Does Evolution Favor True Perceptions?

Donald D. Hoffman*[a], Manish Singh[b], Justin Mark[a]

[a] Dept. of Cognitive Sciences, Univ. of California, Irvine, CA USA 92697; [b] Dept. of Psychology and Center for Cognitive Science, Rutgers Univ., New Brunswick, NJ 08901

ABSTRACT

Does natural selection favor veridical perceptions, those that more accurately depict the objective environment? Vision researchers often claim that it does. But this claim, though influential, has not been adequately tested. In this paper we formalize the claim and a few alternatives. We then discuss the results of evolutionary games and genetic algorithms that indicate that veridical perceptions can be driven to extinction by non-veridical strategies that are tuned to fitness rather than to objective reality. This suggests that natural selection need not favor veridical perceptions, and that the effects of natural selection on perception deserve further study. We then address the question: How can an organism act effectively in an environment that it does not perceive veridically? (Acknowledgement: Brian Marion and Kyle Stevens are collaborators in this research.)

Keywords: Evolutionary games, perception, natural selection, Bayesian inference, genetic algorithms

1. INTRODUCTION

Metaphysical solipsism is the strange doctrine that nothing exists except my own mental states. My perceptual experiences—e.g., of an external space-time world containing inanimate objects and animate creatures with their own mental states—refer to nothing beyond my perceptual experiences. There is no objective reality outside my experiences to which they correspond.

There are few metaphysical solipsists; indeed, any such solipsist must, to be consistent, insist that there is but one. The rest of us, to the contrary, are willing to wager, e.g., when we buy life insurance, that there is a reality that is objective in the sense that it would exist even if we did not.

The majority of us believe, moreover, that our perceptual experiences give genuine insight into the true nature of that objective reality. In particular, most of us believe that the reason we have perceptual experiences of a space-time world containing animate and inanimate objects is because the objective world is indeed spatiotemporal and in fact contains such objects. We admit that our perceptions do not exhaustively describe the objective world: we cannot see, without aid of technology, the very small, the very large, the very distant; nor can we see beyond the 400 to 650 nanometer window of electromagnetic wavelengths to which our eyes are tuned. And we admit that we are subject, on occasion, to illusions and misperceptions. But within these limits we are confident that our perceptions give us a real window into the true nature of the objective world.

---

* ddhoff@uci.edu; phone 949-679-1490; http://www.cogsci.uci.edu/~ddhoff/

It is not just the man on the street that enjoys this confidence. So too do many professionals of vision science. Marr, for instance, argues that "… usually our perceptual processing does run correctly (it delivers a true description of what is there)…" [1] Lehar agrees, saying that "The primary function of perception [is] that of generating a fully spatial virtual-reality replica of the external world in an internal representation." [2]

This confidence, on the part of professionals, that vision "delivers a true description of what is there" is based largely on an argument from evolution, which we can call the True-Vision Argument.

True-Vision Argument: Truer perceptions are, ceteris paribus, fitter, and therefore the genes that code for truer perceptions are more likely to be passed to the next generation. Thus natural selection shapes vision to deliver a true description of what is there.

According to the True Vision Argument, we are the offspring of a line of ancestors who saw progressively more true descriptions of what is there. Marr, for instance, argues that "… one interesting aspect of the evolution of visual systems is the gradual movement toward the difficult task of representing progressively more objective aspects of the world." [1] Palmer concurs, saying that, ''evolutionarily speaking, visual perception is useful only if it is reasonably accurate.'' [3] Geisler and Diehl agree, saying that ''In general, (perceptual) estimates that are nearer the truth have greater utility than those that are wide of the mark.'' [4] Yuille and Bulthoff also agree, saying "visual perception … involves the evolution of an organism's visual system to match the structure of the world…" [5]

Although most vision researchers are confident that human vision delivers a true description of what is there, they concede that vision in more primitive organisms might not. Marr, for instance, suggests "… it is extremely unlikely that the fly has any explicit representation of the visual world around him—no true conception of a surface, for example, but just a few triggers and some specifically fly-centered parameters…" [1] He elaborates that "In a true sense, for example, the frog does not detect *flies*—it detects small, moving black spots of about the right size. Similarly, the housefly does not really represent the visual world about it—it merely computes a couple parameters … which it inserts into a fast torque generator and which cause it to chase its mate with sufficient frequent success." [1]

Why should primitive vision systems have evolved to "not really represent the visual world," and instead to deliver simple triggers and organism-centered parameters? Marr argues, in the case of the fly, that "One reason for this simplicity must be that these facts provide the fly with sufficient information for it to survive." [1] Evolution by natural selection depends on the *fitness* of organisms: Fitter organisms, with fitter vision systems, are more likely to pass their genes to the next generation. The key point is that fitness is distinct from truth. Organisms that evolve fitter vision systems through natural selection need not, ipso facto, evolve to see the truth. They can instead evolve simple triggers, organism-centered parameters, and niche-specific heuristics that guide fitter behavior. Notice that this argument, which is simply the standard argument from natural selection, differs fundamentally from the True Vision Argument, because it does not claim that truer perceptions are necessarily fitter.

How might true perceptions make an organism less fit? If a true perception takes more time or energy to compute, or if it represents aspects of the world that are irrelevant for guiding adaptive behaviors of the organism, then it could adversely affect fitness. Natural selection ultimately judges perceptions by the behaviors they guide. If true perceptions take too much time (letting more predators attack and more prey escape), consume too many calories (requiring more foraging for calories), or provide irrelevant descriptions, then natural selection is likely to judge them harshly.

Are there examples in nature where selection has in fact shaped perceptual systems that fail to deliver a true description of what is there? Many. Perhaps best known are cases of releasing stimuli and supernormal stimuli, which are ubiquitous in nature.[6–8] The jewel beetle *Julodimorpha bakewelli*, for instance, has wing casings that are bumpy, glossy and brown. Beer bottles that are bumpy, glossy and the right shade of brown trigger males of the species to forsake real females and mate instead with the bottles.[9] The perceptual systems of the males fail to deliver a true description of what is there, and instead provide simple triggers that guide mating behaviors that are adaptive in the niche in which they evolved, a niche that did not include beer bottles.

When one species evolves perceptions that fail to deliver a true description of what is there, another species can evolve to take advantage of this failure. One example is “code breaking.” The brown-headed cowbird of North America, for instance, lays an egg in the nest of a bird of another species. When the egg hatches, the young cowbird provides a supernormal stimulus to the mother of the other species, which then feeds it preferentially to its own offspring.[10] The cowbird successfully practices such “brood parasitism” on more than 200 different species of birds; clearly the evolution of nonveridical perceptions is not a rare and isolated event.

In examples like the jewel beetle and brood parasitism, nonveridical perception is strongly suggested by “fixed action patterns,” simple stereotypical behaviors that are triggered by a perceptual stimulus, that normally are adaptive in the niche in which they evolved, but which rigidly play out even if they are not adaptive in the current context.[11]

However, there is no logical requirement that nonveridical perceptions must trigger simple and inflexible behaviors. In principle, they could guide a wide range of complex, flexible and adaptive behaviors. This is a point made by the “interface theory of perception,” which proposes that perceptions can be modeled as species-specific interfaces that function much like the windows interface on a PC.[12–17] The colors, shapes and locations of icons on a windows desktop are not a true description of the colors, shapes and locations of files in a computer. Indeed, files do not have colors and shapes, and placement of an icon on the upper left corner of the desktop does not entail that the corresponding file resides in the upper left corner of the computer. The desktop interface guides useful behavior in part because it hides the truth from the user—it hides the diodes, resistors, circuits, voltages, magnetic fields, and megabytes of software that the user doesn’t need to know about to successfully use the computer to write a paper or edit an image. Instead, the interface provides symbols that allow the user to flexibly engage in a wide variety of complex interactions with the computer. The perceptions of the computer provided by the desktop interface are not true, but they are useful.

The distinction between fitness and truth becomes obvious when one formally specifies a fitness function. To determine a fitness function *f* one must specify, at a minimum, a state of an objective world $w \in W$, an organism $o \in O$, a state of the organism $s \in S$, and an action class $a \in A$.[14–16] Thus a fitness function is, at a minimum, a mapping $f: W \times O \times S \times A \to \mathbb{R}$. Fitness and truth are distinct because *f* and *W* are distinct.

2. BAYESIAN DECISION THEORY AND EVOLUTIONARY GAMES

Given the abundant evidence in nature of organisms with nonveridical perceptions, and given that natural selection is driven by fitness, not by truth, it is natural to ask whether vision researchers are correct in assuming that *Homo sapiens* has evolved to see a true description of what is there.

The standard formal framework in which vision researchers cast this assumption is Bayesian decision theory, or BDT.[18–23] Given some images y_0 and possible scene interpretations $x \in X$, the visual system is assumed to compute the conditional probability measure $p(x \mid y_0)$ using Bayes formula, $p(x \mid y_0) = p(y_0 \mid x)p(x) / p(y_0)$. The images y_0 might, for instance, be left and right eye images at a single instant in time, and x might be a three-dimensional interpretation of those images derived from stereovision. The "posterior" probability measure $p(x \mid y_0)$ gives the probability of various possible scene interpretations, given the images y_0. The "likelihood function" $p(y_0 \mid x)$ is sometimes called a "rendering function", because it gives the probability that an image y_0 will be rendered given that the real scene is x. The probability measure $p(x)$ is the "prior probability" of the scene interpretation x, and represents assumptions built into the visual system. Given the posterior $p(x \mid y_0)$ and a loss or gain function, the visual system picks the best scene interpretation, i.e., the interpretation that minimizes expected loss or maximizes expected gain. This is the interpretation that is perceived. BDT has proved to be a valuable framework for vision research.

The assumption of vision researchers that *H. sapiens* has evolved to see veridically is reflected in the assumption, made standardly in BDT, that *X* corresponds both to the objective world and to the space of possible visual interpretations. BDT of course allows for illusions, i.e., that an observer might see an illusory interpretation x_i rather than the interpretation x that should be seen. But it assumes that X is the right language for a true description of what is there, and that at least one interpretation $x \in X$ is in fact a true description of what is there.

If we wish to test the assumption that *H. sapiens* has evolved to see veridically, and allow the possibility that the answer might be negative, then we must have a formal framework that allows that *none* of the possible perceptions $x \in X$ is in fact a true description of what is there. Thus we must go beyond the BDT framework as it is now used in vision research.

The first step in this direction is to define the notion of a *perceptual strategy*, so that we can compare veridical and nonveridical strategies, and assess which are more fit and therefore more likely to evolve by natural selection.[14–16] To define a perceptual strategy, we assume that the possible perceptions X of an organism are not a priori identical to the possible states of the objective world W. We can then define a perceptual strategy as a mapping $P: W \to X$, where we assume nothing about W and X other than that we can speak sensibly about probabilities of events on each, i.e., that each has an associated σ-algebra or, more generally, an associated σ-additive class.[24] In the dispersion-free case, we can take P to be a measurable function (also known, if X is a subset of real numbers, as a random variable). More generally, P is a Markovian kernel which assigns for each state $w \in W$ a probability measure, $P(w,\cdot)$, on X.[25] We think of $P(w,\cdot)$ as the probability of seeing various interpretations (where interpretations are denoted by the variable $\cdot$) within the set X given that the state of the objective world is w. Formally, P is a linear operator, a generalization of the familiar matrix operators to the continuous case. Viewed as a linear operator, P maps functions and measures on W to functions and measures on X. In the language of information theory, P is a channel. The relationship between perceptual strategies and BDT is illustrated in Figure 1.

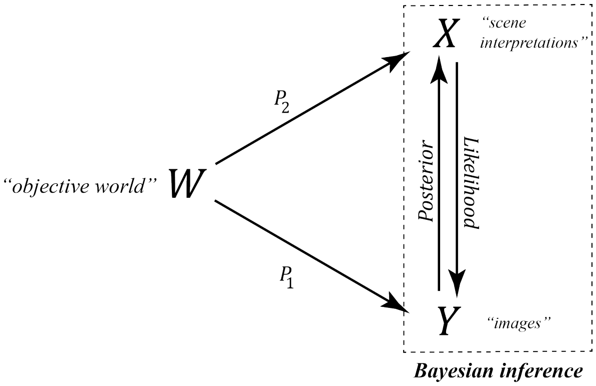

Figure 1. The relationship between BDT and perceptual strategies. BDT corresponds to the contents of the dashed box. The arrows labeled P_1 and P_2 represent perceptual strategies.

Different classes of perceptual strategies can be defined by specifying conditions on P.[14–16] The most restrictive class is the *naïve realist*, which requires that $W = X$ and that P, in the dispersion-free limit, is an isomorphism. The class of *strong critical realist* strategies weakens this by requiring only that $X \subset W$ and that P, in the dispersion-free limit, is an isomorphism on X. The class of *weak critical realist* strategies weakens this further by dropping the requirement that $X \subset W$ but requiring that P, in the dispersion-free limit, is a homomorphism of all structures on W. Most vision scientists are weak critical realists. Finally, the class of *interface* strategies weakens the requirements further by relaxing the homomorphism requirement, and requiring only that P is a Markovian kernel or, in the dispersion-free limit, a measurable function. Thus the set of naïve realist strategies is a proper subset of the strong critical realist strategies, which is a proper subset of the weak critical realist strategies, which is a proper subset of the interface strategies. An interface strategy is a *strict* interface strategy if it is not also weak critical realist strategy.

Which class of strategies is favored by natural selection? This question can be addressed using the tools of evolutionary game theory.[26–27] Artificial worlds can be simulated, having different objective structures and fitness functions. Simulated organisms with different perceptual strategies compete with each other in these worlds. One can determine which strategies ultimately dominate, which go extinct, and which coexist. Simulations thus far give a clear message: If fitness is not monotonically related to objective structures, and if memory and computation are not free of cost, then realist strategies often go extinct when competing with strict interface strategies.[14] In general, fitness need not be monotonically related to objective structures. For instance, the fitness consequences to humans for ingesting salt is not monotonically related to the quantity of salt: too little salt is fatal, a modest amount is ideal, too much is again fatal. Similar remarks hold, mutatis mutandis, for water, sunlight, heat, food, and endless other examples. Indeed, it is difficult to find examples where fitness is monotonically related to objective structures.[1]

3. COMPUTATIONAL EVOLUTIONARY PERCEPTION

We have seen that perceptual strategies are naturally represented as Markovian kernels.[14–15] So too are decisions and actions.[16] This provides a convenient mathematical framework to investigate the co-evolution of perceptions, decisions and actions. [16]

As before, we denote by W the set of possible states of the objective world and by X the set of possible perceptions of an organism. We denote by G the set of possible actions for the organism. We use the symbol G because we think of many of the possible actions as forming a group that acts on X. For instance, if X contains a 3D spatial representation, then G might contain actions such as translation and rotation that act on X. However, we also allow G to contain actions, such as eating, which might not easily be represented as group actions on X.

We have already modeled perceptual strategies by means of a Markovian kernel P, which is a linear operator, or channel, from W to X. To model decisions, we now introduce the Markovian kernel D, which is a linear operator from X to G. We can think of this kernel as assigning to each perceptual state $x \in X$ a probability measure on possible actions $g \in G$; intuitively, an organism decides which action to take based on how it currently perceives. To model actions, we also introduce the Markovian kernel A, which is a linear operator from G to W. We can think of this kernel as assigning to each possible action $g \in G$ a probability measure on possible world states $w \in W$; intuitively, an organism acts by changing the state of the objective world W.

Just as one can multiply two matrix operators together to get a new matrix operator, so one can also take the product of two Markovian kernels to get a new Markovian kernel. For instance, one can take the product of P and D to get the kernel PD, which is a linear operator from W to G. One can take the product AP, which is a linear operator from G to X. One can take the triple product DAP to get a linear operator from X to X. Thus the operators P, D, and A allow one to move from the objective world W to the sensory perceptions X to the actions G and back to the objective world W in a commuting loop called the *PDA loop*. [16] This is illustrated in Figure 2.

---

[1] Here we take items such as salt, water, sunlight, heat and food to be objective just for sake of example; if the perceptions of *H. sapiens* are strict interface, then all of these items are simply species-specific symbols and none are objective properties of the world.

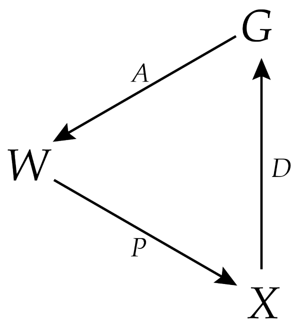

Figure 2. The PDA loop.

Natural selection will co-evolve the perceptual strategies P, the decisions D, and the actions A, which entails as well the co-evolution of X and G. Since the organism itself is part of the objective world W, natural selection is also co-evolving W.

4. GENETIC ALGORITHMS

Prior simulations using evolutionary games suggest that in many situations natural selection favors strict interface strategies.[14] For such strategies the perceptions X of an organism give it no insight into the true structure of the objective world W. Instead, perceptions are tuned by natural selection to guide adaptive behaviors.

But this raises the question: How can an organism act adaptively in a world of which it is utterly ignorant? Is this even possible? Doesn't an organism need to know the world to act properly in it?

In this section we begin to address this question using genetic algorithms (GAs) to model evolution of the *PDA* loop. But a brief example might help allay fears raised by this question. A person can interact effectively with a computer even when utterly ignorant of diodes, resistors, magnetic fields and C++ programming, i.e., when utterly ignorant of the true nature of a computer. The windows interface guides adaptive actions of the user while keeping the user blissfully ignorant of the complex truth.[12] The GAs explored here find that natural selection can effectively evolve organisms whose perceptions and actions work like such an interface, guiding adaptive behaviors while keeping the organism blissfully ignorant of the complex truth.

We start with an engaging GA presented by Mitchell.[28] The world W is a 10 x 10 grid of squares, surrounded on all sides by walls, with soda cans randomly strewn into the grid, so that some squares have one can and others have none. Thus we can model W as a 12 x 12 grid of squares (where the outer border corresponds to the wall), where each square can have one of three states: no can, one can and wall. We can label these states, respectively, as 0, 1 and 2, and let $W(i, j)$ denote the state of the square at location (i, j) . Since all and only squares on the outer border of the 12 x 12 grid are always walls, there are $2^{(10\times10)} = 2^{100}$ possible states of W.

The GA evolves a robot called "Robby," who can only see the contents of the square that he occupies, and the contents of the squares immediately north, south, east and west. Thus Robby's perceptual space X is an ordered set of 5 elements, {current, north, south, east, west}, each of which can take one of 3 values: no can, one can, and wall. Therefore there are at most $3^5 = 243$ possible perceptual states for Robby (a few of these states don't occur since, for instance, a wall cannot appear to Robby's north and south at the same time). Robby does not know W, only X. That is, Robby's perceptions at any time are limited to 5 numbers, and Robby has no idea about a two-dimensional grid or about any spatial relationships of his perceptual elements, {current, north, south, east, west}. The names of these perceptual elements are simply a convenience for us, not an indication that Robby knows anything about their spatial relationships.

The perceptual kernel P from W to X is a many-to-one mapping that forgets all elements of W except 5, viz., the element Robby occupies and the 4 elements that immediately border it. If Robby is at location (i, j) then Robby's perception is simply the set $\{W(i, j)\}$, $W(i, j+1)$, $W(i, j\text{-}1)$, $W(i+1, j)$, $W(i\text{-}1, j)\}$. Note, of course, that the appearance of the coordinates i and j in this mapping does not entail that Robby knows anything about a two-dimensional array; Robby only knows the 5 numbers that he sees, and nothing else.

Robby's job is to explore the grid, avoiding the wall, and picking up as many soda cans as possible. To do this, Robbie has only 7 primitive actions he can take at any instant: step north, step south, step east, step west, step randomly, stay put, and pick up a soda can. Thus G consists of these 7 primitive actions. On each run, Robby is given 200 total time steps at which to take one of these 7 actions.

The decision kernel D from X to G assigns to each perceptual state $x \in X$ a single action $g \in G$. Since there are 243 possible perceptual states and 7 primitive actions, there are 7^{243} different decision kernels possible. This is a huge space of possibilities, and Mitchell uses natural selection, in the form of a GA, to search through this space for decision kernels that embody good foraging strategies.

To specify the action kernel A from G to W we must slightly enlarge our definition of W to include the current location (i, j) that Robbie occupies (an organism is, after all, also part of the objective world). Thus, the possible states of W include the 2^{100} states that we mentioned earlier, multiplied by the 100 possible locations that Robby can occupy. With this enlarged W, the action kernel A assigns to each action $g \in G$ a new state $w \in W$, as follows. Denote Robby's current location by (i, j). If Robby steps north, then the state of W is unchanged except that Robby's location is changed to $(i, j+1)$. Similar changes to W occur, mutatis mutandis, for steps south, east, west, and random. If Robby stays put, then the state of W is unchanged. If Robby picks up a soda can, then the state of W is unchanged except that $W(i, j)$ is set to 0. Note that Robby does not know his current location (i, j), he only knows the 5 numbers that constitute his current perceptual state $x \in X$.

Finally, Mitchell specifies the fitness consequences of actions. Robby gets 10 points for picking up a can, loses 1 point for trying to pick up a can where there is none, and loses 5 points for trying to walk into a wall.

With this set up, Mitchell uses a GA to evolve decision kernels that embody better and better foraging strategies. The "genes" are a specific assignment of one of the 7 primitive actions to each one of the 243 possible perceptual states, and can be coded by a string of 243 numbers, each number varying between 0 and 6. Each such gene codes for a specific decision kernel, which embodies a foraging strategy. Mitchell initially chooses a population of 200 random individuals, each with a random set of 243 numbers that constitute its genes. She then lets each of these random individuals try foraging 100 different times (in each of these tries the individual gets to take 200 primitive actions). She then computes for each random individual the average number of points it got over its 100 foraging expeditions. This is the fitness of its genes, and of its foraging strategy. She then simulates natural selection by choosing two "parents" based on their fitness, with fitter strategies more likely to be chosen. She then mates the two parents by randomly choosing a position to split their genes, and combining the first portion from one parent with the latter portion of the other to create one "child" with a new set of genes; she reverses which portions are combined to create a second child. She then, with small probability, mutates some of the numbers in the new genes, by randomly choosing a few of the numbers and replacing them with a random number between 0 and 6. She adds the two children to a new population, and continues this until the new population reaches 200 individuals. Then she repeats the entire process of finding the fitness of each individual and using this to govern the probabilities of being a parent to the next generation. She does this for 1000 generations. The initial population does poorly. But by the last generation, some very effective foraging strategies have evolved.

In our first modification of Mitchell's GA, we replace soda cans with water. A soda can was either present or absent, but our water comes in eleven different quantities (0, 1, 2, …, 9, 10), and each square of the grid randomly gets a single quantity of water; one square might get 6 units of water, while a square to its left gets 2. We also modify the payoffs. Picking up a soda can gave 10 points of fitness. But now the fitness points for water depends non-monotonically on the quantity of water as follows: (0, 1, 3, 6, 9, 10, 9, 6, 3, 1, 0).[2] Thus, a quantity of 0 and a quantity of 10 both give 0 fitness points, and a quantity of 5 gives 10 fitness points. In Mitchell's GA, Robby saw whether there was a soda can or not in each of the 5 squares it could see. In our GA, Robby does not see the exact quantity of water in the 5 squares it can see. Instead, it just sees two colors, "red" and "green." It must evolve a perceptual strategy that determines which quantities of water to see as green and which to see as red. So the initial set of 243 genes is expanded to include 11 more genes, one for each possible quantity of water. For each of these extra genes there are two possible perceptions that are associated: red or green. Thus there are 2^{11} possible perceptual strategies that are coded for by these genes. Our GA is thus evolving Robby's perceptions as well as his actions.

When we run the GA, the initial population, as expected, performs quite poorly. But after 500 generations we find that sophisticated foraging strategies have evolved, just as in Mitchell's GA. What is new is that we find that two kinds of perceptual strategies have evolved. In the first kind, the water quantities 0, 1, 9, and 10 are seen as red and the other quantities as green. In the second kind the water quantities 0, 1, 9, and 10 are seen as green and the other quantities as red. In the first case, Robby preferentially drinks from squares that look red, and in the second case Robby preferentially drinks from squares that look green. His perceptions have evolved not to represent water quantity (red does not mean more water, green does not mean less) but rather to represent the fitness points that would accrue from trying to drink from that quantity of water. So Robby has evolved perceptions tuned to fitness, not to truth; Robby has evolved a strict interface

---

[2] Informally, we can think of these fitness point as saying that if there is not enough water an organism dies of thirst, too much and it dies by drowning, just the right amount and it thrives.

strategy, not a realist strategy. The difference between the two perceptual strategies is simply a matter of genetic drift: There is no more benefit in coding high fitness quantities of water as red than as green.

In our second GA, we modify things further by not requiring Robby to see in each of the four directions (north, south, east, west). Instead, Robby can evolve to see or not to see in any of these four directions. Robby pays a fitness cost for each direction he can see. Robby also gets two new actions: rotate 90 degrees to the left or right. No cost is assessed for such rotations, but they do count in the total of 200 actions that Robby can take on a given foraging run. We find again that Robby evolves interface perceptual strategies and effective foraging strategies. For instance, as the cost for seeing in different directions increases, Robby evolves to see in fewer directions.

5. CONCLUSION

It is widely assumed, by laymen and experts alike, that our perceptual systems deliver true, though not exhaustive, descriptions of the objective world. This assumption can be studied using mathematical models of evolution by natural selection. The studies thus far, using evolutionary game theory and genetic algorithms, indicate that truer perceptions are not necessarily fitter, and that a simple interface strategy can drive true perceptions to extinction. The ubiquity in nature of phenomena such as supernormal stimuli and releasing stimuli suggest that these mathematical results are not artifacts and that many, perhaps all, organisms have been shaped by natural selection not to see truth but instead to have perceptions that guide behavior that is adaptive in the relevant niche. In short, perception seems to be tuned to fitness rather than truth.

Can an organism act adaptively in an environment that it does not see truly? Our GA results and those of Mitchell show that this is in fact possible. Our GAs evolve an organism that effectively forages in an artificial world that is a two-dimensional grid of squares, even though the organism itself is utterly ignorant of the two-dimensional nature of this world.

These results invite further studies and simulations. The PDA loop formalism discussed here provides a mathematical framework for investigating the co-evolution of perceptions, decisions, and actions. Using this formalism in the context of evolutionary games and genetic algorithms should provide new insights into this co-evolution.

ACKNOWLEDGMENTS

We thank J. Feldman, P. Foley, B. Marion, D. Peshek and K. Stevens for helpful discussions. This work was supported in part by a grant from Procter & Gamble (DH) and by NIH grant EY021494 (MS).

REFERENCES

[1] Marr, D., [Vision: A Computational Investigation into the Human Representation and Processing of Visual Information], Freeman, San Francisco (1982).

[2] Lehar, S., "Gestalt isomorphism and the primacy of subjective conscious experience: a Gestalt Bubble model," Behavioral and Brain Sciences 26, 375-444 (2003).

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

[4] Geisler, W.S., Diehl, R.L., "A Bayesian approach to the evolution of perceptual and cognitive systems," Cognitive Science 27, 379–402 (2003).

[5] Yuille, A., Bulthoff, H. "Bayesian decision theory and psychophysics," In: Knill, D. and Richards, W. (Eds.) [Perception as Bayesian Inference], Cambridge University Press, Cambridge (1996).

[6] Tinbergen, N. [The Study of Instinct], Clarendon Press, Oxford (1951).

[7] Tinbergen, N. [The Herring Gull's World], Collins, London (1953).

[8] Barrett, D. [Supernormal Stimuli: How Primal Urges Overran Their Evolutionary Purpose], Norton, New York (2010).

[9] Gwynne, D. T. and Rentz, D. C. F., "Beetles on the Bottle: Male Buprestids Mistake Stubbies for Females (Coleoptera)," Australian Journal of Entomology, 22(1), 79–80 (1983).

[10] Ortega, C. [Cowbirds and Other Brood Parasites], Univ. of Arizona Press (1998).

[11] Alcock, J. [Animal Behavior, Eighth Edition], Sinauer, Sunderland, MA (2005).

[12] Hoffman, D. D., [Visual Intelligence: How We Create What We See], W. W. Norton, New York (1998).

[13] Hoffman, D.D., "The interface theory of perception," In: Dickinson, S., Tarr, M., Leonardis, A., Schiele, B. (Eds.) [Object Categorization: Computer and Human Vision Perspectives], Cambridge University Press, Cambridge (2009).

[14] Mark, J., Marion, B. and Hoffman, D. D., "Natural selection and veridical perception," J. Theoretical Biol. 266, 504-514 (2010).

[15] Hoffman, D.D. and Singh, M., "Computational evolutionary perception," Perception 41, 1073-1091 (2012).

[16] Singh, M. and Hoffman, D.D., "Natural selection and shape perception," In Dickinson, S. and Pizlo, Z. (Eds.) [Shape Perception in Human and Computer Vision: An

Interdisciplinary Perspective], Springer, New York (2013).

[17] Koenderink J. J., "Vision and information," In Albertazzi, L., Tonder, G.V. and Vishnawath, D. (Eds.) [Perception Beyond Inference. The Information Content of Visual Processes], MIT Press, Cambridge, MA (2011).

[18] Knill, D. and Richards, W. (Eds.) [Perception as Bayesian Inference], Cambridge University Press, Cambridge (1996).

[19] Geisler, W.S. and Kersten, D., "Illusions, perception and Bayes" Nature Neuroscience 5, 508– 510 (2002).

[20] Glimcher, P.W., [Decisions, Uncertainty, and the Brain : The Science of Neuroeconomics], MIT Press, Cambridge, MA (2003).

[21] Jaynes, E.T., [Probability Theory : The Logic of Science], Cambridge University Press, Cambridge (2003).

[22] Mamassian, P., Landy, M. and Maloney, L.T., "Bayesian modeling of visual perception", In Rao, R., Olshausen, B. and Lewicki, M. (Eds.) [Probabilistic Models of the Brain: Perception and Neural Function] MIT Press, Cambridge, MA (2002).

[23] Maloney, L.T. and Zhang, H., "Decision-theoretic models of visual perception and action" Vision Research 50, 2362-2374 (2010).

[24] Gudder, S. [Quantum Probability], Academic Press, Boston (1988).

[25] Revuz, D. [Markov Chains], North-Holland, Amsterdam (2005).

[26] Nowak, M. [Evolutionary Dynamics: Exploring the Equations of Life], Belknap Press of Harvard University Press, Cambridge, MA (2006).

[27] Maynard Smith, J. [Evolution and the Theory of Games], Cambridge University Press, Cambridge (1982).

[28] Mitchell, M. [Complexity: A Guided Tour], Oxford University Press, Oxford (2011).