Published in final edited form as:

Perception. 2012 ; 41(9): 1053–1060.

Does this computational theory solve the right problem? Marr, Gibson, and the goal of vision

William H. Warren

Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, Providence, RI, USA

William H. Warren: Bill_Warren@brown.edu

Abstract

David Marr's (1982) book *Vision* attempted to formulate a thoroughgoing formal theory of perception. Marr borrowed much of the "computational" level from James Gibson: a proper understanding of the goal of vision, the natural constraints, and the available information is prerequisite to describing the processes and mechanisms by which the goal is achieved. Yet as a research program leading to a computational model of human vision, Marr's program did not succeed. This article asks why, using the perception of 3D shape as a morality tale. Marr presumed that the goal of vision is to recover a general-purpose Euclidean description of the world, which can be deployed for any task or action. On this formulation, vision is underdetermined by information, which in turn necessitates auxiliary assumptions to solve the problem. But Marr's assumptions did not actually reflect natural constraints, and consequently the solutions were not robust. We now know that humans do not in fact recover Euclidean structure – rather, they reliably perceive qualitative shape (hills, dales, courses, ridges), which *is* specified by the 2nd-order differential structure of images. By recasting the goals of vision in terms our perceptual competencies, and doing the hard work of analyzing the information available under ecological constraints, we can reformulate the problem so that perception is determined by information and prior knowledge is unnecessary.

David Marr's (1982) book *Vision* was one of the first attempts to formulate a thoroughgoing theory of perception, by synthesizing current ideas from computer vision with what was known from psychophysical experiments and visual neurophysiology. Marr went beyond previous phenomenological analyses and descriptive models in an effort to create explicit, formal theories that were experimentally testable. As he wrote of Ullman's (1978) thesis on structure from motion, "Its importance is that it enables us to formulate a number of empirical questions that would not otherwise arise, and it opens the way for a rational investigation of the phenomenon rather than the confused cataloguing of its phenomenology" (p. 199).

This was a great step forward. Yet as a research program leading to a computational model of human vision, Marr's program did not succeed. The question I wish to pursue is why.

Marr's vision

It is an instructive exercise to re-read Marr 30 years later. I had just finished my Ph.D. when the book first appeared, and a great buzz ran through the vision community. As a newly minted Gibsonian, I had a deep skepticism about the presumed underdetermination of vision and the consequences of reifying representations and computations in the head, as opposed to recognizing them as descriptive conveniences – and still do (Warren, 2005). Yet upon re-reading the book, I was repeatedly struck by how much there is to like. Although, as one

might expect, most of the specific algorithmic proposals have proven inadequate, many of Marr's general points ring true three decades later:

- The goal of perception is “to reliably derive properties of the world from images of it” (p. 25) – in particular to recover properties of visible surfaces (p. 269).
- Perceptual theory must be based on natural constraints, in particular “the ways information is constrained by the world” (p. 104), which permit the recovery of those surface properties.
- The level of computational theory is fundamental to understanding vision (p. 27–28), including a characterization of *what* function or task is to be performed and *why* it's appropriate: because it realizes the goal within the given constraints (p. 22–23).
- A precise description of *how* a visual process proceeds is required if one is ever going to make sense of its neural implementation (pp. 23, 27). This is the motivation for introducing the algorithmic level of description.

The book develops a strong anti-reductionist view of explanation. For Marr, a neural mechanism isn't a theory of vision. Rather, one needs to understand why a mechanism would work: what information it's detecting, the natural constraints under which that variable is informative about the world, and the process by which the surface property is recovered. In contrast to the common wisdom, the constraints running from the computational level to the algorithmic and implementation levels are stronger than those running in the reverse direction. The computational task limits the subset of algorithms for a specific process of information extraction, which in turn provides the job description for a subset of neural realizations; conversely, the physical substrate only constrains the general style of algorithm (e.g. serial or parallel). Consequently, the nature of perception depends more on the computational problem being solved than on the particular neural implementation; indeed, given that there might be many different neural realizations, a general theory of vision cannot be found at the neural level. As Marr famously observed, “Trying to understand perception by studying only neurons is like trying to understand bird flight by studying only feathers: It just cannot be done” (p. 27; see also Ch. 7). With cognitive neuroscience in the ascendancy, it is a useful corrective to be reminded of this point.

One can clearly see Gibson's influence in this formulation of the problem of perception. Indeed, Marr credited him with coming “the nearest anyone came to the level of computational theory” (p. 29) – which Gibson might have called the functional level – and Marr endorsed the idea that the task of vision was “the detection of physical invariants” (p. 30). However, he contended that Gibson “seriously underestimated the complexity of the information-processing problems involved” (p. 29) because he was “misled by the apparent simplicity of the act of seeing” (p. 30). This is a common misreading of Gibson, in which the meaning of “direct” perception is interpreted as “simple” or “immediate,” as if there were no visual process involved (e.g. Ullman 1980); but direct perception means that the perceiver is aware of its environment, not an internal representation thereof (Warren 2005).

Gibson (1966) actually undertook a reanalysis of sensory mechanisms as active systems for detecting information, given the state of the art at the time, and nearly everyone might be accused of having underestimated the difficulty of extracting informative variables from raw light in 1966. Nevertheless it is true that Gibson's primary interest lay at the functional level, for he realized that identifying the information that specifies environmental properties under ecological constraints – *ecological optics* – was foundational for a scientific theory of perception. Marr borrowed this framework from Gibson: a proper understanding of the goal

of vision, the natural constraints, and the available information is prerequisite to describing the processes and mechanisms by which it is extracted. They part company when Marr makes his “conceptual leap” to assert that the visual process is one of forming an internal description or representation of the world, as opposed to one of detecting information to make contact with the world itself.

A distinct tone of triumphalism runs through the book, as though the AI Lab at MIT had solved most of these problems by 1979, and the computational/representational approach would explain biological vision in short order. The ambition of this effort was indeed impressive: Marr discusses solutions for edge extraction, grouping, the correspondence problem, the aperture problem, stereopsis, structure from motion, and shape from texture, shading, and contour, among many others. Ironically, the book appeared at the height of computational vision’s success, and despite its visionary influence, Marr’s program subsequently foundered. Most of the proposed solutions were neither computationally robust nor empirically adequate to biological vision. In the 1980’s enthusiasm shifted to connectionism, and in the 1990’s to Bayesian models, as the zeitgeist in vision. The question, then, is why Marr’s program for computational vision failed.

A familiar answer at the algorithmic and implementation levels is that vision is harder and more complex than the Artificial Intelligentsia thought. A process that takes one-third of the brain, with more descending than ascending connections, may be too complicated to reduce to a set of independent parts composed of a bottom-up sequence of operations. The processes carried out by the visual and motor systems increasingly seem nonlinear, recurrent, dynamic, context-dependent and (as Gibson 1966 intuited) modulated by “top-down” constraints such as attention, the task, or the intended action. In short, Marr’s principle of modular design may be wrong, and if that goes the algorithmic approach becomes unworkable.

A more technical answer is noise. Faced with the inherent noisiness of natural images and biological components, deterministic computational operations turned out to be quite brittle. The field of computer vision subsequently turned to statistical Bayesian models that built variability into their foundations, which in turn inspired biological vision theorists to adopt a Bayesian framework (Kersten and Yuille 2003; Knill and Richards 1996). A robust model of vision will have to take account of noisy processes.

I want to focus at the computational level, and here the answer comes in two parts: First, the goal of vision presumed by the computational approach was too general. And second, this required assumptions about the world that were too specific. As an object lesson, I am going to develop the case of the perception of 3D shape, borrowing freely from recent research on the subject (see also Lappin et al 2011; Todd 2004).

The goal of vision

Marr assumed that the goal of computation was to recover a general-purpose description of a Euclidean, Newtonian world, which could be deployed for any subsequent process or behavior. This is most evident in his conception of the 2½-D sketch. Shape was to be built up out of local measurements of depth and surface orientation, a straightforward Euclidean description. This characterization had its origins in the Gibson of 1950, who proposed that perceived surfaces are based on the perception of local distance and slant to the line of sight, together with their spatial gradients. Similarly, the perception of surface color for Marr was to be the recovery of its physical reflectance. Thus, the description of the world to be perceived was borrowed straight from classical geometry and physics, without reference to the biological functions of vision. This came to be known as the “inverse optics” approach, and it survives in current Bayesian vision models.

This way of defining the problem led directly to the refrain that vision is underdetermined – that there is insufficient information in the image to invert the process and recover a Euclidean description of the scene. This, in turn, necessitated auxiliary assumptions about the world to solve the problem.

However, there is little evidence supporting the presumption that human observers reliably recover Euclidean structure. For over half a century, research has steadily accumulated to show that when properly tested, perceptual judgments of metric Euclidean properties – including depth, length, surface slant, and surface curvature – are systematically distorted (Domini and Caudek 2003; Foley et al 2004; Gilinsky 1951; Gogel 1964; Koenderink et al 2001; Loomis et al 1992; Norman et al 2005; Norman et al 1996; Perotti et al 1998; Tittle et al 1995; Todd and Perotti 1999; Wagner 1985). This is true regardless of the source of shape information (binocular disparity, motion, texture, shading) and even under full-cue conditions. To a first approximation, the distortions are affine, related to a depth scaling along the line of sight (such that objects appear increasingly compressed in depth as viewing distance increases) and a shear in an orthogonal direction (such that objects may appear rotated about an axis perpendicular to the line of sight).

These results imply a surprising lack of size and shape constancy over variation in viewing conditions. Explicit judgments of depth or length can change by a factor of 2 or more with changes in viewing distance, viewpoint, illumination direction, between observers, and even between tasks for a given observer. A startling demonstration of this fact is provided by the dashed lines in the middle of the highway. If you pay close attention as you approach a typical dash, it will appear to expand from under a meter to the length of your car, about 3m (Shaffer, Maynor, & Roy, 2008). And yet, our normal perceptual experience is that the shapes and sizes of things are relatively stable. The world doesn't appear to stretch and shear as we move around or as viewing conditions change. Even though our explicit judgments of metric 3D properties are systematically biased, we hardly notice. The qualitative properties of shape that matter perceptually and behaviorally survive these distortions (Todd and Norman 2003).

Affine distortions are just the beginning. We do not appear to possess a depth map, even a compressed one, that is internally consistent, preserving the order of points in depth across a complex scene (Domini and Braunstein 1998; Koenderink and Doorn 1995; Svarverud et al 2012; Todd and Reichel 1989). For instance, judgments of the depth order of two points are reliable if they happen to lie on a continuous slope of a single surface, such as the same side of a hill. But if two points lie on opposite sides of a depth boundary such as a ridge or valley, or on different objects, performance is surprisingly inconsistent; only local ordinal structure is reliably perceived. Moreover, it seems unlikely that there is a single 2½-D sketch or common representation for all perceptual and motor tasks (Hibbard 2008; Norman et al 2006; Todd 2004; Warren 1998); it increasingly appears that different environmental properties may be extracted for the task or action at hand.

Taken together, there is considerable evidence that inverting the optics to recover a general-purpose representation of Euclidean structure is not the primary goal of human vision.

Assumptions

Presuming the goal of computing Euclidean shape forced computational theorists to adopt auxiliary assumptions about the world in order to render the problem tractable. Marr candidly stated this research strategy as follows:

“In each case the surface structure is strictly underdetermined from the information in images alone, and the secret of formulating the processes accurately lies in

discovering precisely what additional information can safely be assumed about the world that provides powerful enough constraints for the process to run.” (p. 265–6).

Common assumptions included unique correspondence matches, continuity or surface smoothness, surface rigidity, Lambertian reflectance, uniform texture size, and planar contour generators. Such assumptions were supposed to reflect the physical constraints of the real world, but typically they were merely mathematical conveniences that simplified the problem in order to make the process run. For all the invocation of natural constraints, there was little investigation of the natural world and its actual constraints (what Gibson 1979 called *ecological physics*), or analysis of the optical structures made available under those constraints (*ecological optics*). The assumed constraints were thus too specific, narrowing the problem to restricted mini-worlds so that the algorithm would fly. As a consequence, the resulting solutions were not robust and failed to generalize to more complex natural conditions.

In the prevailing Bayesian framework, a similar logical role is played by “priors”. Whereas the computational approach conveniently assumed natural constraints in order to make the algorithms run, the Bayesian approach often assumes prior probability distributions in order to make the inferences go through. In its most circular form, priors are even inferred from psychophysical data to enable the model to explain the data. To avoid such question-begging, empirical investigation is ultimately needed to characterize actual priors, just like it is to identify natural or ecological constraints. Interpreting vision as a process of Bayesian inference imputes knowledge of priors to the visual system; this stands in contrast to ecological constraints, which make available optical regularities to which the visual system may adapt. Marr himself had little enthusiasm for probabilistic approaches because he felt that statistical machinery was not a substitute for the analysis of constraints and information that provides a firm basis for computational theory (p. 201).

Marr’s theory did set an agenda for research on human perception. Many ensuing experiments consisted of violating his simplifying assumptions and finding nevertheless that human judgments are often just as good – or even better – in a complex world. A typical 17th-c Dutch still-life is a readily interpretable image of a 3D scene, full of computationally intractable properties. Luminance changes across the image are due not only to the local orientation of a depicted surface (shading), but also to local surface reflectance, cast shadows, specular highlights, interreflections, transparency, reflections and specularities on transparent surfaces, dust caustics, and so on. Yet despite – or thanks to – this complexity, we clearly see a 3D scene of objects with definite shapes, laid out in a certain configuration, made of specific materials.

Thus, paradoxically, human vision appears to be both worse (non-Euclidean) and better (more robust) than Marr’s algorithms. I believe these two observations are linked. If we relax the goal of vision and jettison inverse optics, this opens up the possibility that there is sufficient information in the light to specify the scene properties we *do* recover.

Reformulating the problem

Let’s return to our morality tale about 3D shape. When we left, our gentle vision theorist was struggling to accept the idea that Euclidean properties such as local depth, slant, and curvature are not accurately or reliably perceived. For the most part, these are not the surface properties that matter perceptually or behaviorally. And yet, perceivers clearly recover *some* aspects of shape that enable them to distinguish, identify, and interact with 3D objects. And in fact, there are many weaker geometric properties that remain invariant under affine distortions and could support the perception of stable shape.

By 1979, Gibson had disavowed the proto-2½-D sketch in favor of a more qualitative description of shape in terms of “convexities” and “concavities”:

“I had made the mistake of thinking that the experience of the layout of the environment could be compounded of all the optical slants of each piece of surface... Convexities and concavities are not made up of elementary impressions of slant but are instead unitary features of the layout.” (p. 166)

There is recent interest in just such qualitative surface features. It is possible to characterize any complex smooth shape, such as Michaelangelo’s muscular “Day” from the Medici tombs, as a configuration of surface patches such as hills, dales, saddles, courses, and ridges. These primitive shapes can be described by the *shape characteristic* or the *shape index*, both of which are based on the relationship between the principal curvatures ($\kappa_{\max}$, $\kappa_{\min}$) of a local surface patch (Koenderink 1990). In contrast to Euclidean properties, such qualitative features can be judged reliably and remain stable under variation in viewing conditions. For example, the shape characteristic of a surface patch can be judged much more reliably than its metric curvedness (Perotti et al 1998). Moreover, these singular points of a surface are highly stable under changes in viewpoint, potentially providing a basis for robust shape constancy (Phillips et al 2003). There is growing reason to believe that our primary perception of 3D shape is closer to a configuration of hills, valleys, courses, and ridges than to a Euclidean structure.

This qualitative description of 3D shape has the distinct advantage that information specific to these surface features is available in the image (Fleming, et al 2004; Koenderink 1990; Koenderink and van Doorn 1976, 1980; Lappin and Craft 2000; Todd 2004). Lappin, et al (2011) call this a “structural correspondence” between objects and images: singular points on locally smooth surfaces have corresponding singular points in patterns of stimulation. For example, the shape characteristic of a surface patch is specified by second-order spatial derivatives of the motion, disparity, and shading fields for the patch. These constitute Gibsonian invariants, facts of ecological optics that constrain the evolution of visual systems. As a consequence, the recovery of qualitative surface features is significantly more tractable than the recovery of Euclidean properties. Thus, with a re-analysis of the goal of vision, the constraints of differential geometry, and the available information, it turns out that perception of 3D structure need not be underdetermined.

The challenge

So 30 years after Marr’s book, it is again time to ask, “Does this computational theory solve the right problem?” (p. 201). And the answer is in: the goal of vision is not to recover a classical physical description of the world. As Gibson (1979) said, “the environment to be perceived... is not the world of physics but the world at the level of ecology” (p. 2).

The goal of vision cannot be assumed *a priori*. The field needs to reassess the functional tasks of vision, based on empirical studies of our perceptual competencies – the properties of the world that we *can* perceive accurately and reliably and the behaviors that we can perform successfully. I suspect that the goal of vision will end up looking less like a general-purpose representation and more like a collection of task-specific functions based on particular informational variables.

If the goals of vision are recast, then there is hope that the theoretical problem can be reformulated in such a way that perception is not underdetermined. This was Gibson’s most fundamental point: the basic challenge for creating a scientific account of vision is to reanalyze the problem so that perception is uniquely determined by information. To paraphrase Gibson (1959), “For each perceptible entity in the world, there must be a

property of stimulation, however complex, that specifies it.” This is precisely the moral of our tale about 3D shape.

Thus, in my view, a science of visual perception must do the hard work of ecological physics: empirical analysis of “the ways information is constrained by the world,” or the constraints in natural environments under which information holds. Concomitantly, it must return to the hard work of ecological optics: analysis of optical structure that, under these constraints, specifies behaviorally relevant properties. There are inspiring exemplars of both in the sensory ecology of invertebrates, bats, and electric fish. Marr clearly understood this need:

“The analytical difficulties arise from deciding what can be safely assumed about the world in order to help the processes interpret images of it. Where this can be done cleanly, more or less by inspection of the real world, we have on the whole been able to develop a clean theory. But where it cannot, I think there is no hope of understanding the processes properly until some other means have been found for determining what is safe to assume about the world and what is not... In the end, these are empirical questions, not so much about our visual systems (although the answers will be reflected in their structural design), as about the statistical structure of the visual world,” (p. 266).

In sum, the aim is to reformulate the problem of vision so that perception is uniquely determined by information within ecological constraints. Of course, despite some successes (e.g. 3D shape, optic flow, time-to-contact, the “ground” theory of size and distance, affordances) this Gibsonian approach faces its own challenges. Ecological physics and ecological optics are difficult projects, both empirically and mathematically. The approach does not imply that any environmental property is perceivable – that itself is an empirical question. There is no guarantee that optical information is available or analyzable with current methods. Perceptual judgments may be reliable and accurate for sufficiently specified properties, but when asked to go beyond the information given, judgments will be variable, subjective, and influenced by other factors (Koenderink et al 2001). Yet a science of perception must be guided by the imperative that, as Lappin and Craft (2000) put it,

“Spatial vision is only possible to the extent that mutual correspondences exist between object structure and image structure” (p. 7)

If we can define the functional tasks of vision more realistically, and can identify information for them under natural constraints, we may be able to hold out for a general theory of perception. This will allow us to simplify the requirements of the perceptual process and to specify the job descriptions for neural mechanisms.

References

- Domini F, Braunstein ML. Recovery of 3-D structure from motion is neither Euclidean nor affine. *Journal of Experimental Psychology: Human Perception and Performance*. 1998; 24:1273–1295.
- Domini F, Caudek C. 3-D structure perceived from dynamic information: A new theory. *Trends in Cognitive Sciences*. 2003; 7:444–449. [PubMed: 14550491]
- Fleming RW, Torralba A, Adelson EH. Specular reflections and the perception of shape. *Journal of Vision*. 2004; 4:798–820. [PubMed: 15493971]
- Foley JM, Ribeiro-Filho NP, DaSilva JA. Visual perception of extent and the geometry of visual space. *Vision Research*. 2004; 44:147–156. [PubMed: 14637364]
- Gibson, JJ. *Perception of the visual world*. Boston, MA: Houghton-Mifflin; 1950.
- Gibson, JJ. *Perception as a function of stimulation*. In: Koch, S., editor. *Psychology: A study of a science*. 1959. p. 456–473.
- Gibson, JJ. *The senses considered as perceptual systems*. Boston, MA: Houghton-Mifflin; 1966.

- Gibson, JJ. The ecological approach to visual perception. Boston, MA: Houghton-Mifflin; 1979.
- Gilinsky AS. Perceived size and distance in visual space. *Psychological Review*. 1951; 58:460–482. [PubMed: 14900306]
- Gogel WC. Visual perception of spatial extent. *Journal of the Optical Society of America*. 1964; 54:411–416. [PubMed: 14130693]
- Hibbard PB. Can appearance be so deceptive? Representationalism and binocular vision. *Spatial Vision*. 2008; 21:549–559. [PubMed: 19017482]
- Kersten D, Yuille A. Bayesian models of object perception. *Current Opinion in Neurobiology*. 2003; 13:150–158. [PubMed: 12744967]
- Knill, DC.; Richards, W. Perception as Bayesian inference. Cambridge: Cambridge; 1996.
- Koenderink, JJ. Solid Shape. Cambridge, MA: MIT Press; 1990.
- Koenderink JJ, Doorn AJv. Relief: Pictorial and otherwise. *Image and Vision Computing*. 1995; 13:321–334.
- Koenderink JJ, van Doorn AJ. Local structure of movement parallax of the plane. *Journal of the Optical Society of America*. 1976; 66:717–723.
- Koenderink JJ, van Doorn AJ. Photometric invariants related to solid shape. *Optica Acta*. 1980; 27:981–986.
- Koenderink JJ, van Doorn AJ, Kappers AML, Todd JT. Ambiguity and the ‘mental eye’ in pictorial relief. *Perception*. 2001; 30:431–448. [PubMed: 11383191]
- Lappin JS, Craft WD. Foundations of spatial vision: From retinal images to perceived shapes. *Psychological Review*. 2000; 107:6–38. [PubMed: 10687401]
- Lappin JS, Norman JF, Phillips F. Fechner, information, and shape perception. *Attention, Perception & Psychophysics*. 2011; 73:2353–2378.
- Loomis JM, DaSilva JA, Fujita N, Fukusima SS. Visual space perception and visually directed action. *Journal of Experimental Psychology: Human Perception and Performance*. 1992; 18:906–921. [PubMed: 1431754]
- Marr, D. Vision. San Francisco: W.H. Freeman; 1982.
- Norman JF, Crabtree CE, Clayton AM, Norman HF. The perception of distances and spatial relationships in natural outdoor environments. *Perception*. 2005; 34:1315–1324. [PubMed: 16355740]
- Norman JF, Todd JT, Norman HF, Clayton AM, McBride TR. Visual discrimination of local surface structure: slant, tilt, and curvedness. *Vision Research*. 2006; 46:1057–1069. [PubMed: 16289208]
- Norman JF, Todd JT, Perotti VJ, Tittle JS. The visual perception of three-dimensional length. *Journal of Experimental Psychology: Human Perception and Performance*. 1996; 22:173–186. [PubMed: 8742260]
- Perotti VJ, Todd JT, Lappin JS, Phillips F. The perception of surface curvature from optical motion. *Perception and Psychophysics*. 1998; 60:377–388. [PubMed: 9599990]
- Phillips F, Todd JT, Koenderink JJ, Kappers AML. Perceptual representation of visible surfaces. *Perception & Psychophysics*. 2003; 65:747–762. [PubMed: 12956582]
- Svarverud E, Gilson S, Glennerster A. A demonstration of ‘broken’ visual space. *PLoS One*. 2012; 7:e33782. [PubMed: 22479441]
- Tittle JS, Todd JT, Perotti VJ, Norman JF. Systematic distortion of perceived three-dimensional structure from motion and binocular stereopsis. *Journal of Experimental Psychology: Human Perception and Performance*. 1995:21.
- Todd JT. The visual perception of 3D shape. *Trends in Cognitive Sciences*. 2004; 8:115–121. [PubMed: 15301751]
- Todd JT, Norman JF. The visual perception of 3-D shape from multiple cues: Are observers capable of perceiving metric structure? *Perception & Psychophysics*. 2003; 65:31–47. [PubMed: 12699307]
- Todd JT, Perotti VJ. The visual perception of surface orientation from optical motion. *Perception & Psychophysics*. 1999; 61:1577–1589. [PubMed: 10598471]
- Todd JT, Reichel FD. Ordinal structure in the visual perception and cognition of smooth surfaces. *Psychological Review*. 1989; 96:643–657. [PubMed: 2798652]
- Ullman, S. The interpretation of structure from motion. Cambridge, MA: MIT Press; 1978.

- Ullman S. Against direct perception. *Behavioral and Brain Sciences*. 1980; 3:373–415.
- Wagner M. The metric of visual space. *Perception & Psychophysics*. 1985; 38:483–495. [PubMed: 3834394]
- Warren WH. Visually controlled locomotion: 40 years later. *Ecological Psychology*. 1998; 10:177–219.
- Warren WH. Direct perception: The view from here. *Philosophical Topics*. 2005; 33:335–361.