

The Construction of Visual Reality

Donald D. Hoffman

Abstract

This chapter examines the standard definitions of illusion and hallucination. These definitions assume a standard theory of perception in which a goal of perception is to estimate true properties of an objective physical world. This standard theory of perception is usually justified on evolutionary grounds: Those creatures that see more truly are ipso facto more fit. However, a closer examination of perceptual evolution using the tools of evolutionary game theory reveals that this standard assumption is incorrect. Perception has not evolved to report truth, but instead to guide adaptive behavior within a niche. In this regard, our perceptions are much like the windows desktop of a computer, which serves to guide useful interactions with the computer while sparing the user from having to know the truth of its structure and function. This understanding of perceptual evolution requires us to replace the standard definitions of illusion and hallucination with new ones that better reflect the central role of perception as a guide to adaptive behavior.

1.1 Illusions: What and Why

Many people, when viewing a windmill in the distance, report that the blades sometimes seem to rotate in the wrong direction. This is an example of a visual illusion. The standard account of such illusions says that each is an incorrect perception seen by most people when they view a specific stimulus. Illusions are rare, but the situations that trigger one person to see an illusion are likely to trigger others to see a similar illusion. Hallucinations, by contrast, are incorrect perceptions that are seen by few people and that occur in the absence of an appropriate stimulus. A person with delirium tremens, for instance, might see a spider that no one else sees.

The standard account of visual illusions naturally raises the question as to why our perceptions should be fallible. What is wrong with our visual system that allows false perceptions to occur?

To answer this question we must understand visual perception as a biological system that has been shaped by natural selection. Each organ of the body has been shaped by natural selection to contribute in specific ways to the fitness of the person. The visual system can be considered as one of the many organs of the body that makes its specific contribution to the fitness of the whole organism.

This still leaves the puzzling question as to why our perceptions are fallible. The standard account of perceptual evolution is that more accurate perceptions are more fit. For instance, the textbook *Vision Science* states that "Evolutionarily speaking, visual perception is useful only if it is reasonably accurate ... Indeed, vision is useful precisely because it is so accurate. By and large, *what you see is what you get*. When this is true, we have what is called

veridical perception ... This is almost always the case with vision ...” (Palmer 1999, p. 6). Geisler and Diehl (2003) say, “In general, (perceptual) estimates that are nearer the truth have greater utility than those that are wide of the mark.” Knill et al. (1996, p. 6) say, “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.”

The idea here is that in the struggle for survival, those of our predecessors that saw the world more truly had a fitness advantage over those that saw less truly. Predecessors with truer perceptions had a better chance of becoming our ancestors. Over many generations this shaped *Homo sapiens* to have more accurate perceptions. We are the offspring of those who saw more truly, and in consequence our perceptions are usually veridical.

From this evolutionary perspective, one answer to the question as to why our perceptions are fallible is simply that evolution is not yet done with us. We are a species in process, not a species that is the end product of an evolutionary great chain of being.

While this last answer is, as far as it goes, correct, it is far from a complete account of why perception is fallible and visual illusions occur. A more complete account requires us to understand that (1) vision is a constructive process and (2) evolution has shaped this constructive process not to deliver truth but to guide adaptive behavior. When these points are understood, we find that we must redefine the notion of illusion. We also find that illusions are an unavoidable feature of perception, and cannot be eradicated by further evolution.

1.2 Vision as Construction

Roughly half of the brain’s cortex is engaged in vision. Billions of neurons and trillions of synapses are engaged when we simply open our eyes and look around. This is, for many of us, a surprise. We think of visual perception as being a simple process of taking a picture. There is an objective physical world that exists whether or not we look, and vision is just a camera that takes a picture of this preexisting world. We can call this the camera theory of vision. Most of us, to the extent that we think about vision at all, assume that the camera theory of vision is true.

That billions of neurons are involved in vision is a surprise for the camera theory. So much computational power is not necessary to take a picture. Cameras existed long before computers.

The eye is, of course, like a camera. It has a lens that focuses an image on the retina at the back of the eye. But this is just the starting point of the visual system. From there, billions of neurons are engaged in cortical and subcortical processing. Why all this processing power?

The story that has emerged from research in cognitive neuroscience is that vision is a constructive process. When we open our eyes, our visual system constructs in a fraction of a second all the shapes, depths, colors, motions, textures and objects that we see. The computational power required for such construction is massive, but the construction is done so quickly that we are fooled into thinking that there is no construction at all, that we are simply taking a snapshot of the world as it is.

Why does the visual system bother to do all this construction? Why doesn't it simply take a picture and be done? That would certainly require less computation, and would reduce the chance of illusions.

The visual system constructs our visual worlds because it must. The starting point of vision is the photoreceptor mosaic in the retina of the eye. Each eye has roughly 120 million photoreceptors, and the activation of each photoreceptor is proportional to the number of photons it catches. One can think of the retina as starting with an array of 120 million numbers, describing the number of photons caught by each photoreceptor. There are no colors, shapes, objects, textures, motions, or depths. There is only a description that says something like, "This photoreceptor caught 5 photons, this one caught 12, this one caught...". From this array of 120 million numbers, the visual system must proceed to construct all the colors, shapes, objects and depths that constitute our visual world.

This point is painfully clear to computer scientists building robotic vision systems. The input to such a system is an array of numbers from a video camera. If the computer is going to see anything more than just this meaningless array of numbers, then it must have sophisticated programs that set about to construct visual worlds from the video input. Writing such programs has proved exceedingly difficult, and has led to great respect for the constructive powers of biological vision systems.

For any image given to the computer, there are always an infinite number of visual worlds that could be constructed that are compatible with that image. For instance, an infinite number of 3D constructions are always, in principle, compatible with any given 2D image: An ellipse in an image could be the projection of a circle seen at an angle, or the projection of any one of an infinite number of different ellipses at different angles. This makes the construction process nontrivial.

A concrete example of visual construction is illustrated in Figure 1. On the left is a collection of green lines. On the right is the same collection of green lines, but with black lines attached. Notice that on the right the green appears to fill in the space between the lines, to create a glowing green worm with fairly sharp edges. The glowing green and the sharp edges are all constructed by your visual system, an effect known as neon color spreading (Redies and Spillman 1981). You can check that you are constructing the neon worm: If you cover the black lines the worm disappears.

[Fig 1 near here]

Neon color spreading is often used as an example of a visual illusion. It fits the standard definition of an illusion. Most observers see the color spreading when they look at such a figure, and the perceived green spreading where there is no green ink is, most would agree, an incorrect perception. So here we see a case where the constructive power of the visual system leads to a visual illusion. Indeed each standard visual illusion is in fact a case where we catch the visual system in the act of a construction that is apparently incorrect (for hundreds of illusions and their explanations, see Hoffman 1998, Seckel 2009). Illusion and construction are inextricably linked.

Now the standard view of visual constructions is that they are, in the normal case, *reconstructions*. There is an objective physical world with depths, shapes, and colors, and the constructions of the visual system are, in the normal case, fairly accurate reconstructions of the true physical properties. According to this standard view, the reason that visual constructions are usually accurate reconstructions is due to evolution by natural selection. The more accurately an organism's visual system reconstructs the objective physical properties of its environment, the more fit the organism and the better its chances of surviving long enough to reproduce.

So the standard view of visual illusions is that they are the result of visual constructions that are not accurate reconstructions of the objective physical world. Evolution by natural selection has made sure that such incorrect constructions are rare.

1.3 Perceptual Evolution

One problem with the standard view of visual illusions is that it is based on an incorrect understanding of evolution by natural selection. As we noted earlier, Geisler and Diehl (2003) say, "In general, (perceptual) estimates that are nearer the truth have greater utility than those that are wide of the mark." Most vision researchers agree that truer perceptions have greater utility, and therefore contribute to greater fitness of the organism.

But this assumption, though perhaps plausible, is in fact incorrect. Truth and utility are distinct concepts, and conflating them is a fundamental error. Utility depends on the organism and the world. One cannot assign a utility to the true state of the world unless one specifies an organism. For instance, being 5000 feet below sea level has high utility for a benthic fish, but is fatal for a person. The same objective feature of the world has radically different utility for people and fish. Mathematically we can write that utility, u , is a function from the objective world, W , and an organism, O , into the real numbers, R .

$$u: W \times O \rightarrow R \quad (1)$$

So utility and truth are related as shown in (1), and therefore are not the same concepts.

Now it might still be the case that although utility and truth are distinct, nevertheless it happens to be an empirical fact that truer perceptions have greater utility. But this needs to be demonstrated. It cannot simply be assumed to be true.

One way to test this assumption is through the mathematical theory of evolution, known as evolutionary game theory (Maynard Smith 1974, Nowak 2006). Using computer simulations, one can create a wide variety of objective worlds and of organisms with different kinds of perceptual systems. The organisms can compete with each other in evolutionary games, and one can determine whether the organisms that see more truly are in fact the ones that tend to outcompete other organisms and have more offspring.

Results of such simulations have recently been published by Mark et al. (2010). They simulate a variety of worlds with varying numbers of resources, and allow organisms to compete. Some see the whole truth, others part of the truth, and still others none of the truth. The organisms in the simulations that see none

of the truth have perceptions that are tuned to utilities rather than to the objective structure of the world. For instance, a particular world might have several territories, each having a resource, such as food or water or salt, that can vary in quantity from 0 to 100. The utility of the resource is varied across simulations. Sometimes utility might be a monotonic function of the quantity of the resource, and other times it might be a Gaussian or some other non-monotonic function.

What Mark et al. find is that true perceptions are not, in general, more fit. In most cases of interest, an organism that sees none of the truth, but instead sees abstract symbols related to utility, drives the truth perceivers to swift extinction. Natural selection does not usually favor true perceptions. It generally drives them to extinction.

One reason is that perceptual information does not come free. There are costs in time and energy for each bit of information that perception reports about the environment. For every calorie an organism spends on perception, it must kill something and eat it to get the calorie. As a result, natural selection pressures perception to be quick and cheap. Getting a detailed description of the truth is too expensive in time and energy. It is also not usually relevant, since utility, not truth, is what perception needs to report.

1.4 Interface Theory of Perception

Simulations using evolutionary game theory show that perceptual systems that report the whole truth or just part of the truth are not as fit as those that report utility (Mark et al. 2010). How shall we understand these fitter perceptual systems? Are there intuitions that can help us understand why they are more fit?

The key idea is that perception serves to guide adaptive behavior. Guiding adaptive behavior is not the same as constructing veridical perceptions.

An example of the difference is the windows desktop of PC (Hoffman 1998, 2009). The desktop interface is not there to present a veridical report of the diodes, resistors, magnetic fields, voltages and software inside the PC. Instead it is there to allow the user to be ignorant of all this, and yet still interact effectively with the PC to get work done.

If the icon for a file is orange, rectangular and in the center of the display, this does not mean, of course, that the file itself is orange, rectangular and in the center of the PC box. The color of the icon is not the true color of the file; files have no colors. The rectangular shape of the icon is not the true shape of the file. The position of the icon on the screen is not the true position of the file in the computer. No property of the icon on the screen is veridical. But this does not mean that the windows interface is useless, or misleading, or an illusion. It means that the purpose of the windows interface is to guide useful interactions with the PC while allowing the user to be free of the burden of knowing its complex details.

So, with the windows interface example, we see that reporting the truth is not the only way to be helpful, and that in fact reporting the truth can be an impediment to progress, rather than a help. Perception can be useful even though it is not veridical. Indeed, perception is useful, in part, precisely because it is not veridical, and does not burden us with complex details about objective

reality. Instead, perception has been shaped by natural selection to be a quick and relatively inexpensive guide to adaptive behavior.

The view of visual perception that emerges from this evolutionary understanding can be summarized as follows. Perceived space and time are simply the desktop of the perceptual interface of *Homo sapiens*. Objects, with their colors, shapes, textures and motions, are simply the icons of our space-time desktop. Space, time, objects, colors, shapes and motions are not intended to be approximations to the truth. They are simply a species-specific interface that has been shaped by natural selection to guide adaptive behaviors that increase the chance of having kids.

One objection that often comes to mind at this point is the following. If that bus hurdling down the road is just an icon of your perceptual interface, why don't you step in front of the bus? After you are dead, and your interface theory with you, we will know that perception is not just an interface, and that it is indeed a report of the truth.

The reason not to step in front of the bus is the same reason one would not carelessly drag a file icon to the trashcan. Even though the shape and color of the file icon do not resemble anything about the true file, nevertheless if one drags the icon to the trash one could lose the file and many hours of work. We know not to take the icons literally. Their colors and shapes are not literally correct. But we also know to take the icons seriously.

Our perceptions operate the same way. They have been shaped by natural selection to guide adaptive behavior. We had better take them seriously. Those of our predecessors who did not take them seriously were at a selective disadvantage compared to those who did take them seriously. If you see a cliff, don't step over. If you see a spider, back away. If you see a moving bus, don't step in front of it. Take your perceptions seriously. But this does not logically require that you take them to be literally true.

Another objection that often comes to mind regards consensus. If a bus is hurdling down the road, any normal observer will agree that they indeed see a bus hurdling down the road. So, since we all agree about the bus, since there is consensus, doesn't that mean that we are all seeing the same truth?

But consensus does not logically imply that we are all seeing the truth. It simply implies that we have similar perceptual interfaces, and that the rules of visual construction that we use are similar. Just because an icon appears as orange and rectangular on different desktops and to different users does not mean that orange and rectangular are the true color and shape about the file. It just means that the various desktops have similar conventions that they observe.

1.5 Biological Examples

It is one thing to argue from simulations of evolutionary game theory, and from analogies with computer interfaces, that visual perception is simply a species-specific user interface that has been shaped by natural selection to guide adaptive behavior and to hide the complexities of the truth. It is quite another thing to present concrete evidence that this is how perception really works in living biological systems.

Such concrete evidence is abundant. Some of the most salient examples are seen in the phenomena of mimicry, camouflage, supernormal stimuli and ecological traps. Each of these phenomena can be understood as resulting from

natural selection shaping perception to be a quick and inexpensive guide to adaptive behavior, rather than a veridical report.

Many dragonflies, for instance, lay their eggs in water. For millions of years their visual systems have guided them to bodies of water appropriate for oviposition. This is an impressive feat, and might suggest that their visual systems have evolved to report the truth about water. Experiments reveal instead that they have evolved a quick and cheap perceptual trick (Horvath et al. 1998). Water slightly polarizes the light that reflects from it, and dragonfly visual systems have evolved to detect this polarization. Unfortunately for the dragonfly, *Homo sapiens* has recently discovered uses for crude oil and asphalt, and these substances polarize light to an even greater degree than does water. Dragonflies find pools of oil even more attractive than bodies of water, and end up dying in large numbers. They also are attracted to asphalt roads. Pools of oil and asphalt roads are now ecological traps for these dragonflies. Apparently their visual system evolved a quick trick to find water: Find something that polarizes light, the more polarization the better. In the environment in which they evolved, this trick was a useful guide to behavior, and allowed them to avoid constructing a complex understanding of the truth.

Mimicry and camouflage can be understood as arms races between organisms in which one organism exploits a vulnerability in the perceptual interface of a second, and in which the second organism in turn sometimes evolves its perceptual interface to remedy that particular vulnerability. Since perception has not evolved to report truth, but is instead a quick and cheap interface that has evolved to guide adaptive behavior, there will always be a myriad of vulnerabilities that can be exploited. Hence we find an endless and entertaining variety in the strategies of mimicry and camouflage.

1.6 A New Theory of Illusions and Hallucinations

Vision has evolved to guide adaptive behavior, not to report truth. Our perceptions of space, time, objects, colors, textures, motion and shapes are useful because they are not true, just as the icons of a computer desktop are useful because they are not true, but simply serve as guides to useful behavior.

Given that none of our perceptions are true, then we must revise the standard definition of illusions, which says that each illusion is an incorrect perception seen by most people when they view a specific stimulus. The key to a new theory of illusions is to think about the evolutionary purpose served by perceptual systems: They have evolved to be guides to adaptive behavior.

This suggests the following new definition. An illusion is a perception, experienced by most people in a specific context, that is not an adaptive guide to behavior.

The windmill illusion, for instance, in which one misperceives the movement of the blades, is an illusion because such a perception is not an adaptive guide. One could be injured by a blade whose movement is misperceived (although, fortunately, the windmill illusion usually disappears if one gets close to the windmill). Similarly, the neon color spreading shown in Figure 1 is an illusion because it is not an adaptive guide, and leads the observer to see a surface with certain chromatic properties when it is not adaptive to do so.

We must also revise the standard definition of hallucination, which says that hallucinations are incorrect perceptions that are seen by few people and that occur in the absence of an appropriate stimulus. An evolutionary framework suggests the following new definition. A hallucination is a perception experienced by few people, that occurs in the absence of an appropriate context, and that is not an adaptive guide to behavior.

The key move in the new definitions of illusion and hallucination is to replace the central role of incorrect perception in the old definitions with the new central role of guiding adaptive behaviors. Our perceptual constructions have been shaped by evolution to be cheap and quick guides to adaptive behaviors in the niches that constituted our environment of evolutionary adaptiveness. Occasionally a situation arises that triggers in most members of the species perceptual constructions that are not adaptive guides to behavior. These are illusions. And occasionally a perceptual system of a member of the species engages in an idiosyncratic perceptual construction that is not an adaptive guide to behavior. This is a hallucination.

The new definitions of illusion and hallucination incorporate an evolutionary understanding of normal perception. They alert us that, when we try to understand the nature and provenance of illusions and hallucinations, it is important to consider how our perceptual systems evolved to serve as guides to adaptive behavior in our environment of evolutionary adaptiveness.

References

- Geisler, W.S., Diehl, R.L., (2003). A Bayesian approach to the evolution of perceptual and cognitive systems. *Cognitive Science* 27, 379-402.
- Hoffman, D.D. (1998). *Visual Intelligence: How We Create What We See*. New York: W. W. Norton.
- Hoffman, D.D. (2009). *The interface theory of perception: Natural selection drives true perception to swift extinction*. In: *Object categorization: Computer and Human Vision Perspectives*. Edited by Dickinson, S., Tarr, M., Leonardis, A. and Schiele, B. Cambridge, UK: Cambridge University Press, pp. 148-165.
- Horvath, G., Bernath, B., Molnar, G. (1998). Dragonflies find crude oil more attractive than water: Multiple-choice experiments on dragonfly polarotaxis. *Naturwissenschaften*, 85, 292-297.
- Knill, D.C., Kersten, D., Yuille, A. (1999). *Introduction: A Bayesian formulation of visual perception*. In: *Perception as Bayesian Inference*. Edited by Knill, D. and Richards, W. Cambridge, UK: Cambridge University Press, pp. 1-21.
- Mark, J., Marion, B., Hoffman, D.D. (2010). Natural selection and veridical perception. *Journal of Theoretical Biology*, 266, 504-515.
- Maynard Smith, J., 1974. The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology*, 47, 209-221.

Nowak, M.A., 2006. *Evolutionary Dynamics: Exploring the Equations of Life*. Cambridge, MA: Belknap Harvard University Press.

Palmer, S. (1999). *Vision Science: Photons to Phenomenology*. Cambridge, MA: MIT Press.

Redies, C., Spillmann, L. (1981). The neon color effect in the Ehrenstein illusion. *Perception*, 10, 667-681.

Seckel, A. (2009). *Optical Illusions: The Science of Visual Perception*. Richmond Hill, Canada: Firefly Books.

Figure Captions

Figure 1. Neon color spreading. The green glowing worm on the right side of the figure is a perceptual illusion.