The Standard Theory of Conscious Perception

Carolyn Dicey Jennings (cjennings3@ucmerced.edu)
Cognitive and Information Sciences, 5200 N. Lake Road
Merced, CA 95343 USA

Abstract

In this paper I argue that the prioritization of sensory input by top-down attention is constitutive of and essential to conscious perception. Specifically, I argue that top-down attention is required to provide informational integration at the level of the subject, which can be contrasted with integration at the level of features and objects. Since the informational content of conscious perception requires integration at the level of the subject, top-down attention is necessary for conscious perception as we know it. I present this argument through a theory, which I call the "Standard Theory." According to this theory, top-down attention brings about subject-level integration for sensory input by prioritizing that input with respect to a subject-level standard.

Keywords: top-down attention; conscious perception; Tononi; integrated information; standard theory

Introduction

In the cognitive sciences, top-down attention is understood to prioritize select neural processing according to the subject's current task. Although this function of top-down attention may not appear at first to play an important role in the bringing about of conscious perception, I will argue here that such prioritization is constitutive of and essential to the structure of conscious perception, expanding on arguments published elsewhere (Jennings, 2015). I present this argument through a theory, which I call the "Standard Theory." According to this theory, top-down attention integrates sensory input by prioritizing this input with respect to a subject-level[1] standard, using a spatiotemporal framework common to all types of sensory input to do so.

The Standard Theory is original in its details but not in its vision. William James connects attention to conscious perception in his *Principles of Psychology*.

> Millions of items of the outward order are present to my senses which never properly enter into my experience. Why? Because they have no interest for me. *My experience is what I agree to attend to.* Only those items which I notice shape my mind – without selective interest, experience is an utter chaos. Interest alone gives accent and emphasis, light and shade, background and foreground – intelligible perspective, in a word. It varies in every creature, but without it the consciousness of every creature would be a gray chaotic indiscriminateness, impossible for us even to conceive. (James, 1981, p. 403)

James claims here that it is interest on the part of the subject that determines the content of perception, both in quantity and in quality. His claim rests on evidence of sensory selection ("Millions of items...never properly enter into my experience") together with the idea that such selection must be governed by interests. Of course, quantitative selection need not be achieved by a subject, since it can take place through bottom-up filtering and selection. Thus, the fact that we do not perceive all of the "items" surrounding us does not by itself show us that interest-based attention is the gateway to conscious perception. Nonetheless, James' intuition that interest-based attention provides for the qualitative content of conscious perception finds support in the Standard Theory.

Maurice Merleau-Ponty, on the other hand, entertains the notion that attention is the key to the common spatiotemporal framework of conscious perception. In the Introduction to *Phenomenology of Perception*, he makes the following observation.

> Attention first of all presupposes a transformation of the mental field, a new way for consciousness to be present to its objects. Take the act of attention whereby I locate a point on my body which is being touched...A *vaguely located spot*, this contradictory phenomenon reveals a pre-objective space where there is indeed extension, since several points on the body touched together are not confused by the subject, but as yet no univocal position, because no spatial framework persists from one perception to another. The first operation of attention is, then, to create for itself a *field*, either perceptual or mental, which can be 'surveyed' (Merleau-Ponty, 1962, pp. 33-34, translation amended)

Here, Merleau-Ponty claims that although prior to the act of attention experience is organized (e.g. two unattended points on the body are easily distinguished), attention allows for a new type of organization. Namely, attention transforms conscious experience from a pre-objective space to an objective space by invoking a common spatiotemporal framework.

Merleau-Ponty argues for this role of attention by looking at the phenomenon of directed search. Directed search implies that the subject has a partly determinate and partly indeterminate idea of its target (Merleau-Ponty, 1962, p. 33). When I search for a pen on the surface of my desk, for

[1] "Subject-level" processing is taken to be the subset of cognitive processing that is open to modification by "the subject"—that to which we normally ascribe conscious experience, thought, and high-level behavior, but which is not necessarily equivalent to "the organism."

example, I use the determinate idea that the pen is somewhere on my desk to guide my search and the indeterminacy of its precise location to drive the search. That determinacy can guide the search reveals that the target has conceptual qualities, and is not a pure sensory object (contrary to the "empiricist" view, as Merleau-Ponty construes it). Yet, that indeterminacy drives the search reveals that the target has qualities that extend beyond our current knowledge (contrary to the "rationalist" view, as Merleau-Ponty construes it). Attention, as the mechanism of search, will have to be the sort of thing that can connect the realm of structured thought to the realm of pure sensory objects to reveal a target of this sort, according to Merleau-Ponty.

The Standard Theory gives a working account of these phenomenological intuitions. Namely, the Standard Theory shows how top-down attention transforms sensory input into conscious percepts by prioritizing the former with respect to a 'subject-level standard'–a standard rooted in or held by the subject. According to the Standard Theory, this transformation brings about both the informational content of perception, as flagged by James, and the determination of a shared perceptual space, as flagged by Merleau-Ponty. Thus, the Standard Theory is an account of how the subject brings about the structure of conscious perception through the activity of top-down attention.

The Problem of Subject-Level Integration

To start, it may be helpful to clarify my working definitions of top-down attention and conscious perception. As I use it, "top-down attention" should be understood as the prioritization of select mental or neural processing according to the subject's current interests. This can be contrasted with prioritization that takes place *solely* through filtering and bottom-up selection.[2] "Conscious perception," on the other hand, normally refers to the experience of informational sensory content within a spatiotemporal framework. Although sensory input may include some spatiotemporal information, it is thought to differ from conscious perception in lacking informational content. That is, unlike sensory input, conscious perception presents the world as being a certain way to the subject. This "being a certain way" involves both general or shared attributes between the many items of content and particular instantiations of those attributes held by each item. The stargazer, for example, experiences the stars as each having relative brightness: they share the perceptual quality of brightness, but each of them instantiates a particular degree of brightness. As Tyler Burge puts it: "a perception–a representational perceptual state instance, or the content of a perceptual state instance–must always involve the context-dependent singular application of (general) perceptual attributives" (Burge, 2010, p. 381). Perhaps we could provisionally characterize conscious perception as the experience of sensory elements bound to a spatiotemporal structure, where "elements" signifies the particularities or instantiations of generalities, and where this spatiotemporal structure may include localized events, patterns, or simply spatiotemporal depth.[3]

A full characterization of such informational content has been undertaken by Giulio Tononi as part of his integrated information theory of consciousness[4]:

> Every time we experience a particular conscious state out of such a huge repertoire of possible conscious states, we gain access to a correspondingly large amount of information. This conclusion is in line with the classical definition of information as reduction of uncertainty among a number of alternatives (Shannon and Weaver, 1949)...the information generated by the occurrence of a particular conscious state lies in the large number of different conscious states that *could potentially* have been experienced but were not. (Tononi, 2005, p. 111)

So the informational content of one's current conscious experience occurs through contrast with what is not experienced. Tononi motivates this theory by describing the visual experience of total darkness: such an experience only has visual content if "content" measures what is present against what is possible. That is, a room devoid of light has visual content only when compared with other potential visual experiences one could have. Normal visual sensation, in contrast, relies on the presence of light. The fact that we experience total darkness, then, shows us that conscious content is inherently informational. In Tononi's view, even the conscious perception of a homogenous plane of light has a lot of information for us, since its content is generated in contrast to other potential experiences.

Tononi further notes that the information of conscious content is *integrated* in a way that is not captured by "Shannon Information."[5] As captured in the paragraph above, we might say of the content of perceptual experience that it is informational *for us*. Tononi suggests that

> To measure information integration, it is essential to know whether a system of elements constitute a causally integrated system, or they can be broken down into a number of independent or quasi-independent subsets among which no information can be integrated. (Tononi, 2005, p. 113)[6]

[2] A more complete account of this form of attention can be found in Jennings, 2012.

[3] Although I take it that these remarks hold for all varieties of conscious perception, I focus here on visual perception.

[4] While Tononi emphasizes the informational content of consciousness, in general, I focus only on the informational content of conscious perception, in particular.

[5] "Shannon Information" was put forward as a partial account of information—the data without its meaning (Shannon, 1948; Floridi, 2009).

[6] Note that Tononi's use of the term "element" does not necessarily include participation in a complex as a defining feature.

That is, Tononi claims that to get anything like conscious content we require not only information but integrated information, which we can measure by looking at causal integration.

But integration with respect to what? Tononi says that integration is implicated by the fact that we cannot experience shape without color (Tononi, 2005, p. 112). He also says that integration is implicated when a "conscious state is experienced as an integrated whole" (Tononi, 2005, p. 112). Yet these forms of integration do not seem extensive enough to capture both actual and potential experience, which is central to his account—one might have integrated perceptual features and even an entirely integrated experience without achieving integration across actual and potential experience. Without this level of integration the conscious experience of a completely dark room would not be able to yield information through contrast with other potential experiences, for instance. I thus suggest that Tononi's account of conscious content is best understood as information that is integrated with respect to the subject (a view that sometimes appears to be supported by Tononi). Although I adopt Tononi's basic perspective in this paper, note that I am not arguing that all of conscious experience is captured by integrated information, but only that all of our perceptual experiences can be understood this way.[7]

This understanding of conscious perception (as the experience of integrated information within a spatio-temporal framework) raises a problem. Specifically, assuming conscious perception is rooted in neural processing and that this requires structural correspondence between the two, neural processing must involve a part-whole structure that can support the experience of a unified complex of sensory elements. The problem arises when we note that the early processing of sensory input contributes only half of this requisite structure: the early processing of sensory input divides that input into feature-specific processing without it yet having membership in a unified complex of processing. The problem of discovering how this unified complex comes about is known as the 'problem of unity,' which has now been solved at the level of features and objects, but not yet at the level of the subject. Importantly, although recognition of this problem assumes the possibility of a natural explanation of conscious perception, it need not confuse the content of conscious perception with its vehicles. That is, one can agree with Susan Hurley that "the properties of subpersonal processes, of vehicles of content, cannot simply be projected into personal-level mental content, or vice versa" (Hurley, 2002, p. 3) and still find particular vehicles wanting. As stated, it is the presumption of structural correspondence, and not a simple projection of properties, that drives the claim that content-vehicles will have to instantiate a part-whole relationship that corresponds to the division and unity of integrated information. I assume structural correspondence because I take this to be the minimal constraint on the "is rooted in" relation.

Burge, mentioned above, claims that the development of perceptual content occurs through perceptual constancies (Burge, 2010, pp. 407-413), but I look further back to the source of these constancies and find a role for the subject through top-down attention. As I will argue below, a solution to problem of integration at the level of the subject will necessarily involve a subject-level standard, which can only be applied through top-down attention. The solution to this problem lies, in other words, in the adoption of the Standard Theory.

The Standard Theory

From the work of Treisman and Gelade (1980); Treisman (1988); Wolfe, Cave, and Franzel (1989); and others we have an account of how objects are bound from the features of separable feature sets (e.g. color and luminance), which answers the problem of unity at the level of objects. Yet, if we want to fill in the account of perceptual unity then we need an explanation of the unity that exists across objects and feature sets, which allows for integration at the level of the subject. The Standard Theory provides this missing explanation. The basic claim of the Standard Theory is that top-down attention provides for conscious perception by transforming sensory input into conscious percepts. Top-down attention achieves this by prioritizing (and thus differentiating) the sensory input according to a subject-level standard, through which the sensory-input-turned-percepts are "unified" or integrated at the level of the subject.

To start, let's reflect on conscious perception at the level of objects and feature sets to examine the claim that perceptual experience has a level of integration that goes beyond objects and feature sets. If you look across the space in front of you, you will probably note many specific shades of color (e.g. the green of a plant, the red of a couch). Each of these shades of color should be easily attributed to its object, since color is normally integrated with other features of an object. Yet, these objects share the feature of having color. Thus, these shades of color are also integrated at the level of the feature set, which exists across visual objects. Similarly, the sounds around you each have a specific pitch (e.g. the rumble of a car, the hum of a fan), easily attributed to separable objects due to integration at the level of objects, while nonetheless sharing the quality of having pitch, due to integration of pitch across aural objects.

The above reflections reveal the existence of integration at the level of both objects and feature sets. To see the integration that occurs beyond objects and feature sets, note the difference within experience that occurs when one switches focus from one object to another (e.g. from the plant to the couch, or from the car to the fan): the current feature or object of interest gains a sort of proximity, whereas other features and objects now seem more distant. This is not a change in acuity or spatial distance–the change occurs even if one keeps one's eyes still. It is instead

[7] See Jennings (in press) for those forms of conscious experience that depart from this model.

experienced as an epistemic or valuational proximity: the feature or object of interest appears more available for knowledge, more valuable, or more meaningful, whereas other features or objects appear less available for knowledge, less valuable, or less meaningful. However exactly one experiences the difference between the proximal and distal objects, this structural difference reveals integration that transcends objects and feature sets. Since this integration changes with the subject's interests, it qualifies as subject-level integration (see Jennings, 2015 for a more detailed account).

That perceptual experience has this structure was pointed out by Gurwitsch in his *Field of Consciousness* (Gurwitsch, 1964). Gurwitsch claims that perceptual experience always contains a "theme" and a "thematic field," where the separation between these can be as rough-grained as that between two different sensory modalities. In "ganzfeld" experiences, for example, where a single sensory modality contains only homogenous information (as of a homogenous plane of light), that modality may yet experienced as a theme relative to other sense modalities (Wackermann et al., 2008). Even in Balint's syndrome, where a patient with bilateral parietal lesions only recognizes one object at a time, there is a contextual basis for the object that can serve as its thematic field (Michel & Henaff, 2004, p.11).

How is this structure achieved? A spatiotemporal framework common to objects and feature sets clearly plays a role. The green of the plant and the red of the couch are not confused, in part, because these objects and features are separated in space and time (as is central to Treisman's account). Yet, a spatiotemporal framework common to objects and feature sets cannot provide the relevant standard of comparison for integration at the level of the subject. What can provide this standard? A clear common factor in all perceptual experience is the experiencing subject. Since the determination of subjective value depends upon the subject, it makes sense that subject-level integration, according to the subject's current interests and values, would be brought about by the subject.

This is where the Standard Theory comes in. According to the Standard Theory, conscious perception requires the integration of early sensory processing by the subject according to a subject-level standard through top-down attention. It is based on a relatively simple argument:

1) An essential feature of conscious perception is that it is informational for a subject

2) Early sensory processing is not integrated in a way that could allow for this type of information without further processing

3) Some process must bring about this integration for early sensory processing to bring about conscious perception

4) In order to bring about this integration, early sensory processing will have to be differentiated according to a subject-level standard

5) Only top-down attention differentiates sensory processing according to subject-level standards

6) Thus, top-down attention is necessary for conscious perception.

Further evidence for this argument comes from observing the variation between perceivers. For two perceivers watching someone bike past, one may perceive the color of the bike but not that of the cyclist's clothing, while the other perceives the color of the cyclist's clothing but not that of the bike. The most natural explanation of this difference in what is perceived is that the perceivers have different interests. That is, the interests of the subject in each case determine what is and is not consciously perceived.

I find it useful to think of the maintenance of the boundary between what is consciously perceived and what is not consciously perceived as being a form of consumption by the subject. That is, just as the boundary between what digestion and waste is internally regulated, the boundary between perception and sensation is internally regulated. When one digests peanut butter, the peanut butter is (eventually) transformed into parts of the body. How is this accomplished? It is not simply a process of the peanut butter being taken in by the body; bodies have considerable variation in what they will digest, and this variation is partly regulated by the digestive system itself. Just as the digestive system (partly) determines what will be incorporated into the body, the claim here is that the subject (partly) determines what will be integrated into the subject-level.

If these observations are correct, the integration of early sensory processing into perceptual experience requires action by the subject. Yet, sensory processing must also be differentiated in order to be integrated as information. This, I claim, is the role of prioritization. Importantly, integration and differentiation are twin concepts. That is, for parts to be integrated with respect to some standard is for the parts to be differentiated, and for the parts to be differentiated with respect to some standard is for them to be integrated. Thus, the role of attention in differentiating and the role of the subject in integrating are linked: the subject is the standard of integration with respect to which attention differentiates.

To see the link between informational differentiation and integration, consider that for something to count as information it must be differentiated with respect to some standard. The number "59" is not information, even though it can be differentiated from other numbers, unless it is set against some standard. The phrases "My father is aged 59 years" or "I saved 59 dollars on my wedding dress," do count as informational because they supply a context and an implied standard. For something to be informational for a subject, the subject must either hold the standard or the standard must be rooted in the subject. That is, the differentiated parts will have to point back to the subject according to the subject-level standard. Otherwise the information will not be integrated or unified with respect to the subject. Top-down attention is a process of prioritization by the subject or according to the subject's current interests. Thus, for the subject to differentiate those items it integrates, it will have to use top-down attention. This doesn't mean, of course, that bottom-up attention cannot

also play a role, so long as top-down attention is involved to some degree.

The most basic form of integrated information, in my view, can be found in someone that perceives only in terms of positive and negative valence. Let's imagine such a person named "Michéle," who likes sitting on cool beaches (in Scotland, say) but detests sitting on hot beaches (in Florida, say). When presented with a particular beach in Scotland, Michéle might have a perceptual experience of positive valence at a particular intensity, shape, and duration, which she calls "West Sands." It may be that "West Sands" for Michéle applies to beaches other than what humans call "West Sands" so long as those beaches produce just the same experience of positive valence. When presented with a particular beach in Florida, Michéle might have a perceptual experience of negative valence at a particular intensity, shape, and duration, which she calls "Naples Bay."[8] It may be that "Naples Bay" (for Michéle) applies to other beaches, or even to other things that Michéle detests, so long as those things produce just the same experience of negative valence. Michéle will not have a perceptual experience of anything that does not have a ranking of positive or negative valence for her. Michéle may have no feelings at all towards umbrellas, for example, in which case she would not perceive them. For Michéle to perceive something, that thing must be ranked according to her subject-level standard. Otherwise that thing would not be differentially integrated with respect to Michéle (qua subject).

My claim is that all perceptual experience must be based on a subject-level standard akin to Michéle's beach interests, by virtue of which it has subject-level integration. The story so far supplies only the most universal and basic form of integrated information. The conscious perception of most humans has layers of differentiation and integration beyond a single subject-level standard. It is the integration of these further layers, I claim, that requires a common spatiotemporal framework.

Returning to an example used at the start of this section, there is unity of color in one's visual field. How might this occur? My suggestion is that it occurs through the pre-subjective prioritization of sensory input according to habitual tasks that reliably fit the subject-level standard. That is, for stable subject-environment couplings it may be the case that the processing of sensory input becomes tuned to goals that can reliably bring about value according to the subject-level standard. Groups of neurons may become tuned, for example, to the inputs that normally enable the realization of subjective value and respond preferentially to them. Thus, feature maps, or the mapping of inputs to pre-subject-level standards that regularly contribute to the subject-level standard, will take place without top-down attention for habitual or long-term goals. This explains how it is that certain features of stimuli are pre-subjectively prioritized by the visual cortex – the inheritance of neural tuning created by millennia of other subjects matching goals like ours to a world like ours does much of the dividing work for us. Thus, one need not find top-down attention necessary for each instance of integration, such as the integration of individual feature sets, but only for integration with the subject's current interests.

To achieve this level of integration, across objects and feature sets, it will be necessary to have a common framework for comparison. That is, the application of a unifying standard to multiple sets of features, for example, will require a structure through which to compare the values of each set without losing the internal structure of each set. This is the role of the common spatiotemporal framework found in nearly every perceptual experience. A common spatiotemporal framework is required so that the prioritizations found in the feature maps can be compared with respect to the subject-level standard without losing the particular information contained within those feature maps. This common spatiotemporal framework need only be general enough to combine all the feature sets, and not so general as absolute or allocentric space-time. In such a system the different feature dimensions will share a framework or matrix that allows for prioritization across dimensions but that sacrifices neither the bottom-up salience values contained in each dimension nor the spatiotemporal location values of the information sources. This spatiotemporal framework helps us to solve the problem of subject-level integration in that it provides a mechanism of comparison for the application of an integrating standard. Yet, even once we have a framework for comparison, we will still need top-down attention to match the values from pre-subject-level standards to the subject-level standard to obtain subject-level integration.

According to the Standard Theory, all perceptual experience minimally contains the structure yielded by prioritization with respect to a subject-level standard, even if it does not also contain other layers of prioritization that need to be fitted to a common spatiotemporal framework, as illustrated in the case of Michéle. Thus, although the working definition of conscious perception supplied at the start of this paper involved the experience of sensory elements bound to spatiotemporal structure, the Standard Theory does not depend on this limited understanding of conscious perception to make its case. Any type of perceptual content, in being integrated with respect to the subject, will rely on top-down attention, according to the Standard Theory.

Conclusion

Throughout this paper I have illustrated how the content of conscious perception relies upon top-down attention. I have claimed that although the integration of individual feature-sets can be automatized, such that neural structures can become tuned to particular types of features and collections of features according to the habitual or long-term goals of the subject, integration across features and objects will only be achieved with the real-time application of a subject-level standard through top-down attention. Thus, if the Standard

[8] For the record, I love Naples Bay.

Theory is correct, top-down attention is necessary for conscious perception as we know it, which minimally involves the division of theme from thematic field, but also normally involves deeper layers of differentiation and integration.

.

References

Burge, T. (2010). Origins of Objectivity. Oxford University Press.

Floridi, L. (2009). Philosophical conceptions of information. In G. Sommaruga (Ed.), Formal Theories of Information. Springer-Verlag.

Gurwitsch, A. (1964). Field of Consciousness. Duquesne University Press.

Hurley, S. (2002). Consciousness in Action. Harvard University Press.

James, W. (1981). The Principles of Psychology. Harvard University Press.

Jennings, C. D. (2012). The subject of attention. Synthese, 189(3), 535–554.

Jennings, C. D. (2015). Attention and perceptual organization. Philosophical Studies, 172(5), 1265-1278.

Jennings, C. D. (in press). Consciousness without attention. Journal of the American Philosophical Association.

Merleau-Ponty, M. (1962). Phenomenology of Perception. Routledge. (Trans. C. Smith)

Michel, F., & Henaff, M. (2004). Seeing without the occipitoparietal cortex: Simultagnosia as a shrinkage of the attentional visual field. Behavioral Neurology, 15.

Shannon, C. E. (1948). A mathematical theory of communication. Bell System Technical Journal, 27, 379–423.

Tononi, G. (2005). Consciousness, information integration, and the brain. In S. Laureys (Ed.), Progress in Brain Research. Elsevier.

Treisman, A. M. (1988). Features and objects: The Fourteenth Bartlett Memorial Lecture. Journal of Experimental Psychology, 40(2), 201–237.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. Cognitive Psychology, 12(1), 97–136.

Wackermann, J., Putz, P., & Allefeld, C. (2008). Ganzfeld-induced hallucinatory experience, its phenomenology and cerebral electrophysiology. Cortex, 44, 1364–1378.

Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. Journal of Experimental Psychology, 15(3), 419–433.