Please use DOI when citing or quoting. DOI: 10.1086/705525

Perception, Representation, Realism and Function

Alison Ann Springle

University of Pittsburgh

Department of Philosophy

alisonspringle@gmail.com

Abstract:

According to orthodox representationalism, perceptual states have constitutive veridicality or accuracy conditions. In defense of this view, several philosophers, but most notably Burge (2010), employ a realist strategy that turns on the purported explanatory ineliminability of representational posits in perceptual science. I argue that Burge's version of the realist strategy fails as a defense of orthodox representationalism. However, it may vindicate a different kind of representationalism.

According to orthodox representationalism, perceptual states, and not just post-perceptual states (e.g. judgement, belief), are representational in the sense of having constitutive truth, veridicality, or accuracy conditions. In defense of orthodox representationalism, a number of philosophers appeal to the broadly Marrian (computational) explanatory framework for visual perception (e.g. Burge 2010, Rescorla 2015, Hill 2014, Schellenberg 2018, Matthen 2005).

However, the grounds for this appeal are nowhere more clearly articulated than in Burge (2010). According to Burge, realism about some of the framework's representational posits is justified because those posits are explanatorily ineliminable, i.e. eliminating posits bearing the constitutive features of representation would result in explanatory loss.

Burge's argument has generated many responses, but many of these responses, sympathetic and critical alike, neglect the heart of his argument.

Rescorla (2015), for instance, attempts to buttress Burge by appealing to the representational posits of specifically Bayesian perceptual psychology, the explanatory success of this paradigm, and its mathematical rigor. But Rescorla's arguments do nothing to buttress a key feature of Burge's grounds for realism, namely, the empirical grounds for treating perceptual systems as having a representational function. In fact, Burge denies that perceptual systems make "inferences," as a Bayesian model would have it, in any non-metaphorical sense, and he's clear that the mere fact that a science uses the term "representation," "veridicality," etc., doesn't show that the science is committed to a *robust* notion of representation. (27) Indeed, Burge thinks that much of what sensory and perceptual scientists call "representations" are merely "deflationary" representations, i.e. the products of systems that are responsible for "information registration," where this information may play a functional (typically biological) role, but where we can account for the activity of these systems and the states they produce without invoking veridicality or accuracy conditions.[1] To wit, deflationary sensory representations are sensory

[1] See also Orlandi (2014).

states which sciences may refer to as representations, but for which accuracy conditions are explanatorily *eliminable*. According to Burge, only where a sensory system has a representational function are accuracy conditions explanatorily ineliminable. So only where ineliminable representational posits also ineliminably function representationally does science provide grounds for realism about (robust) perceptual representation.

As I'll explain in (§1), the mere success of the computational (Marrian, Bayesian) framework is not sufficient evidence of ineliminable representational function because that framework *presupposes* that perception functions to veridically represent, and there are competing frameworks that can also boast of success that claim that perception instead has an action-guiding function. I'll also explain why we cannot, and why Burge does not, conclude *a priori* that perception's action-guiding function entails its having the function to veridically represent.

According to Burge, representational function is only ineliminable for explanations of sensory systems that exhibit the perceptual constancies. Many philosophers have discussed (e.g. Vincente 2012, Aguilera 2016) and some have also criticized (e.g. Ganson, Bronner, and Kerr 2014; Olin 2016) Burge's conception of the perceptual constancies in connection with his defense of representationalism without recognizing the connection he draws between the constancies and what he calls "objectifying capacities," and between objectifying capacities and representational function. As I'll explain in (§2), the constancies enter Burge's picture as empirical evidence of the existence of sensory systems that function to (descriptively, veridically) represent, for Burge thinks that the constancies constitute objectifying capacities.

In (§3) I'll argue that the constancies probably aren't objectifying capacities, and *ipso facto* probably aren't empirical evidence of sensory systems that function to veridically represent. However, in (§4) I'll suggest that Burge's realist strategy may support a form of representationalism that exchanges constitutive veridicality conditions for constitutive appropriateness conditions.

§1.

As Frisby and Stone note, computational theories of vision begin with the question "what is vision for?" and the answer to this question "has shifted over time, reflecting the changing emphasis within the research field of vision" (2010, 545) According to David Marr, when perceptual processing runs correctly, it "delivers a true description of what is there" (1982, 30) and "the true heart of visual perception is the inference from the structure of an image about the structure of the real world outside" (68). Similarly, Knill et. al assert that: "Visual perception … involves the evolution of an organism's visual system to match the structure of the world and the coding scheme provided by nature" (1996, 6). However, over the last two decades, the Marrian approach has been challenged by a different paradigm in computer vision: Active Vision, so called "because its key idea is to build seeing systems geared to specific sorts of visually guided actions"(Frisby and Stone 2010, 545).

The Active approach to computer vision has developed systems for, among other things, automatic surveillance. A computer visual system might be designed to guide a video camera rig to track a potential suspect with "the visual processing built to do that having no idea of what the tracked object is." This serves as an example of a vision module designed to perform a certain action (camera control) "without doing anything else (such as recognizing the tracked object)" (Frisby and Stone, 545). Hence, it appears that in perceptual action-control and perceptual tracking, it is possible to bypass intermediate descriptive representations (e.g. object identification). Active Vision may provide a better model of biological vision than traditional models. At the very least, the ideas of Active Vision have guided a vigorous program of research in biological vision (Findlay and Gilchrist 2003).

So where Marr answers the question “what is vision for” with “to know what is where by looking,” the active vision paradigm answers: “for guiding actions.” Of course, these answers are not obviously in conflict: perception might function to veridically represent in order to achieve its action-guiding function (see e.g. Palmer 1999, 6; see also Marr 1982, 340). But as Burge knows, this isn’t a given.

Burge cites Kathleen Akins’s (1996) empirically driven arguments against the "traditional view of the senses" as a major impetus to his (2010, xviii). On the traditional view, "sensory systems must be veridical in some sense of the word," that is, the senses ought to report the "what, when, and where" of the world's events” and "provide an accurate account of just how things are: the brain must be able to tell, from the signals it receives, how things stand in the world." Akins’s case study is thermoreception. If thermoreception has the function of providing accurate information about the world, then, *ceteris paribus*, the sensory response that we think tracks temperature ought to be produced (except for in cases of malfunction, which should presumably be rare) only in the presence of temperature stimuli. And if the temperature rises, the response should reflect this, i.e. if the temperature rises only slightly, the sensory system should *ipso facto* register only a slight increase. It should not systematically exaggerate or embellish on the information it’s receiving from the world. Rather, it should strive to be objectively accurate.

As it turns out, however, this isn’t the way human thermoreception works at all. Rather than striving to accurately report on the world, Akins argues that most sensory systems are concerned with how the organism is being affected. She dubs such systems “narcissistic.” And when it comes to navigating the world successfully, narcissism probably has a selective advantage over "veridical" measuring systems, for their "narcissism" makes them fit to serve the demands of what Akins calls the "*sensory-motor project*": "the pre-ontological narcissistic encoding of the information required to act and get around in the world." The upshot of Akins's argument is that there's a gap between the demands of this project and those of what she terms the "*ontological project*": "to conceive of types and tokens, places, and objects as existing at all given our sensory access to the world" (369).

The gap between these projects is mirrored in Burge's distinction between biological functions, which are essentially practical, and representational functions, which, he thinks, have essentially to do with veridicality. Burge's distinction rests on the dissociability of these functions, i.e. their conditions of success. A system can fail representationally by being quite inaccurate with respect to the environment, and that inaccuracy itself may serve biological success. On the flip-side, a system might succeed representationally by being extremely accurate, and this could be a biological shortcoming of the system. So, there's a gap and potentially a tension between biological success and representational success.

This gap underwrites Burge's key objection to teleosemantic accounts of representation, as these attempt to ground representation in biological function.[2] Burge claims that there's "a root mismatch between representational error and error of biological function" (2010, 301). The reason is that, as Stitch (1990, 62) observed, "natural selection doesn't care about truth; it cares about reproductive success," or, as Burge puts it "Evolution does not care about veridicality. It does not select for veridicality *per se*" (303). Indeed, though Burge makes no mention of such work, several perceptual psychologists have recently employed evolutionary game theory to investigate the likelihood of evolving objectifying capacities-- i.e. of descriptive function being

[2] This issue is omitted in discussions of Burge's criticisms of bio-functionalism. See e.g. Vincente (2012), Aguilera (2013).

selected as a means to successfully guiding action. Their overwhelming verdict: it's extremely unlikely (e.g. Hoffman 2009, Jager 2007, O'Connonr 2014).

According to Burge, veridicality conditions for a putative content analyzed in terms of biological function *simpliciter* are explanatorily eliminable; the explanatory work they do is "redundant" (2010, 9). Veridicality conditions merely provide an intentional gloss on the explanatorily essential posits: natural signs (information carriers) which are not, in themselves, representational, and biological functions which do not, essentially, have anything to do with veridicality, and thus cannot, *on their own,* confer descriptive norms and ipso facto explanatorily essential veridicality conditions. Appealing to information and biological function *alone* secures only deflationary representation, for nothing about biology guarantees that accuracy should be the goal of an evolved system.

Why then should we think that perceptual systems function to (descriptively) represent? Burge's answer: the constancies. [3]

§2

Burge's (2010) aims to provide a de-intellectualized account of the conditions that must be met for an individual to achieve objective representation as of particulars in the environment. Burge rejects what he calls "Compensatory Individual Representationalism" (CIR), according to which perception delivers something short of objective representation, i.e., subjective representation; in order to achieve objective representation, the *individual* must supplement what perception gives her. Against CIR, Burge argues that there is no sensory representation short of objective representation. Genuine representation requires explanatorily ineliminable veridicality

[3] Gilchrest (2006, 126) notes that "machine vision's emphasis on veridicality resonated with the traditional themes of constancy in psychology, bringing this issue back to the foreground."

conditions, and the only sensory states for which veridicality conditions are ineliminable are those that issue from sensory systems that we must understand as having descriptive function. Which systems are those? Perceptual systems. In contrast to mere sensory registration (mere sensation), perception "requires perceptual constancies" (399).

Key for Burge is the relationship between the constancies and objectification. According to Burge, the constancies constitute the "central instances of perceptual objectification" (397). Since genuine representation begins with objectification, objective representation isn't the result of sensory representations that are supplemented by the individual. For sensory outputs that are not objective must be nonrepresentational, and hence mere sensations.

Objectification is realized via the constancy mechanisms that perform transformations on proximal stimulation that, purely perceptually (non-conceptually, i.e. independent of judgement), achieve the stable perception of actual macro-level physical properties of distal stimuli. Burge understands these mechanisms as capacities for "separating out" information about the creature and how it's being affected from information about the world, and, critically, as capacities for getting the objective world right: they determine, via internalized physical constraints, what physical categories (attributes) to posit upon the receipt of various patterns of subjective proximal simulation. They thereby track *objective* properties of the distal cause. Burge thinks that perceptual science has *discovered* these objectifying capacities and that the representational function of perceptual sensory systems are manifest in scientific descriptions of constancy mechanisms, hence his assertion (e.g. xi-xii) that perceptual representations are natural kinds.

But the case for considering the perceptual constancies "objectifying capacities" is weak.

§3.

Hatfield (2009) critiques the view that perceptual systems aim at full constancy, i.e. to represent an object with the properties that it has--its veridical or objective size, shape, position, and color. One of his targets is the physicalist construal which is "most clearly present in formulations of perceptual constancies as seeking to solve an inverse problem: that of inferring back from the proximal stimulus to its physical cause" where the aim of constancy is to "produce a mental representation of the physical properties of surfaces" (185). In other words, Hatfield targets Burge's construal of the constancies.

Perceptual reports differ systematically depending on the way the experimenter explains the perceptual task to subjects, and/or on what questions they are asked. The earlier literature didn't explicitly control for such factors, resulting in ambiguity in the data. V.R. Carlson's (1977) review of the literature corrected for this by distinguishing "projective" (projective shape), "apparent" (how the shape looks, sans cognitive correction), and "objective" (the real shape they believe the object to have) instructions.

The results were surprising: projective instructions yield close to projective results, apparent instructions yield underconstancy or perfect constancy, and objective instructions yield constancy or overconstancy (size is overestimated; shape is overcorrected). The results were also robust: investigators from Robert Thouless (1931, 1932) onwards (Shape: Epistein et al. 1977; Size: Epistein 1963; Granrud 2004, 79; Wagner 2006, 130) have observed underconstancy in the perception of size and shape: reported phenomenal size is near to but less than real size, and reported phenomenal shape is nearer to real shape than projected shape-- even in full cue conditions.

Hatfield suggests that part of the reason underconstancy seems surprising is that scientific and philosophical investigators have failed to appreciate the phenomenological implications of full constancy. For instance, "…with full constancy, objects of the same size should appear with exactly the same size despite being at different distances" and "...we should see a hallway with no phenomenal convergence whatsoever" (188-9). Of course, this isn't the way the world spatially appears. Rather, Hatfield hypothesizes that the entire visual space contracts with distance, yielding diminished phenomenal size with increasing distance even under full cue conditions (ibid, 198) This contracted space also affects phenomenally presented shape (ibid, 203).

Hatfield's hypothesis was initially studied by Franz Hillebrand (1902). Hillebrand found that in order to make black cords on a white table top appear parallel, and to adjust a suspended curtain of threads to produce phenomenally parallel walls, subjects set the chords and curtain threads to diverge as they ran near to far. Walter Blumenfeld (1913) subsequently conducted an experiment in which subjects had to adjust individual lights in an otherwise dark alley so that they formed either phenomenally parallel lines from far to near or equidistant horizontal intervals at various distances. In reduced-cue conditions, Blumenfeld found that the parallel appearing lights diverged even more than had Hillebrand's cords. From their findings, both Hillebrand and Blumenfeld inferred phenomenal contraction.

Neither Hillebrand nor Blumenfeld took their findings to support a non-Euclidean geometry for visual space. But in the 1940s, Rudolf K. Luneburg reinterpreted their results in accordance with non-Euclidean geometry, proposing that visual space is a hyperbolic space of constant negative curvature. This was in part because Luneburg (1947, 39) held that only if visual space is of constant curvature can it meet the conditions usually associated with physical possibility: that objects can freely rotate and move in space without altering their shapes, permitting true size constancy. Tarow Indow's (1962, 1982) investigations contributed additional data consistent with Luneburg's proposed non-Euclidean geometry for visual space. However, while Indow (1991) agreed with Luneburg about the desirability of constant curvature, he concluded from further data that planes perpendicular to the line of sight are non-Euclidean in structure, so that any curvature of visual space would be in the depth dimension to the viewer (Hatfield 2009, p 198-99).

Similarly, Mark Wagner's (2006, ch. 7) extensive empirical investigations of the global structure of visual space lead him to conclude that it's a general fact about human visual experience that it contracts with respect to visual space in the depth dimension, i.e. in the direction running away from the observer. Wagner cast further doubt on Luneburg's suggestion that visual space meets the axioms that define a constant physical space, including free mobility of objects without distortion by reviewing a literature showing deviations from Euclidean structure (e.g., Koenderink et al. 2000) while producing no single constant metric, and showing dynamic effects on judgment of size and orientation that depend on the presence or absence of stimuli besides the target object (Thorndyke 1981; Schoumans et al. 2002; Indow 1991, 450). Wagner concluded that there's no single metric for visual space: "no single geometry can fully encompass human visual experience" (2006, 223; see also Suppes 1977). Rather, multiple metrics may arise depending on the visual task, and variable metrics may apply to a single phenomenal experience.

Likewise, because Hatfield's proposed contracted space does not exhibit free mobility, it does not meet the usual conditions on a physically possible space. Hatfield consequently denies the physicalist assumption that perception aims to phenomenally present physical sizes and shapes "as they are."[4]

Why then do perceiver's reports often respect full constancy under objective instructions? Hatfield proposes that contracted spatial experience enables subjects to recognize or judge the real, objective physical shapes and sizes with some accuracy, and that the products of such judgements or acts of recognition are phenomenologically present (the phenomenological presence of the conceptual) (ibid, 190). Critically, Hatfield holds that our ability to respond to true, objective, or physical sizes and shapes arises from *judgmental or recognitional abilities* that employ concepts and reasoning. Though phenomenally integrated into occurrent experience, these abilities are psychologically distinct from visual appearances themselves (187).

Hatfield's explanation takes a cue from Granrud (2004). Granrud argues that apparent size diminishes with distance (beyond 3 meters) for both adults and children, but between the ages of 5 and 10, children *learn to reason* about the effects of distance on apparent size. Though they show slight underconstancy at near distances (5 or 6.1 meters) and exhibit underconstancy for far distances when asked to report on "apparent size," when asked to report "objective size", they exhibit full or overconstancy. But notably, children exhibit no difference between apparent and objective instructions when they did not show that they could reason about the effects of distance. To explain this data, Granrud proposes that "adults and children have the same perceptual experience" (78) of apparent size diminishing wish distance, while "far-distance size constancy results from a cognitive judgement and is not a feature of perception" (89). The view that the constancies require cognitive contributions is consistent with a number of empirical

[4] Both Hatfield and Wagner Wagner (2006, 177) suggest that successful action guidance is compatible with contracted visual space. Hatfield even suggests that contracted visual space may be selectively preferred for successful spatial behavior.

studies on human perceptual constancies that assume cognitive contributions. In other words, there's strong reason to think that spatial constancies qua objectifying are not capacities of perception proper, and there's even more reason to be skeptical in the case of color constancies (see e.g. Chirimuuta 2008, 2015; Hatfield 2009, 2003).

Explanations of the constancies in simpler creatures that are unlikely to rely on judgement, on the other hand, don't give them the appearance of objectifying capacities. Hutto and Myin (2012) make this case on the basis of Neander's (2006) detailed reportage on the state of the art science concerning the brain and behavior of anurans (frogs and toads), noting that Burge "definitively places frogs (as well as fish and octopi) on the list of perceivers. All these creatures unquestionably make the cut as full-blooded (if cold-blooded) perceivers because they can perceive basic constancies (Burge 2010, 420)" (119). Anurans also exhibit what Burge considers the most reliable indicator that any individual is an objectifying perceiver: the capacity to localize—i.e. determine the direction and distance—of a "distal source of stimuli without serial sampling" (Burge 2010, 427). Frogs localize both by orienting toward or away from interesting stimuli, and by targeting and snapping at prey.

Neander concludes that the working assumptions of the relevant sciences do not include a commitment to anuran perceptual representation in any obvious or robust way. She reports that neuroethologists "rarely if ever utter sentences beginning with 'the content of a +T5(2) is . . .' or '+T5(2)s mean . . .', or '+T5(2)s [activated neuron of a certain relevant sort] represent. . . ." Indeed, Neander states that she must "concede to those who dislike talk of content in the case of such simple systems that there is little explicit talk of error concerning what is represented" (ibid, 183). Following Barlow (1953), Chirimuuta (2015, 110) draws a similar conclusion about the nature of the frog's "bug detector."

Treating the constancies not as objectifying capacities but as capacities for robust responsiveness appears to be in line with actual scientific practice, while Burge's interpretation appears unparsimonious. As Hutto and Myin put the point: "The perceptual sciences, it seems, would not collapse if it turned out that true perceivers fall into the class of aspectual *respondents* but *not that of attributive claimants*… perception doesn't depend on, or entail, the existence of attributive states of mind" (121).

One might object: But we can't understand the organization and activity of sensory regions in an animal's brain without considering the parts of the environment with which its activity covaries. No doubt this covariation is a critical part of the way the brain executes its perceptual functions. But this doesn't entail robust representation: information is not robust representation (Ramsey 2007, ch. 4), and causal covariation is ubiquitous and non-normative. Indeed, while teleosemanticists consider causal covariation "information" and take it to provide at least some critical part of representation, information, covariation, or indication is representation only when it has the right function. If Burge is right that this must be the function to veridically represent, and if I'm right that scientific explanation does not entitle us to attribute this function to perceptual systems, Burge's realist strategy for defending representationalism is a no-go.

But this may be too quick.

§4.

According to Neander, while neuroethologists rarely invoke standards of veridicality, they *do* invoke standards of "appropriateness " (184). Hutto and Myin rightly claim that this isn't grounds for orthodox representationalism to celebrate. But what about a different sort of representationalism?

Elsewhere, I develop and defend an account of what I call "instructive representationalism" (Springle, Unpublished Manuscript). According to instructive representationalism, perceptual experience is constituted by the application of instructive categories to or on situations, where the application of an instructive category constitutes an instructive representation. Instructive categories constitute a kind of mode of presentation that is distinct from the more familiar descriptive mode of presentation. The latter are third-personal modes of presentation that present things (including actions) as having some property (qualitative, quantitative, causal, metaphysical, or normative), and the descriptive representations their applications constitute are evaluable as true/false (or veridical/falsidical; accurate/inaccurate).

Instructive categories, in contrast, are first-personal practical modes of presentation that instructively stand-for tokens of types of actions perceivers might produce in response to situations so as to promote their flourishing, i.e. to meet their needs. Actions are presented *as* means, but not as *being* means, i.e. not as objects (events) independent of the perceiver and characterized by a practical nature. Instructive categories relate perceivers directly to actions that are in their power to produce; perceptual experiences present perceivers with their own relational potentialities-- their own possible performances in relation to environmental particulars. Instructive representations function not to describe the world but to position perceivers to take the appropriate means with respect to the world, depending on what their actual needs are. Instructions are successful not by virtue of the actions they instruct being manifested, but by virtue of being appropriate with respect to the situations on which they are occasioned. An instruction is appropriate iff were the perceiver to execute the instructed action with respect to the situation, doing so would satisfy the need with which the action-type, qua means, is constitutively connected.

Instructive representationalism is a version of what Mazviita Chirimuuta (2015) calls "Perceptual Pragmatism": it's an alternative to philosophical views that assume that perception succeeds when it accurately describes the world according to which perception functions to "help you live by guiding your activity in the world" (110). Just as perceptual pragmatism is a "philosophical articulation of the minimal perceptual epistemology that is in the background of the empirical research" which "… gets to a core commitment of perceptual science" (ibid, 110), instructive representationalism, unlike orthodox representationalism, posits no functions beyond the practical, action-guiding functions to which perceptual science is firmly committed.

References

Aguilera, B. (2016) “Is Perception Representational? Tyler Burge on Perceptual Functions.” In *An Anthology of Philosophical Studies*, Vol.7, P. Hanna (Ed.), ATINER, Grecia, pp. 59-71.

Akins, K. (1996) “Of Sensory Systems and the Aboutness of Mental States” *Journal of Philosophy* Vol. 93 (7) pp. 337—372.

Barlow, H. B. (1953) “Summation and Inhibition in the Frog’s Retina.” *Journal of Physiology* Vol. 119 (1) pp. 69–88.

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

Carlson, R. (1977) “Instructions and Perceptual Constancy Judgements.” In William Epstein (ed.) *Stability and Constancy in Visual Perception: Mechanisms and Processes,* New York: Wiley. pp. 217-54.

Chirimuuta, M. (2015) O*utside Color: Perceptual Science and the Puzzle of Color in Philosophy*. MIT Press.

Chirimuuta, M. (2008) “Reflectance Realism and Colour Constancy: What would count as scientific evidence for Hilbert’s ontology of colour?” *Australasian Journal of Philosophy* Vol. 86(4) pp. 563-582

Dretske, F. (1986). “Misrepresentation.” In Bogan, R. (ed.). *Belief: Form, Content, and Function.* New York: Oxford University Press.

Epistein, W., Hatfield, G., Muise, G. (1977) “Perceived Shape at a Slant as a Function of Processing Time and Processing Load.” *Journal of Experimental Psychology: Human Perception and Performance,* Vol. 3 pp. 473-83.

Epistein, W., Bontrager, H., and Park, J. (1963) “The induction of non-veridical slant and the perception of shape.” *Journal of Experimental Psychology,* Vol. 63 pp. 472-9.

Findlay, J. and Gilchrist, I. (2003) A*ctive Vision: The Psychology of Looking and Seeing.* Oxford University Press, NY.

Frisby, J and Stone, J. (2010) *Seeing: The Computational Approach to Vision.* 2nd Ed. MIT Press.

Ganson, T. Bronner,G. & Kerr, A. (2014) “Burge’s Defense of Perceptual Content.” *Philosophy and Phenomenological Research.* Vol. 88 (3) pp. 556-73

Gilchrist A.L. (2006). *Seeing Black & White*. Oxford University Press.

Granrud, C. (2004) “Visual metacognition and the development of size constancy.” In Daniel T. Levin Ed., *Thinking and Seeing: Visual Metacognition in Adults and Children*, Cambridge Mass.: MIT Press, pp. 75-95.

Hatfield, G. (2009) *Perception and Cognition: Essays in the Philosophy of Psychology*. Oxford: Oxford University Press.

-------. (2003) “Objectivity and Subjectivity Revisited: Color as a Psychobiological Property.” In *Colour Perception: Mind and the Physical World*, ed. By Rainer Mausfeld and Dieter Heyer. Oxford: Oxford University Press. 187-202.

Hillebrand, F. (1902) “Theorie der cheinbaren Grösse beim binokularen Sehen.” *Denskschrift der Kaiserlichen Akademie der Wissenschaften Wien, Mathematische-Naturwissenschaftliche Classe* Vol. 72 pp. 255-307.

Hill, C. (2014) *Meaning, Mind, and Knowledge.* Oxford: Oxford University Press.

Hoffman, D. (2009) “The Interface Theory of Perception: Natural Selection Drives True Perception to Swift Extinction.” *In Object Categorization: Computer and Human Vision Perspectives,* S. Dickinson, M. Tarr, A. Leonardis, B. Schiele (Eds.) Cambridge,UK: Cambridge University Press, 148–165.

Hutto, D.D. & Myin, E. (2012) *Radicalizing Enactivism: Basic Minds Without Content*, MIT Press.

Indow, T., Inoue, E., Matsuchima, K. (1962) “An Experimental Study of the Luneburg Theory of Binocular Space (2) The Alley Experiments.” *Japanese Psychological Research* Vol. 4 pp. 17-24

Indow, T. (1991) “A Critical Review of Luneburg’s Model with Regard to Global Structure of Visual Space.” *Psychological Review* Vol. 98 pp. 430-53

Jager, G. (2007) “The Evolution of Convex Categories.” *Linguistics and Philosophy* Vol. 30 (5) pp.551-564

Koenderink, J., van Doorn, A., Lappin J. (2000) “Direct Measurement of the Curvature of Visual Space.” *Perception* Vol. 29 pp. 69-79.

Luneburg R. (1947) *Mathematical Analysis of Binocular Vision.* Princeton, Princeton University Press.

Marr, D. 1982. V*ision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: Freeman.

Matthen, M. (2005) *Seeing, Doing, & Knowing A Philosophical Theory of Sense Perception*. Oxford University Press: Oxford.

Neander, K. (2006) “Content for Cognitive Science.” In *Teleosemantics*, ed. G. Macdonald and D. Papineau (eds).Oxford University Press.

O’Connor, C. (2014) “Evolving Perceptual Categories.” *Philosophy of Science* Vol. 81(5), pp. 840–851.

Olin, L. (2016) “Burge on Perception and Sensation” *Synthese* Vol. 193 pp. 1479-1508.

Orlandi, N. (2014) *The Innocent Eye: Why Vision is not a Cognitive Process.* New York: Oxford University Press.

Ramsey, W. M. (2007) *Representation Reconsidered.* Cambridge University Press.

Rescorla, M. (2013) “Bayesian Perceptual Psychology” in T*he Oxford Handbook of the Philosophy of Perception*. Mohan Matthen (ed) Oxford University Press.

Schellenberg, S. (2018) *The Unity of Perception: Content, Consciousness, and Evidence*. New York: Oxford University Press.

Schoumans, N., Kappers, A., Koenderink, J. (2002) “Scale Invariance in Near Space: Pointing Under the Influence of Context.” *Acta Psychologica* Vol. 110 pp. 63-81.

Springle, A. (Unpublished Manuscript) “Instructive Representationalism.”

Stich, S. (1990) *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: MIT Press.

Suppes, P. (1977) “Is Visual Space Euclidean?” *Synthese* Vol. 35 pp. 397-421.

Thorndyke, P. (1981) “Distance Estimation from Cognitive Maps.” *Cognitive Psychology* Vol. 13 pp. 526-50.

Thouless, R. (1931) “Phenomenal Regression to the Real Object.” *British Journal of Psychology* Vol. 21 pp. 339-56.

Thouless, R. (1932) “Phenomenal Regression to the “Real” Object. II.” *British Journal of Psychology* Vol. 22 pp. 1-30.

Vincente, A. (2012) “Burge on Representation and Biological Function” *Thought* (1) pp. 125-133.

Wagner, M. (2006) *Geometries of Visual Space*. Mahwah, N.J. Erlbaum.