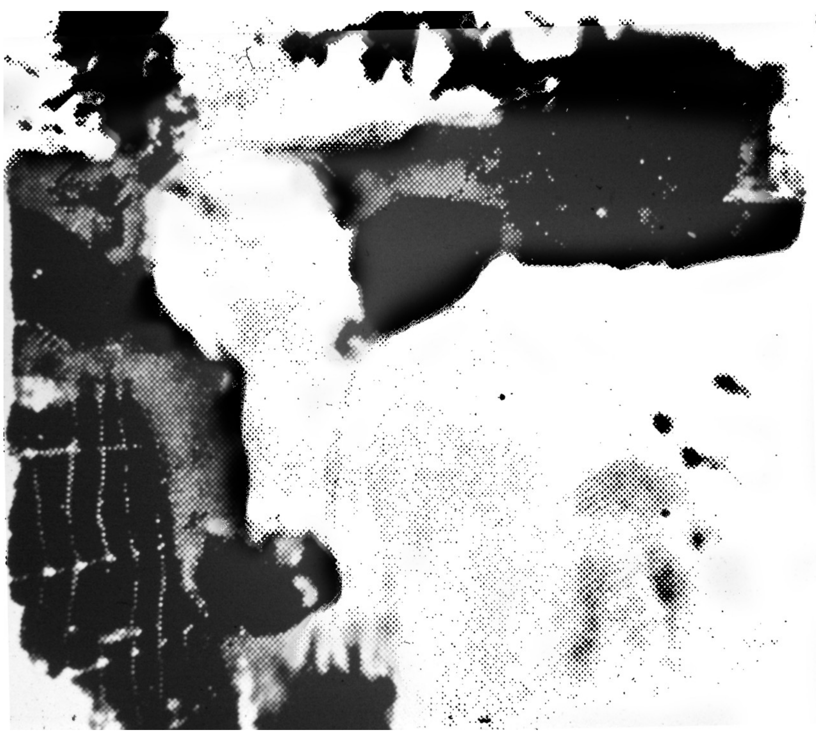

Figure 4: An image demonstrating multiple levels of organization.

Based on a sample of 5000 undergraduate psychology students, about 5 percent of viewers immediately understand it correctly, about 90 percent do so with a few hints, and about 5 percent never can see what it is. (These percentages may vary here depending on just how this slide reproduces in print.) Assuming you are in the 90%, you are initially aware of a black and white image of a number of irregular blobs. If I tell you it is a newspaper photograph of a nice black and white cow, standing in front of a fence and some trees, having its head, with black ears and nose, to your left and turned towards you, and its white flank to the right with a few black marks, you probably "get it" and see the image for what it is – a representation of a cow.

The point of interest here is that from the time you didn't see the cow to the time you did, something definitely changed in your experience. In fact, once you "see" the cow it is very hard to go back and not see it (the perceptual change has hysteresis as would be predicted by the control-theoretic aspects of the model.) Yet, notwithstanding the difference in your experience, if you were to draw accurately exactly what you saw both before and after "seeing" the cow, your drawings would be identical. None of the blobs changed. Two different levels of representation in the "stack" are at play here. One is presenting fine-grained, local properties. The other is presenting a phenomenal organization of the image that alters its phenomenal character at a more global level. One deals in properties of local areas such as black and white, the other with properties such as grouping of related regions. This illustrates the separability of these levels of representation, and the fact that if we do not achieve one level of representation, the other is still available to us. There is a separate "what it is like" to the experience of each of them.

How might these different levels of representation be characterized? Consider first just the "bottom up" input from the sensory channels. The model's representation of this information would certainly seem to fit the definition of "wide" or "broad" content, since it is a representation derived from information that is directly inherent in the current state of the environment. There is at least one clear sense in which the model's representation of this information already is intentional even at the lowest level. This is because what it represents is a predictive hypothesis for testing against the incoming sensory data. The model always represents not the current sensory input, but rather a justified belief[18] about the current state of the world, built up over prior states, and where justification occurs when current input matches predicted state within some limit. If it is true that the model is the content of awareness, then awareness is everywhere intentional in the sense of being a belief about the way the world is and how it will become. That it qualifies as a belief seems obvious from the model's role as a guide to the planning of impending actions of the motor system. This is not the less true because we are considering a non-conceptual level of organization.

What of higher levels of representation dominated by top-down input? In the hypothesis presented here this represents information of non-sensory (or at least not current sensory) origin. This may include anything from hardwired evolutionary function-

18 In the sense of an assumption about the state of the world that may be true or false.

ality to learning and memory, to reason and volition; essentially *a priori*, semantic knowledge about the world and intentional types. There is of course precious little information about how any of this is represented prior to phenomenal instantiation, and most of it is almost certainly represented in different ways. Notwithstanding, there is good empirical evidence, both neurophysiological and phenomenal, that it does influence the phenomenal model of the world. Those influences effect changes in the information content of the biological/phenomenal realization, and hence in the information state of the model. What we can then ask is how this interacts with the bottom-up driving of the information state and how the result resembles positions on offer in the philosophical literature.

Assuming that whatever is available from current sensory input is already in the model, new information available from the top-down input seems likely to be what is often called "narrow content"; that is, intentional content that does not depend on the details of the current environment. Thus, knowledge of naive physics that tells us how things fall, without being about the trajectory of the particular ball currently approaching my body-image in my internal model, or knowledge of a cow as a "type of thing" with certain properties that can be tested against the collection of features currently before me. In veridical perception the interaction of the top-down and bottom-up information is that the top-down information is "modulatory" in that it provides a hypothesized higher-level and larger scale organization of features and properties which, *if justified by subsequent input*, is instantiated into the current state of the predictive model, thus improving the scope and accuracy of its representation of the world. This instantiation then can be considered as producing a token of a top-down input type and, in the process "widening" the narrow-content input and adding additional intentionality to the representation. As this proceeds in an ever more fine-grained manner from top to bottom we could perhaps think of it it in terms of "degree" of wideness and of intentionality of representation at different levels rather than asking whether the representation is wide or narrow or what its intentional content might be. It is not obvious to me however that trying to force the control-theoretic picture of the model's operation into this terminology provides any particular improvement to understanding.

Putting the argument presented earlier concerning awareness as the final representation together with the view of representation diagrammed in Section 3, figure 1, it appears that some positions on phenomenal character and properties, and on representational properties, are implied. First, awareness stands as the representation at the lower right in this figure. Phenomenal properties are then the realization of the information space in the upper right, and phenomenal character is the conscious experience of these properties. Further, while phenomenal properties represent properties of the mind-independent distal object, and, if connected to them by causal information flow, represent them veridically to the extent that a valid mapping exists between the two information spaces, these are clearly intrinsic properties of the observer. The mappings (vertical arrows) between the information states and the things that realize them are incomparable and at least potentially idiosyncratic, and the particular properties of the realizations must be intrinsic. It is the information states, not the realizations that determine veridicality of representation as used here. Their incomparable realizations are the media of representation.

At the lowest levels of the hierarchy, as discussed above, there is a sort of primitive intentionality to the model in the sense that it is “about” the expected state of things in a predictive sense. Even in the absence of any top-down influence one would suppose that there was still an expectation about the state of the world based simply on the time-history of the input variables and the “knowledge” that the world does not usually change abruptly at random. This is the essence of Kalman filtering (Becker 2021). This sort of knowledge is likely hard-wired into the system by evolution, and may be, but need not be, influenced by top-down input from higher levels of representation. In this case, to the extent that there is any difference between them, representational properties and intentionality would seem to be equivalent, and there is no obvious sense of priority or how one might ground the other. At the upper levels of representation in the hierarchy, the issue may be less clear cut. The narrow types presumed to provide input to the phenomenal model enter phenomenal experience when instantiated as tokens. However, still speaking only of ordinary sensory perception, these enter phenomenal experience as modulations and organizations of bottom-up phenomenal character. Here I would lean towards regarding phenomenal properties as prior to (phenomenal) intentionality. In the case of hallucination and dreams on the other hand, the top-down input would be driving rather than modulatory and better regarded as prior to phenomenal character and representation. (Although we need to know more about the possibly hierarchical nature of generation of content for hallucinations and dreams.)

It has been argued by some that only low-level properties from the receptors are phenomenally represented. (Dretske 1995) Others have argued that high-level properties are represented. (Bayne 2009) The predictive model position as just discussed certainly argues for awareness of higher-level properties, however it is not obvious that in ordinary perceptual experience they are manifest other than as influences on the manner in which lower-level properties are instantiated in our awareness. Thus, in this case it is possible that awareness of higher-level properties depends upon lower-level properties. However, if it is accepted that this same mechanism is the final common path for phenomenal experience of dreams and hallucinations, it seems plausible that top-down inputs could manifest as higher-level phenomenal character in the absence of bottom-up driving, or in cases of unusually intense top-down driving. We might then expect to see representations down to some level of detail, but not necessarily penetrating to levels of fine-grained resolution.

Cognitive Penetration

Cognitive penetration – the influence of higher-level, non-phenomenal, semantic information on phenomenal perception – has been a debated topic among philosophers.(Cavedon-Taylor 2018; Zeimbekis and Raftopoulos 2015; Vetter and Newen 2014) Allowing for terminological differences, it is supported by considerable evidence from neurobiology as discussed earlier, and is an ordinary part of control-theoretic systems. Of relevance to the present discussion is the argument that cognitive penetration is most easily explained in terms of the sort of controlled predictive model presented here. This is especially evident in time-varying, dynamic illusions that merge the constraints of sensory input and voluntary subject actions with cognitive assumptions, producing non-veridical representations of the world. A dynamic version of the Mach Strip illusion presented below illustrates this.

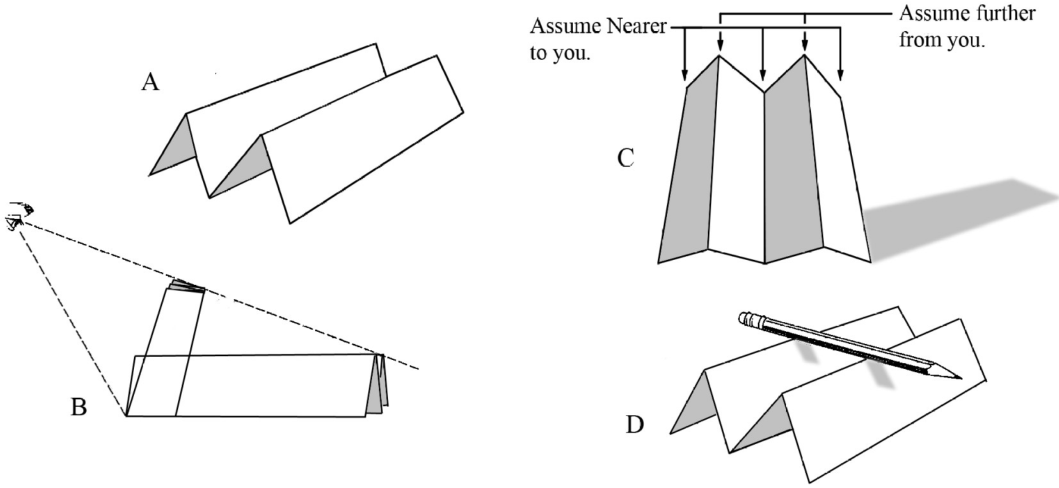

Figure 5: (A) How the paper is folded. (B) How to view it, showing the actual and the illusory positions. (C) How you see it. (D) How the pencil is positioned.

First fold a sheet of paper across the middle and then fold each half the other direction to make a construction as shown in “A”. Set it on the floor in front of your chair and view it lengthwise as shown in “B.”

The image on your retina is then a two-dimensional projection, and so it is ambiguous with respect to the two orientations shown in “B.” You now need to make it “stand up” as shown in “C.” This requires making a mental shift in your assumptions about its orientation very similar to that required to flip a Necker Cube. This can take a bit of practice because your brain knows the real orientation and wants to hold on to it. It will help to close one eye to reduce depth cues and to think of the locations of the upper corners as shown in “C.” Some people see it immediately, others need to work on it a bit. After you once see it, it gets easier and with a little practice you can flip it up and down at will.

The next step is to slowly move your body laterally from right to left and back. You will see the paper shape move, twist, and deform to track your movements. Your brain may reject this and flip it back down again, since it contradicts you knowledge about the paper, but with a little practice you can hold it upright while moving freely to and fro and side to side. You may also notice other effects such as non-veridical illumination as discussed below.

As a final demonstration, once you can easily maintain the upright phenomenal orientation while moving, place a pencil on the paper as shown in “D.” This may again cause you difficulty in maintaining the orientation, but a bit of practice will again put you right. Once again, observe the effects of body movement on your phenomenal experience. I will not spoil it by telling you what you will see, but the answer to “why” is that you are using incorrect assumptions about the paper, but not about the pencil.

Discussion of the result

Assuming you now can run the experiment, you no doubt noticed that the object shifted, twisted and deformed in exactly the required way to both maintain veridical constraints such as the alignment of points on the paper with background points, while at the same time honoring your assumption about its three-dimensional orientation. Everything, in fact, remains consistent with the image on the retina. Depending on your illumination, you might also have noticed the that object appeared to “glow” slightly or have an internal light source in places. This is a similar illusion that is required to make the assumed three-dimensional orientation consistent with the way the parts of the object shadow itself and produce various luminance values on the retina, given what you know about the sources of light.

In short, the actual retinal image, feed-forward from the motor system, and your assumptions about the external world are being kept in correspondence in real time by what must be an enormous amount of computation that you could not possibly be doing consciously. Indeed, everyone sees the same results whether or not told about them in advance, so the result is not a matter of conscious choice, despite that the initiating assumption clearly is. This servoing of an internally-maintained best-model of the environment to the results of motor action and sensory input is a normal process that goes on all the time in ordinary perception. Usually, however, the assumptions about the external world are taken from stored knowledge, both learned and evolved, and from ongoing successful validation of the current state of the running model. Hence normally our phenomenal experience of the model matches our expectations and we notice nothing unusual.

This is a system that has evolved as a control system for the organism, enabling it to operate efficiently on the basis of noisy and incomplete sensory input, and it appears to use exactly the same principles as those of modern control theory. The system probably did not evolve to serve conscious awareness, and conscious awareness is probably not necessary to its efficient operation, although obviously conscious assumptions may influence the model. Other standard demonstrations of voluntary, cognitive effects on phenomenal perception, including flipping of a Necker cube, figure-ground reversals, and ambiguous figure reorganizations, also illustrate various properties expected from the control-theoretic view, such as hysteresis, but the notable feature of demonstrations of the Mach Strip variety is to underscore the tight integration of such effects with other neural functionality such as motor control. This is an entirely expected result on the hierarchical, predictive model hypothesis but much more difficult to incorporate into other explanations of cognitive penetration.

Adverbialism

It has been suggested to me that the position taken here that awareness is itself the final representation in a causal chain has similarities to adverbialism, and I think that is worth consideration. Adverbialism opposes the notion that phenomenal experience is the result of an act of awareness operating on a representation, but that is a different matter from it *being* the representation. It does not seem that experiencing red ballishly is inconsistent with that being how I represent that a red ball is before me. Siegel discusses two concepts of perceptual content by using the analogy of the "contents of a bucket" and the "contents of a newspaper story".(Siegel 2016) That is, a distinction between containing a thing and containing information about it. We might see the contents of the "newspaper" as the state of the information space and the contents of the "bucket" as the state of that which realizes it. We could consider awareness to be the bucket and the two senses of content to correspond to the state of awareness (state of the bucket) and the state of the information space it realizes (state of the newspaper.) This nicely captures awareness as a representation according to the definition, while retaining the adverbialist notion of phenomenal experience as a property of the observer. Assuming that there is a brain state which either is, or is the neural correlate of, phenomenal awareness, we could describe the brain state and phenomenal awareness either as two buckets containing different but nomically-linked realizations of the same newspaper story, or if you prefer, two aspects of the content of the same bucket. In either case, awareness-as-a-representation is the *final* representation in a causal chain originating in the distal object, but, as the adverbialist prefers, not the object of an act of awareness.

Intrinsic vs. extrinsic character

How do issues such as reducibility and intrinsic vs. extrinsic character map onto the concepts of the internal model? Again, we may fail to find an answer in all-or-none terms. At the lowest level where we are dealing in the most fine-grained way with local properties (such as color and luminosity in the case of vision) we find the least temporal smoothing (the prior state of the model is viewed as an intentional statement or belief about the expected next state.) We also find the least penetration by cognitive factors (although they seem not to be non-existent.) This suggests that here there is a close relation between phenomenal properties and representational content, and that in this case representational content reduces to phenomenal property. In control theory terms, sensory input is "driving." There is not much here to support that the representation is of anything but intrinsic phenomenal character. If we look back at Figure 1, the model at every level qualifies as a representation to the extent that there is an isomorphism between the state of its information space and the state of the information space of a distal, mind-independent object. However, because the mappings (vertical arrows) between the information states and the things that realize them are incomparable and at least potentially idiosyncratic, the particular properties of the realizations must be intrinsic. It is the information states, not the realizations that enter into defining representation as used here. Their realizations are the medium of representation.

Direct vs. Indirect Realism

Direct Realists often emphasize the existence of a relation between a perception and the object of perception, although they may employ this idea in different ways ((Alston 1999; Huemer 2001)), and the neural world-model obviously possesses a relation to the distal object. The distal object and its neural model are on either end of a causal, physical relation in which changes in the distal object, or its spatial relation to other objects, or to the observer, result in changes in the neural model. If as suggested, the content of perceptual awareness can be identified with the information content of this internal neural model, then it too bears this same relation to the distal object. However admitting that this mechanism comprises a relation seems not to add anything to the discussion that carries any concept of "directness."

Although Direct Realism comes in many flavors, we can discuss two general types. The first simply asserts that the distal object is itself in some fashion an actual part of our perception. This type of Direct Realism, often referred to as some variant of "Naive Realism" seems to have no obvious point of contact with the ideas discussed above. So far as I can tell, it suffers from being unfalsifiable, and is motivated only by the conviction that the world "out there" must in some sense "look like" our conscious experience of it. In the context of the present work I think there is little more to be said about it.

A second type of Direct Realism (often confusingly referred to as "Representationalism,") acknowledges the necessity of the physics and neurophysiology, and even of internal representations, but asserts that perception is nonetheless direct because it is identified with the distal object and we are unaware of the medium of the representation, and experience only the information the medium embodies. (Dretske 1995; Tye 1995)If it were the case that we are aware of an internal representation in the act-object sense, but unaware of the medium of the representation (and I believe it is the case according to some indirect Realist positions,) I would still have an objection to labeling such a situation as "direct." It seems to me simply to be the application of a different label to the same situation that changes nothing, while the "indirect" label for the situation seems more apt. In this case, I would say pick your favorite label and move on.

However, the position of the present argument is not quite the same. Since awareness as argued above must itself be a representation, and in fact the final one in a continuous chain, it is not the case that awareness is of an *intervening* representation. It is however awareness of an internal representation nonetheless. It is not obvious in this case that the question of whether or not we are aware of the medium is even well-posed, and may be incoherent. Even if it is well-posed, however, I don't think the answer would distinguish in any meaningful way between the position being direct or indirect.

In general, awareness of hallucinations, dreams, imaginations and the like would not be considered direct perceptions by most, but it should be evident at this point that the view of perceptual awareness that is supported by the neurobiological evidence is essentially a construct of this same sort. That is, ordinary perceptual awareness of the world is synthesized from epistemic knowledge and beliefs at varying time-scales and on multiple hierarchical levels, forming hypotheses which are continuously justified by statistical fitting to current sensory data. On this hypothesis, the information content of our awareness is derived by synthesis and correction of stored information, ranging from longer-term beliefs and knowledge to the most recent prior state of the model itself, but in any case, stored information. In this regard perceptual awareness is like an hallucination which appears to be a distal object but is something else. This position would make it difficult to understand any claim of a direct relation between perception and the distal object.

Justification of perceptual belief

Given all of this, it is reasonable to ask how perceptual belief can be justified. Philosophers use terms such as 'belief' and 'knowledge' in very precisely-defined ways, conceived on quite other bases than neurobiological function. I think that at least if we restrict ourselves to non-verbal, perceptual knowledge and beliefs, verification of prediction ("predictive validity") has a close relation to what is often thought of as justification.

The dynamics of the model-maintenance processes embody a very important mechanism, which is feedback-confirmation of the model's predictions. If the model is correct and the relation between the predictive model and the physical world is to be accurately maintained, then actions on the part of the observer should result in predicted changes in his or her perceptual awareness. The changes should be in accord with beliefs or knowledge about the naive physics of the world, and with expected contacts with the objects in it. The model generates expectations of how light and shadow will move across a surface, or how an object's alignment with more distant objects will change under observer motion, or with anticipated visual or auditory stereopsis, and many other moment-to-moment sensory inputs. If these predictions are within the range of acceptable error, the model's predictive validity is confirmed. When it is dis-confirmed the model is rejected and hopefully successfully modified. This process spans a time-course of milliseconds at the lowest hierarchical levels of the model and minutes or longer at the highest. It is by this means that from childhood we learn perceptual understanding. Predictive validity is a hard-wired, fundamental feature of our brains. In neural systems, predictive validity of stimulus associations is the basis of Classical Conditioning, and predictive validity of action is the basis of instrumental learning. In the case of perception, whenever predictive validity is confirmed we learn beliefs about the world, and how to act within it based on our sensory input and prior information (Lalazar and Vaadia 2008). The predictive correlation between distance senses and connative actions, as confirmed by contact senses, is of paramount importance since every individual's body is different, and hence the correlation *must* be learned. Predictive validity of vision and other distance senses being confirmed by touch thus exerts a powerful influence on our perceptual beliefs.[19] Everyone has always known this:

"Is this a dagger which I see before me,
The handle toward my hand? Come, let me clutch thee.
I have thee not, and yet I see thee still.
Art thou not, fatal vision, sensible
To feeling as to sight? or art thou but
A dagger of the mind, a false creation,
Proceeding from the heat-oppressed brain?"

19 The primacy of contact senses in grounding distance senses, at least during early learning, my have evolutionary roots in the antiquity of contact senses.

– Shakespeare, “Macbeth.”

12 Summary

An information-based definition of representation based on the flow of the inherent information contained in physical objects is advocated as applicable to both physical and phenomenal objects, and it is argued that according to this definition awareness itself must be considered the final representation in a continuous causal chain of representations originating in the distal object.

It is argued that the information content of perceptual awareness is that of a continuously-evolving model of the world that is the nervous system's best estimate of the current state of the external world, in which “bottom-up”, non-intentional representation interacts with “top-down”, intentional states to produce a hierarchical representation from which both, and their mixed states, are simultaneously accessible. This estimate is a belief about the state of the world that is justified by continuous testing against raw sensory input, and constantly adjusted to a best fit to it, and actively maintained by Bayesian control theory principles.

It accounts for and anticipates changes due to our actions in the world by virtue of input from the motor system, and it employs epistemic beliefs and knowledge, including in particular perceptual belief instantiated in the current state of the model itself, to constrain the time-evolution of its synthesized model of the world. This view is examined in the context of current neurobiological investigations of sensory and perceptual processes.

Some implications of this position are examined with respect to concepts in the philosophy of perception which appear to bear relation to it. I suggest that this approach, although not fundamentally at odds with many concepts in philosophy, may still offer a more coherent way to view mind-dependent aspects of perception than some traditional philosophical positions.

Bibliography

Adams, Frederick. 2003. “The Informational Turn in Philosophy1.” *Minds and Machines* 13 (4): 471–501. https://doi.org/10.1023/A:1026244616112.

Albers, Anke Marit, Peter Kok, Ivan Toni, H. Chris Dijkerman, and Floris P. de Lange. 2013. “Shared Representations for Working Memory and Mental Imagery in Early Visual Cortex.” *Current Biology: CB* 23 (15): 1427–31. https://doi.org/10.1016/j.cub.2013.05.065.

Aleman, André, Geert-Jan M. Rutten, Margriet M. Sitskoorn, Geraud Dautzenberg, and Nick F. Ramsey. 2001. “Activation of Striate Cortex in the Absence of Visual Stimulation: An FMRI Study of Synesthesia.” *NeuroReport* 12 (13): 2827.

Alston, William P. 1999. “Back to the Theory of Appearing.” *Nous* 33 (s13): 181–203. https://doi.org/10.1111/0029-4624.33.s13.9.

Bailey, Andrew. 2007. “Representation and a Science of Consciousness.” *Journal of Consciousness Studies.* 14 (1): 62–76.

Barlow, H. B. 1994. “What Is the Computational Goal of the Neocortex.” In *Large-Scale Neuronal Theories of the Brain*, edited by Christof Koch and J. Davis, 1–22. MIT Press.

Barlow, Horace. 1987. “Cerebral Cortex as Model Builder.” In *Matters of Intelligence*, edited by Lucia M. Vaina, 395–406. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-009-3833-5_18.

Barton, Robert A., and Chris Venditti. 2014. “Rapid Evolution of the Cerebellum in Humans and Other Great Apes.” *Current Biology* 24 (20): 2440–44. https://doi.org/10.1016/j.cub.2014.08.056.

Bayne, Tim. 2009. “Perception and the Reach of Phenomenal Content.” *The Philosophical Quarterly* 59 (236): 385–404.

Becker, Alex. 2021. “Online Kalman Filter Tutorial.” 2021. https://www.kalmanfilter.net/.

Berkeley, George. 1709. *An Essay Towards a New Theory Of Vision*. Edited by David Wilkins. 1st ed.

Blakemore, Sarah-Jayne, and Jean Decety. 2001. “From the Perception of Action to the Understanding of Intention.” *Nature Reviews Neuroscience* 2 (8): 561–67. https://doi.org/10.1038/35086023.

Brown, Harriet, Karl J. Friston, and Sven Bestmann. 2011. “Active Inference, Attention, and Motor Preparation.” *Frontiers in Psychology* 2. https://doi.org/10.3389/fpsyg.2011.00218.

Cavanagh, Patrick. 2011. “Visual Cognition.” *Vision Research*, Vision Research 50th Anniversary Issue: Part 2, 51 (13): 1538–51. https://doi.org/10.1016/j.visres.2011.01.015.

Cavedon-Taylor, Dan. 2018. “Naïve Realism and the Cognitive Penetrability of Perception.” *Analytic Philosophy* 59 (3): 391–412. https://doi.org/10.1111/phib.12131.

“Cerebellum.” 2020. In *Wikipedia*. https://en.wikipedia.org/w/index.php?title=Cerebellum&oldid=965469343.

Cerminara, Nadia L., Richard Apps, and Dilwyn E. Marple-Horvat. 2009. “An Internal Model of a Moving Visual Target in the Lateral Cerebellum: Cerebellar Internal Model of Target Motion.” *The Journal of Physiology* 587 (2): 429–42. https://doi.org/10.1113/jphysiol.2008.163337.

Chalmers, David J. 1997. *The Conscious Mind: In Search of a Fundamental Theory*. Revised edition. New York: Oxford University Press.

———. 2004. “The Representational Character of Experience.” In *The Future for Philosophy*, edited by B. Leiter. Oxford: Oxford University Press. http://www.consc.net/papers/representation.html.

Cheron, Guy, Axelle Leroy, Ernesto Palmero-Soler, Caty De Saedeleer, Ana Bengoetxea, Ana-Maria Cebolla, Manuel Vidal, Bernard Dan, Alain Berthoz, and Joseph McIntyre. 2014. "Gravity Influences Top-Down Signals in Visual Processing." *PLOS ONE* 9 (1): e82371. https://doi.org/10.1371/journal.pone.0082371.

Chong, Edmund, Ariana M. Familiar, and Won Mok Shim. 2016. "Reconstructing Representations of Dynamic Visual Objects in Early Visual Cortex." *Proceedings of the National Academy of Sciences* 113 (5): 1453–58. https://doi.org/10.1073/pnas.1512144113.

Cichy, Radoslaw M., Jakob Heinzle, and John-Dylan Haynes. 2012. "Imagery and Perception Share Cortical Representations of Content and Location." *Cerebral Cortex* 22 (2): 372–80. https://doi.org/10.1093/cercor/bhr106.

Clark, Andy. 2013. "Whatever next? Predictive Brains, Situated Agents, and the Future of Cognitive Science." *Behavioral and Brain Sciences* 36 (3): 181–204. https://doi.org/10.1017/S0140525X12000477.

Collier, John. 1990. "Intrinsic Information." In *Information, Language and Cognition*, edited by P. Hanson, 1:390–409. Vancouver Studies in Cognitive Science. Oxford University Press.

Cullen, Kathleen E., and Jessica X. Brooks. 2015. "Neural Correlates of Sensory Prediction Errors in Monkeys: Evidence for Internal Models of Voluntary Self-Motion in the Cerebellum." *Cerebellum (London, England)* 14 (1): 31–34. https://doi.org/10.1007/s12311-014-0608-x.

David, A. S., P. W. Woodruff, R. Howard, J. D. Mellers, M. Brammer, E. Bullmore, I. Wright, C. Andrew, and S. C. Williams. 1996. "Auditory Hallucinations Inhibit Exogenous Activation of Auditory Association Cortex." *Neuroreport* 7 (4): 932–36.

De Weerd, Peter. 2006. "Perceptual Filling-in: More than the Eye Can See." *Progress in Brain Research* 154: 227–45. https://doi.org/10.1016/S0079-6123(06)54012-9.

De Weerd, Peter, and et al. 1998. "Perceptual FIlling-in a Parametric Study.Pdf." *Vision Research*, 2721–34.

Dorsch, Fabian. 2011. *Experience_and_Reason*. Rero Doc.

Dresler, Martin, Stefan P. Koch, Renate Wehrle, Victor I. Spoormaker, Florian Holsboer, Axel Steiger, Philipp G. Sämann, Hellmuth Obrig, and Michael Czisch. 2011. "Dreamed Movement Elicits Activation in the Sensorimotor Cortex." *Current Biology* 21 (21): 1833–37. https://doi.org/10.1016/j.cub.2011.09.029.

Dretske, Fred I. 1995. *Naturalizing the Mind*. The Jean Nicod Lectures. Bradford Books, M.I.T. Press.

Edwards, Grace, Petra Vetter, Fiona McGruer, Lucy S. Petro, and Lars Muckli. 2017. "Predictive Feedback to V1 Dynamically Updates with Sensory In-

put." *Scientific Reports* 7 (1): 16538. https://doi.org/10.1038/s41598-017-16093-y.
ffytche, D. H., R. J. Howard, M. J. Brammer, A. David, P. Woodruff, and S. Williams. 1998. "The Anatomy of Conscious Vision: An FMRI Study of Visual Hallucinations." *Nature Neuroscience* 1 (8): 738–42. https://doi.org/10.1038/3738.
Fiser, József, Pietro Berkes, Gergő Orbán, and Máté Lengyel. 2010. "Statistically Optimal Perception and Learning: From Behavior to Neural Representations." *Trends in Cognitive Sciences* 14 (3): 119–30. https://doi.org/10.1016/j.tics.2010.01.003.
Fleming, Roland W. 2014. "Visual Perception of Materials and Their Properties." *Vision Research* 94 (January): 62–75. https://doi.org/10.1016/j.visres.2013.11.004.
Francis, B. A., and W. M. Wonham. 1976. "The Internal Model Principle of Control Theory." *Automatica* 12 (5): 457–65. https://doi.org/10.1016/0005-1098(76)90006-6.
French, Robert. 2018. "Chapter Three - A Defense of Representational Realism." In *Direct versus Indirect Realism*, edited by John Smythies and Robert French, 15–31. Academic Press. https://doi.org/10.1016/B978-0-12-812141-2.00003-9.
French, Robert E. 1987. *The Geometry of Vision and the Mind Body Problem*. Peter Lang.
Friston, Karl. 2005. "A Theory of Cortical Responses." *Philosophical Transactions of the Royal Society B: Biological Sciences* 360 (1456): 815–36. https://doi.org/10.1098/rstb.2005.1622.
———. 2010. "The Free-Energy Principle: A Unified Brain Theory?" *Nature Reviews Neuroscience* 11 (2): 127–38. https://doi.org/10.1038/nrn2787.
———. 2012. "Prediction, Perception and Agency." *International Journal of Psychophysiology*, Predictive information processing in the brain: Principles, neural mechanisms and models, 83 (2): 248–52. https://doi.org/10.1016/j.ijpsycho.2011.11.014.
Gibson, James J. 1950. *The Perception of the Visual World*. The Riverside Press, Cambridge Mass. http://www.archive.org/details/perceptionofvisuOOjame.
Gilbert, Charles D., and Wu Li. 2013. "Top-down Influences on Visual Processing." *Nature Reviews. Neuroscience* 14 (5). https://doi.org/10.1038/nrn3476.
Goldstein, E. Bruce. 1981. "The Ecology of J. J. Gibson's Perception." *Leonardo* 14 (3): 191. https://doi.org/10.2307/1574269.
Golub, Matthew D, Byron M Yu, and Steven M Chase. 2015. "Internal Models for Interpreting Neural Population Activity during Sensorimotor Con-

trol." Edited by Timothy Behrens. *ELife* 4 (December): e10015. https://doi.org/10.7554/eLife.10015.

Grunwald, Peter D., and Paul M. B. Vitanyi. 2008. "Algorithmic Information Theory." *ArXiv:0809.2754 [Cs, Math, Stat]*, September. http://arxiv.org/abs/0809.2754.

Haar Romeny, Bart M. ter, Luc M. J. Florack, Jan J. Koenderink, and Max A. Viergever. 1991. "Scale Space: Its Natural Operators and Differential Invariants." In *Information Processing in Medical Imaging*, edited by Alan C. F. Colchester and David J. Hawkes, 239–55. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer. https://doi.org/10.1007/BFb0033757.

Harrison, Stephenie A., and Frank Tong. 2009. "Decoding Reveals the Contents of Visual Working Memory in Early Visual Areas." *Nature* 458 (7238): 632–35. https://doi.org/10.1038/nature07832.

Hatfield, Gary. 2011. "Philosophy of Perception and the Phenomenology of Visual Space." *Philosophic Exchange*, 31–66.

Helmholtz, H von. 1896. "Handbuch Der Physiologischen Optik, (Hamburg/Leipzig: Voss)." In *Treatise on Physiological Optics*.

Heydt, Rüdiger von der. 2015. "Figure–Ground Organization and the Emergence of Proto-Objects in the Visual Cortex." *Frontiers in Psychology* 6 (November). https://doi.org/10.3389/fpsyg.2015.01695.

"Hierarchical Control System." 2020. In *Wikipedia*. https://en.wikipedia.org/w/index.php?title=Hierarchical_control_system&oldid=984846978.

"Hierarchical Control System - an Overview | ScienceDirect Topics." n.d. Accessed October 31, 2020. https://www.sciencedirect.com/topics/engineering/hierarchical-control-system.

Hikosaka, Okihide, Kae Nakamura, Katsuyuki Sakai, and Hiroyuki Nakahara. 2002. "Central Mechanisms of Motor Skill Learning." *Current Opinion in Neurobiology* 12 (2): 217–22. https://doi.org/10.1016/S0959-4388(02)00307-0.

Hinton, Geoffrey E. 2007. "Learning Multiple Layers of Representation." *Trends in Cognitive Sciences* 11 (10): 428–34. https://doi.org/10.1016/j.tics.2007.09.004.

Hochstein, Shaul, and Merav Ahissar. 2002. "View from the Top: Hierarchies and Reverse Hierarchies in the Visual System." *Neuron* 36 (5): 791–804. https://doi.org/10.1016/S0896-6273(02)01091-7.

Hogan, Michael J. 2004. "The Cerebellum in Thought and Action: A Fronto-Cerebellar Aging Hypothesis." *New Ideas in Psychology* 22 (2): 97–125. https://doi.org/10.1016/j.newideapsych.2004.09.002.

Hohwy, Jakob, Andreas Roepstorff, and Karl Friston. 2008. "Predictive Coding Explains Binocular Rivalry: An Epistemological Review." *Cognition* 108 (3): 687–701. https://doi.org/10.1016/j.cognition.2008.05.010.

Horikawa, T., M. Tamaki, Y. Miyawaki, and Y. Kamitani. 2013. “Neural Decoding of Visual Imagery During Sleep.” *Science* 340 (6132): 639–42. https://doi.org/10.1126/science.1234330.

Howard, Robert, Anthony David, Peter Woodruff, John Mellers, Ian Wright, Mike Brammer, Edward Bullmore, and Steven Williams. 1997. “Seeing Visual Hallucinations with Functional Magnetic Resonance Imaging.” *Dementia and Geriatric Cognitive Disorders* 8 (2): 73–77. https://doi.org/10.1159/000106610.

Huang, Yanping, and Rajesh P. N. Rao. 2011. “Predictive Coding.” *Wiley Interdisciplinary Reviews. Cognitive Science* 2 (5): 580–93. https://doi.org/10.1002/wcs.142.

Huemer, Michael. 2001. *Skepticism and the Veil of Perception*. Studies in Epistemology and Cognitive Theory. Lanham, Md: Rowman & Littlefield Publishers.

Hutto, D D. 2009. “Mental Representation and Consciousness.” In *Encyclopedia of Consciousness*, edited by William P. Banks, 2:15. Oxford: Elsevier.

Igawa, M., Y. Atsumi, K. Takahashi, S. Shiotsuka, H. Hirasawa, R. Yamamoto, A. Maki, Y. Yamashita, and H. Koizumi. 2001. “Activation of Visual Cortex in REM Sleep Measured by 24-Channel NIRS Imaging.” *Psychiatry and Clinical Neurosciences* 55 (3): 187–88. https://doi.org/10.1046/j.1440-1819.2001.00819.x.

Imamizu, Hiroshi, Satoru Miyauchi, Tomoe Tamada, Yuka Sasaki, Ryousuke Takino, Benno Pütz, Toshinori Yoshioka, and Mitsuo Kawato. 2000. “Human Cerebellar Activity Reflecting an Acquired Internal Model of a New Tool.” *Nature* 403 (6766): 192–95. https://doi.org/10.1038/35003194.

Ito, Masao. 2008. “Control of Mental Activities by Internal Models in the Cerebellum | Nature Reviews Neuroscience,” April, pages 304-313.

Jeannerod, M. 1994. “The Representing Brain: Neural Correlates of Motor Intention and Imagery.” *Behavioral and Brain Sciences* 17 (02): 187. https://doi.org/10.1017/S0140525X00034026.

Jehee, Janneke F. M., and Dana H. Ballard. 2009. “Predictive Feedback Can Account for Biphasic Responses in the Lateral Geniculate Nucleus.” *PLoS Computational Biology* 5 (5). https://doi.org/10.1371/journal.pcbi.1000373.

Jerath, Ravinder, Molly W. Crawford, and Vernon A. Barnes. 2015. “Functional Representation of Vision within the Mind: A Visual Consciousness Model Based in 3D Default Space.” *Journal of Medical Hypotheses and Ideas* 9 (1): 45–56. https://doi.org/10.1016/j.jmhi.2015.02.001.

Kandel, et al, E., ed. 2014. “High-Level Visual Processing: Cognitive Influences.” In *Principles of Neural Science, Ch. 28*, Fifth Edition. New York,

NY: McGraw-Hill Education. neurology.mhmedical.com/content.aspx?aid=1101680178.
Kawato, Mitsuo. 1999. "Internal Models for Motor Control and Trajectory Planning." *Current Opinion in Neurobiology*, no. 9: 718–27.
Kent, Ernest, and James S. Albus. 1984. "Servoed World Models as Interfaces Between Robot Control Systems and Sensory Data." *Robotica* 2 (December). https://www.nist.gov/publications/servoed-world-models-interfaces-between-robot-control-systems-and-sensory-data.
Kerkoerle, Timo van, Matthew W. Self, Bruno Dagnino, Marie-Alice Gariel-Mathis, Jasper Poort, Chris van der Togt, and Pieter R. Roelfsema. 2014. "Alpha and Gamma Oscillations Characterize Feedback and Feedforward Processing in Monkey Visual Cortex." *Proceedings of the National Academy of Sciences* 111 (40): 14332–41. https://doi.org/10.1073/pnas.1402773111.
Knauff, Markus, Jan Kassubek, Thomas Mulack, and Mark W. Greenlee. 2000. "Cortical Activation Evoked by Visual Mental Imagery as Measured by FMRI." *NeuroReport* 11 (18): 3957.
Knill, David C., and Alexandre Pouget. 2004. "The Bayesian Brain: The Role of Uncertainty in Neural Coding and Computation." *Trends in Neurosciences* 27 (12): 712–19. https://doi.org/10.1016/j.tins.2004.10.007.
Koenderink, J. 1984. "Geometrical Structures Determined by the Functional Order in Nervous Nets." *Biol. Cvybern.* 50: 43–50.
Koenderink, J., and A. J. van Doorn. 1982. *Cybernetic Model for Simultaneous Order of Signals in the Optic Nerve. Proc Cvb Soc Vienna, 1980. Reprinted in in: Progress in Cybernetics and Systems Research IX.153-159*. Edited by Pask G. Trappl R. Ricciardi L. Hemisphere Publ.Comp., Washington.
Koenderink, J. J. 1987. "The Construction of a Simultaneous Functional Order in Nervous Systems. I. Relevance of Signal Covariances and Signal Coincidences in the Construction of a Functional Order." *Biol. Cybernet.* 57: 115-125.
Koenderink, J. J., and A. J. van Doorn. 1987. "Representation of Local Geometry in the Visual System." *Biological Cybernetics* 55 (6): 367–75. https://doi.org/10.1007/BF00318371.
Koenderink, Jan J. 1984. "Simultaneous Order in Nervous Nets from a Functional Standpoint." *Biological Cybernetics* 50: 35–41.
Kok, Peter, and Floris P. de Lange. 2015. "Predictive Coding in Sensory Cortex." In *An Introduction to Model-Based Cognitive Neuroscience*, edited by Birte U. Forstmann and Eric-Jan Wagenmakers, 221–44. New York, NY: Springer New York. https://doi.org/10.1007/978-1-4939-2236-9_11.

Komatsu, Hidehiko. 2006. "The Neural Mechanisms of Perceptual Filling-In." *Nature Reviews. Neuroscience* 7 (3): 220–31. https://doi.org/10.1038/nrn1869.

Kornmeier, Jürgen, and Michael Bach. 2004. "Early Neural Activity in Necker-Cube Reversal: Evidence for Low-Level Processing of a Gestalt Phenomenon." *Psychophysiology* 41 (1): 1–8. https://doi.org/10.1046/j.1469-8986.2003.00126.x.

———. 2012. "Ambiguous Figures – What Happens in the Brain When Perception Changes But Not the Stimulus." *Frontiers in Human Neuroscience* 6. https://doi.org/10.3389/fnhum.2012.00051.

Kosslyn, Stephen M., Nathaniel M. Alpert, William L. Thompson, Vera Maljkovic, Steven B. Weise, Christopher F. Chabris, Sania E. Hamilton, Scott L. Rauch, and Ferdinando S. Buonanno. 1993. "Visual Mental Imagery Activates Topographically Organized Visual Cortex: PET Investigations." *Journal of Cognitive Neuroscience* 5 (3): 263–87. https://doi.org/10.1162/jocn.1993.5.3.263.

Kosslyn, Stephen M., and William L. Thompson. 2003. "When Is Early Visual Cortex Activated during Visual Mental Imagery?" *Psychological Bulletin* 129 (5): 723–46. https://doi.org/10.1037/0033-2909.129.5.723.

Kwon, Oh-Sang, and David C. Knill. 2013. "The Brain Uses Adaptive Internal Models of Scene Statistics for Sensorimotor Estimation and Planning." *Proceedings of the National Academy of Sciences* 110 (11): E1064–73. https://doi.org/10.1073/pnas.1214869110.

Lalazar, Hagai, and Eilon Vaadia. 2008. "Neural Basis of Sensorimotor Learning: Modifying Internal Models." *Current Opinion in Neurobiology* 18 (6): 573–81. https://doi.org/10.1016/j.conb.2008.11.003.

Lamme, Victor A.F., and Pieter R. Roelfsema. 2000. "The Distinct Modes of Vision Offered by Feedforward and Recurrent Processing." *Trends in Neurosciences* 23 (11): 571–79. https://doi.org/10.1016/S0166-2236(00)01657-X.

Land, M. F. 2014. "Do We Have an Internal Model of the Outside World?" *Philosophical Transactions of the Royal Society B: Biological Sciences* 369 (1636): 20130045–20130045. https://doi.org/10.1098/rstb.2013.0045.

Lange, Richard D., and Ralf M. Haefner. 2016. "Inferring the Brain's Internal Model from Sensory Responses in a Probabilistic Inference Framework." *BioRxiv*, December, 081661. https://doi.org/10.1101/081661.

Layton, Oliver W., Ennio Mingolla, and Arash Yazdanbakhsh. 2014. "Neural Dynamics of Feedforward and Feedback Processing in Figure-Ground Segregation." *Frontiers in Psychology* 5 (September). https://doi.org/10.3389/fpsyg.2014.00972.

Lee, Tai Sing. 2002. “Top-down Influence in Early Visual Processing: A Bayesian Perspective.” *Physiology & Behavior* 77 (4): 645–50. https://doi.org/10.1016/S0031-9384(02)00903-4.
———. 2016. “The Visual System’s Internal Model of the World.” *Proc IEEE Inst Electr Electron Eng. 2015* August; 103(8): 1359–1378.: 43.
Lee, Tai Sing, and David Mumford. 2003. “Hierarchical Bayesian Inference in the Visual Cortex.” *Journal of the Optical Society of America A* 20 (7): 1434. https://doi.org/10.1364/JOSAA.20.001434.
Lee, Tai Sing, David Mumford, Richard Romero, and Victor A.F. Lamme. 1998. “The Role of the Primary Visual Cortex in Higher Level Vision.” *Vision Research* 38 (15–16): 2429–54. https://doi.org/10.1016/S0042-6989(97)00464-1.
Lee, Tai Sing, and My Nguyen. 2001. “Dynamics of Subjective Contour Formation in the Early Visual Cortex.” *Proceedings of the National Academy of Sciences* 98 (4): 1907–11. https://doi.org/10.1073/pnas.98.4.1907.
Marek, Scott, Joshua S. Siegel, Evan M. Gordon, Ryan V. Raut, Caterina Gratton, Dillan J. Newbold, Mario Ortega, et al. 2018. “Spatial and Temporal Organization of the Individual Human Cerebellum.” *Neuron* 100 (4): 977-993.e7. https://doi.org/10.1016/j.neuron.2018.10.010.
McIntyre, J., M. Zago, A. Berthoz, and F. Lacquaniti. 2001. “Does the Brain Model Newton’s Laws?” *Nature Neuroscience* 4 (7): 693–94. https://doi.org/10.1038/89477.
Muckli, Lars. n.d. “What Are We Missing Here? Brain Imaging Evidence for Higher Cognitive Functions in Primary Visual Cortex V1.” *International Journal of Imaging Systems and Technology* 20 (2): 131–39. https://doi.org/10.1002/ima.20236.
Naselaris, Thomas, Cheryl A. Olman, Dustin E. Stansbury, Kamil Ugurbil, and Jack L. Gallant. 2015. “A Voxel-Wise Encoding Model for Early Visual Areas Decodes Mental Images of Remembered Scenes.” *NeuroImage* 105 (January): 215–28. https://doi.org/10.1016/j.neuroimage.2014.10.018.
Nir, Yuval, and Giulio Tononi. 2010. “Dreaming and the Brain: From Phenomenology to Neurophysiology.” *Trends in Cognitive Sciences* 14 (2): 88. https://doi.org/10.1016/j.tics.2009.12.001.
O’Callaghan, Claire, Kestutis Kveraga, James M. Shine, Reginald B. Adams, and Moshe Bar. 2016. “Convergent Evidence for Top-down Effects from the ‘Predictive Brain’ 1.” *Behavioral and Brain Sciences* 39. https://doi.org/10.1017/S0140525X15002599.
O’Herron, Philip, and Rüdiger von der Heydt. 2009. “Short-Term Memory for Figure-Ground Organization in the Visual Cortex.” *Neuron* 61 (5): 801–9. https://doi.org/10.1016/j.neuron.2009.01.014.

Parkkonen, Lauri, Jesper Andersson, Matti Hämäläinen, and Riitta Hari. 2008. "Early Visual Brain Areas Reflect the Percept of an Ambiguous Scene." *Proceedings of the National Academy of Sciences* 105 (51): 20500–504. https://doi.org/10.1073/pnas.0810966105.

Penfield, Wilder, and Theodore Rasmussen. 1950. *The Cerebral Cortex of Man; a Clinical Study of Localization of Function*. The Cerebral Cortex of Man; a Clinical Study of Localization of Function. Oxford, England: Macmillan.

Pessoa, L., E. Thompson, and A. Noë. 1998. "Finding out about Filling-in: A Guide to Perceptual Completion for Visual Science and the Philosophy of Perception." *The Behavioral and Brain Sciences* 21 (6): 723–48; discussion 748-802.

Pessoa, Luiz, and Peter De Weerd. 2003. *Filling-in: From Perceptual Completion to Cortical Reorganization*. Oxford University Press.

Petro, Lucy S., and Lars Muckli. 2016. "The Brain's Predictive Prowess Revealed in Primary Visual Cortex." *Proceedings of the National Academy of Sciences* 113 (5): 1124–25. https://doi.org/10.1073/pnas.1523834113.

Pickering, Martin J., and Andy Clark. 2014. "Getting Ahead: Forward Models and Their Place in Cognitive Architecture." *Trends in Cognitive Sciences* 18 (9): 451–56. https://doi.org/10.1016/j.tics.2014.05.006.

Picton, Terence W, and Donald T. Stuss. 1994. "Neurobiology of Conscious Experience." *Current Opinion in Neurobiology* 4 (2): 256–65. https://doi.org/10.1016/0959-4388(94)90082-5.

Pitt, David. 2020. "Mental Representation -SEP." Edited by Edward N Zalta, Uri Nodelman, Colin Allen, and R Lanier Anderson. *Stanford Encyclopedia of Philosophy*, 23.

Raftopoulos, Athanassios, and Vincent C. Muller. 2006. "The Phenomenal Content of Experience." *Mind <html_ent Glyph="@amp;" Ascii="&"/ > Language* 21 (2): 187–219. https://doi.org/10.1111/j.0268-1064.2006.00311.x.

Raisch, J\"org, and Thomas Moor. 2005. "Hierarchical Hybrid Control of a Multiproduct Batch Plant" 322: 99–216.

Rao, Rajesh P.N. 1999. "An Optimal Estimation Approach to Visual Perception and Learning." *Vision Research* 39 (11): 1963–89. https://doi.org/10.1016/S0042-6989(98)00279-X.

Rees, Geraint, and Rimona weil. 2009. "HowDoestheBrainFillintheVisualWorld.Pdf." *ACNR* 9 (4).

Revonuso, A. 2006. *Inner Presence*. The MIT Press. https://mitpress.mit.edu/books/inner-presence.

Sá Pereira, Roberto Horácio de. 2016. "Meaning Representationalism: Between Representationalism and Qualia Realism." *Grazer Philosophische Studien* 93 (4): 548–70. https://doi.org/10.1163/18756735-09302005.

Schmahmann, Jeremy D. 2010. "The Role of the Cerebellum in Cognition and Emotion: Personal Reflections Since 1982 on the Dysmetria of Thought Hypothesis, and Its Historical Evolution from Theory to Therapy." *Neuropsychology Review* 20 (3): 236–60. https://doi.org/10.1007/s11065-010-9142-x.

Scholte, H. Steven, Jacob Jolij, Johannes J. Fahrenfort, and Victor A. F. Lamme. 2008. "Feedforward and Recurrent Processing in Scene Segmentation: Electroencephalography and Functional Magnetic Resonance Imaging." Research-article. Http://Dx.Doi.Org/10.1162/Jocn.2008.20142. October 14, 2008. https://doi.org/10.1162/jocn.2008.20142.

Schulman, Ari N. 2017. "Is Perception a 'Controlled Hallucination'?" *BQO* (blog). September 6, 2017. https://www.bigquestionsonline.com/2017/09/06/perception-controlled-hallucination/.

Seghier, M., M. Dojat, C. Delon-Martin, C. Rubin, J. Warnking, C. Segebarth, and J. Bullier. 2000. "Moving Illusory Contours Activate Primary Visual Cortex: An FMRI Study." *Cerebral Cortex* 10 (7): 663–70. https://doi.org/10.1093/cercor/10.7.663.

Seidler, R. D, D. C Noll, and G Thiers. 2004. "Feedforward and Feedback Processes in Motor Control." *NeuroImage* 22 (4): 1775–83. https://doi.org/10.1016/j.neuroimage.2004.05.003.

Seth, Anil. n.d. "How Does Your Brain Construct Your Conscious Reality?" NPR.Org. Accessed December 7, 2018. https://www.npr.org/2016/07/15/654730916/how-does-your-brain-construct-your-conscious-reality.

Shadmehr, Reza, and Henry H. Holcomb. 1997. "Neural Correlates of Motor Memory Consolidation." *Science* 277 (5327): 821–25. https://doi.org/10.1126/science.277.5327.821.

Siegel, Susanna. 2016. "The Contents of Perception." In *Stanford Encyclopedia of Philosophy*, edited by Edward N Zalta, Uri Nodelman, Colin Allen, and R Lanier Anderson, Winter 2016, 44.

Smith, Fraser W., and Lars Muckli. 2010. "Nonstimulated Early Visual Areas Carry Information about Surrounding Context." *Proceedings of the National Academy of Sciences* 107 (46): 20099–103. https://doi.org/10.1073/pnas.1000233107.

Somers, David C., Anders M. Dale, Adriane E. Seiffert, and Roger B. H. Tootell. 1999. "Functional MRI Reveals Spatially Specific Attentional Modulation in Human Primary Visual Cortex." *Proceedings of the National Academy of Sciences* 96 (4): 1663–68. https://doi.org/10.1073/pnas.96.4.1663.

Spillmann, Lothar, and Peter De Weerd. 2003. *Mechanisms of Surface Completion: Perceptual Filling-In of Texture*. Oxford University Press. http://

www.oxfordscholarship.com/view/10.1093/acprof:oso/9780195140132.001.0001/acprof-9780195140132-chapter-5.
Stokes, Mark, Russell Thompson, Rhodri Cusack, and John Duncan. 2009. "Top-Down Activation of Shape-Specific Population Codes in Visual Cortex during Mental Imagery." *Journal of Neuroscience* 29 (5): 1565–72. https://doi.org/10.1523/JNEUROSCI.4657-08.2009.
Sülzenbrück, Sandra, and Herbert Heuer. 2011. "Type of Visual Feedback during Practice Influences the Precision of the Acquired Internal Model of a Complex Visuo-Motor Transformation." *Ergonomics* 54 (1): 34–46. https://doi.org/10.1080/00140139.2010.535023.
Synofzik, Matthis, Axel Lindner, and Peter Thier. 2008. "The Cerebellum Updates Predictions about the Visual Consequences of One's Behavior." *Current Biology* 18 (11): 814–18. https://doi.org/10.1016/j.cub.2008.04.071.
"The Underestimated Cerebellum Gains New Respect From Brain Scientists." n.d. NPR.Org. Accessed July 12, 2020. https://www.npr.org/sections/health-shots/2018/10/25/660504533/the-underestimated-cerebellum-gains-new-respect-from-brain-scientists.
Thoroughman, Kurt A., and Reza Shadmehr. 2000. "Learning of Action through Adaptive Combination of Motor Primitives." *Nature* 407 (6805): 742–47. https://doi.org/10.1038/35037588.
Toet, J., A.; Blom, J.; Koenderink. 1987. "The Construction of a Simultaneous Functional Order in Nervous Systems. III. The Influence of Environmental Constraints on the Resulting Functional Order." *Biol. Cybernet.* 57: 331–40.
———. 1988. "The Construction of a Simultaneous Functional Order in Nervous Systems. IV. The Influence of Physical Constraints on the Resulting Functional Order." *Biol. Cybernet.* 58: 275–86.
Toet, J. J., A.; Blom, J.; Koenderink. 1987. "The Construction of a Simultaneous Functional Order in Nervous Systems. II. Computing Geometrical Structures." *Biol. Cybernet.* 57: 127–36.
Tye, Michael. 1995. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind.* The MIT Press.
———. 2017. "Qualia_Stanford_Encyclopedia of Philosophy." https://plato.stanford.edu/entries/qualia/.
Venn, John. 1866. *The Logic of Chance 1st Ed.* London, Cambridge, Macmillan. https://ia802702.us.archive.org/33/items/logicchanceanes01venngoog/logicchanceanes01venngoog.pdf.
Vetter, Petra, and Albert Newen. 2014. "Varieties of Cognitive Penetration in Visual Perception." *Consciousness and Cognition* 27: 62–75. https://doi.org/10.1016/j.concog.2014.04.007.

Wang, M., D. Arteaga, and B. J. He. 2013. "Brain Mechanisms for Simple Perception and Bistable Perception." *Proceedings of the National Academy of Sciences* 110 (35): E3350–59. https://doi.org/10.1073/pnas.1221945110.
Wikipedia contributors. 2021. "Information Flow." In *Wikipedia*. https://en.wikipedia.org/w/index.php?title=Information_flow_(information_theory)&oldid=1000165672.
Williams, Mark A., Chris I. Baker, Hans P. Op de Beeck, Won Mok Shim, Sabin Dang, Christina Triantafyllou, and Nancy Kanwisher. 2008. "Feedback of Visual Object Information to Foveal Retinotopic Cortex." *Nature Neuroscience* 11 (12): 1439–45. https://doi.org/10.1038/nn.2218.
Wokke, Martijn E., Annelinde R. E. Vandenbroucke, H. Steven Scholte, and Victor A. F. Lamme. 2013. "Confuse Your Illusion: Feedback to Early Visual Cortex Contributes to Perceptual Completion." *Psychological Science* 24 (1): 63–71. https://doi.org/10.1177/0956797612449175.
Wolpert, D. M., and M. Kawato. 1998. "Multiple Paired Forward and Inverse Models for Motor Control." *Neural Networks* 11 (7): 1317–29. https://doi.org/10.1016/S0893-6080(98)00066-5.
Wolpert, Daniel M, Zoubin Ghahramani, and J. Randall Flanagan. 2001. "Perspectives and Problems in Motor Learning." *Trends in Cognitive Sciences* 5 (11): 487–94. https://doi.org/10.1016/S1364-6613(00)01773-3.
Wolpert, Daniel M, R. Chris Miall, and Mitsuo Kawato. 1998. "Internal Models in the Cerebellum." *Trends in Cognitive Sciences* 2 (9): 338–47. https://doi.org/10.1016/S1364-6613(98)01221-2.
Young, Malcolm. 2000. "The Architecture of Visual Cortex and Inferential Processes in Vision." *Spatial Vision* 13 (2): 137–46. https://doi.org/10.1163/156856800741162.
Zago, Myrka, and Francesco Lacquaniti. 2005. "Visual Perception and Interception of Falling Objects: A Review of Evidence for an Internal Model of Gravity." *Journal of Neural Engineering* 2 (3): S198. https://doi.org/10.1088/1741-2560/2/3/S04.
Zalta, Edward N. 2001. "Fregean Senses, Modes of Presentation, and Concepts." *Nous* 35 (s15): 335–59. https://doi.org/10.1111/0029-4624.35.s15.15.
Zeimbekis, John, and Athanassios Raftopoulos. 2015. *The Cognitive Penetrability of Perception: New Philosophical Perspectives*.
Zeki, S., and A. Bartels. 1999. "Toward a Theory of Visual Consciousness." *Consciousness and Cognition* 8 (2): 225–59. https://doi.org/10.1006/ccog.1999.0390.
Zeki, Semir. 1993. *A Vision of the Brain*. A Vision of the Brain. Cambridge, MA, US: Blackwell Scientific Publications.