

A reflexive dispositional analysis of mechanistic perception

John Dilworth

Received: 24 November 2004 / Accepted: 10 October 2006 / Published online: 17 November 2006
© Springer Science+Business Media B.V. 2006

Abstract The field of machine perception is based on standard informational and computational approaches to perception. But naturalistic informational theories are widely regarded as being inadequate, while purely syntactic computational approaches give no account of perceptual content. Thus there is a significant need for a novel, purely naturalistic perceptual theory not based on informational or computational concepts, which could provide a new paradigm for mechanistic perception. Now specifically *evolutionary* naturalistic approaches to perception have been—perhaps surprisingly—almost completely neglected for this purpose. Arguably perceptual mechanisms enhance evolutionary fitness by facilitating *sensorily mediated causal interactions* between an organism Z and items X in its environment. A ‘reflexive’ theory of perception of this kind is outlined, according to which an organism Z perceives an item X just in case X causes a sensory organ z_i of Z to cause Z to acquire a disposition toward the *very same item X* that caused the perception. The rest of the paper shows how an intuitively plausible account of mechanistic perception can be developed and defended in terms of the reflexive theory. Also, a compatibilist option is provided for those who wish to preserve a distinct informational concept of perception.

Keywords Dispositions · Evolutionary theories · Functionalism · Informational semantics · Intentionality · Low versus high level perception · Mechanistic perception · Naturalistic theories

Presumably a general, purely naturalistic account of perception would attempt to explain all kinds of perception—whether specifically human, more generally biological, or purely mechanical or robotic perception—exclusively in causal terms, i.e. in terms of causal laws, structures, and functional causal dispositions. But in biological cases, the issue is complicated by the wide variety of

J. Dilworth (✉)
Department of Philosophy, Western Michigan University, Kalamazoo, MI 49008, USA
e-mail: dilworth@wmich.edu

physiological differences between the sense-organs of various species. Also, in human and other higher mammalian cases, the issue is even further complicated by issues of intentionality, and of perceptual representation versus misrepresentation.

Hence it would be highly desirable, if at all possible, to try to find some relatively simple or basic kind of perceptual structure that could be mechanically implemented, in relative independence from such physical complexity and intentionality issues. Thus ideally one would wish for a very basic, scalable functional structure, that would be recognizably perceptual even in very simple mechanistic cases, yet which could also provide a central, purely naturalistic explanatory component in any biological cases of perception, including higher intentionality cases.

However, some caution is needed in the investigation. It might be that some specifically biological or evolutionary aspects of perception are central to the concept, so that it could only be applied in purely non-biological, mechanistic cases by analogy or extension. Thus, even if biological perception constitutes a natural kind that is identifiable independent of our concepts, it might be that purely mechanistic uses are in some significant ways dependent on human interests and concepts. For example, there is a common folk psychological view that strictly speaking no mechanical device could itself literally perceive anything, even though humans might use such devices to extend their own perceptual capacities, and it is at least conceivable that there might be theoretical support for such a view.

Nevertheless, a fully general, naturalistic theory of perception presumably should be able to settle such issues. If perception is explained using a broadly teleological concept of evolutionary proper function, such as by Millikan (1984), Dretske (1995) and Price (2001), then it would follow that mechanical devices lacking the required evolutionary history could not perceive anything. But if instead a purely non-teleological theory could be developed, it could both be compatible with a Darwinian account of the evolutionary fitness-promoting characteristics of perceptual mechanisms, and also be mechanically realizable. Such a theory will be described here, and applied to mechanistic cases.

To be sure, there is already a large and thriving interdisciplinary scientific area dealing with what is often called 'machine perception', which applies methods from computer science, information processing, robotics, artificial intelligence, and other more specialized fields to a wide range of practical problems concerning the collection and processing of sensory or perceptual data (Cantoni, 2005). However, these fields already presuppose answers to the more foundational naturalistic issues concerning perception to be investigated here, such as the almost universal assumption that perception necessarily involves the processing of semantic information, when currently there is no purely naturalistic and non-controversial analysis of such information that is available (Cummins, 1996).

Also, as a related concern, widely used computational models of perception are also philosophically highly controversial, in that philosophers such as John Searle (1983, 1992) deny that any purely computational process could instantiate the kinds of original intentionality found in human perception and thought. Thus there is still plenty of room for a different kind of biologically based, naturalistic foundational approach to mechanistic perception, which at least initially would bypass informational and computational issues altogether.

Sensory input and sensors

It will be useful to begin with a factor in perceptual theories that is widely agreed on. It is generally agreed that, whatever else is involved in perception, at least it involves processes in which *sensors* or *sensory organs* are used to provide a causal input from the environment to organisms, a process that many describe as one of collecting relevant information or data about the environment. Indeed, the best-known naturalistic approach to perception, that of informational semantics as developed by Dretske, Fodor et al., characterizes the relevant kind of sensory information in terms of the *nomic covariance* relations that typically hold between sensory inputs and outputs (Dretske, 1981; Fodor, 1990). On this view the function of the human eye is that of a transducer, which transforms input radiant energy into nomic covarying output electrical impulses in the optic nerve. Then the characteristic informational semantics claim is that such nomic covariance relations can be used to underwrite a concept of the *information* conveyed by such sensory transducers to the brain.

However, it is now widely agreed even by the supporters of such theories, including Dretske and Fodor themselves, that sensory transduction by itself does not provide a complete account of perception or perceptual content (Cummins, 1996; Dretske, 1995; Fodor, 1990). Possession of working sensory organs or mechanical sensors can generally be agreed to be a necessary condition of any kind of perception, but it is not sufficient. But exactly what more is required? I shall argue that significant clues are provided by evolutionary considerations about the adaptive role both of perception itself, and of the integral contributions of sensory organs to that adaptive role.

The fitness-enhancing evolutionary role of perception via sensory organs

On standard evolutionary views, structures or traits in organisms would tend to become more common insofar as they contribute to the evolutionary fitness of species that randomly happen to possess or acquire those traits. Thus an evolutionary account of the development of perception and of sensory organs must explain their contributions to evolutionary fitness. What follows is a very schematic account of that development, explained more fully in Dilworth (2004, 2005a, b, c).

To begin, the basic evolutionary picture—prior to the consideration of any perceptual factors—is of a world such that the only causally available way in which a species *S* can improve its evolutionary fitness is by *causal interactions* of its members *Z* with other items *X* in their environment. Two kinds of items *X* are of particular importance, namely those items *H* that are potentially harmful to *Z* on the one hand, such as predators, and on the other hand items *B* potentially beneficial to *Z*, such as food. *Z* and its species will tend to improve their fitness insofar as they are successful at minimizing the harmful effects of interactions with predators *H*, and maximizing the beneficial effects of interactions with food *B*.

Now prior to perceptual factors coming into existence, all of the relevant causal interactions between an organism *Z* and environmental items *X* would be *direct*, in the sense that they would have to involve actual physical interactions between *Z* and *X*, such as a case of *X* attempting to eat *Z* if *X* is a harmful predator *H*, or a case of *Z* attempting to eat *X* if *X* is a beneficial food item *B*. (For simplicity we assume that sensory factors of touch and taste are absent in the direct cases being discussed, even though they could occur along with such direct causal interactions).

The evolutionary disadvantage of such cases of direct causal interaction is that they permit only a highly restricted range of possibilities of variation during the causal interactions of members of each species. In virtually all cases of a predator X starting to directly causally interact with a prey animal Z by attempting to eat it, a single stereotypical result would ensue, namely that X would end up having completely eaten Z. Thus, for a given species of predator X and prey animal Z, restriction of all of their interactions to such direct causal interactions might quickly result in the extinction of Zs by Xs, plus perhaps the subsequent extinction of Zs as well if Xs are their main food supply.

Now my claim is that perception is a fundamental fitness-enhancing evolutionary mechanism, which works by providing *mediating* causal factors—sensory organs—that permit an organism Z and environmental items X to *indirectly* as well as directly interact with each other. The evolutionary advantage of such mediated interactions is that they potentially permit a *much broader range* of interactions to occur between Zs and Xs, some of which would have a much greater fitness-enhancing potential for Z with respect to X than would any direct interaction with X. For example, the development of sensory organs for seeing and hearing would permit organisms Z to become aware of the approach of predators X—by seeing or hearing them while they are still some distance away—and hence permit a Z to flee from an X before the X ever has a chance to directly eat the Z. Thus, on the view being proposed here, biological cases of perception by an organism Z are all cases of *sensorily mediated causal interaction* between Z and some environmental item X.

If this account is correct, one would predict that sensory mechanisms throughout the biological kingdom would tend to have become increasingly common over millions of years of evolution, because of their fitness-enhancing tendencies for organisms that happen to have acquired them. Also, those mechanisms would increasingly tend to have been used in ways that provide enhanced levels of fitness for the relevant surviving organisms, relative to what they could have achieved purely via non-sensorily mediated direct causal interactions with environmental items. Arguably the actual facts of evolution strongly support the accuracy of these predictions, in that sensory mechanisms are indeed ubiquitous throughout the biological kingdom, plus most extant organisms do make very effective uses of their sensory apparatus, to produce beneficial results for their species, which would be very hard to achieve without those sensory mechanisms.

Here is a summary so far of this evolutionary account of perception. First, on this evolutionary view perception must be distinguished from sensory input, in that perception involves the whole sensorily mediated interaction between X and Z rather than just the input part of it. Second, particular perceptual episodes as such are not inherently information-providing, or otherwise epistemically or functionally positive or correct, because the fitness-promoting aspects of sensorily mediated biological perception pertain only to the overall wide range of evolutionary variation made possible by sensory mechanisms, rather than to any particular cases or episodes of perception. Hence, in applying perceptual concepts to purely mechanical, non-biological devices (see Section ‘Applying reflexive perceptual concepts to non-biological mechanisms’), there need be no inherent requirement that sensors must improve the functioning of the relevant devices, nor that mechanistic perception, when achieved, would necessarily be correct, or necessarily even qualify as an information-processing activity.

The reflexive theory of perception

The preceding discussion of the evolutionary role of perception, though brief, provides enough structure and relevant factors to more precisely characterize the relevant theory of perception—to be called the *reflexive* theory of perception—as follows. To begin, the view takes *causal interaction* between an environmental item X and an organism Z as being a basic factor in evolutionary change. A causal interaction is a closed causal loop in which, in the simplest and most direct cases, X causes Z to cause some corresponding change in X itself. As a very simple example, if predator X bites down on some part of organism Z's anatomy, then X's teeth would penetrate Z's flesh to the point at which the elastic reaction of Z's tissue to X's pressure would be equal and opposite to the force applied, so that X's biting causes Z to cause X's biting to be resisted by the opposite elastic force reactively produced by Z.

Now the claim is that no such direct causal interaction cases are perceptual at all; instead, genuine perceptual cases always involve *sensory mediation*. If organism Z has an array of sensory organs z_1, z_2, \dots, z_n for sight, hearing, touch etc, then perceptual cases are, in the simplest instances, ones in which an environmental item X causes such a sensory organ z_i to cause Z to cause some corresponding change in X itself. This is still an *interactive* or *reflexive* chain, in which X causes Z to cause a change in X itself, but it is mediated by the closed causal chain being routed through Z's sense-organ z_i . Recall that the evolutionary fitness reason as to why such sensory mediation is essential to perception is because of the greatly increased flexibility and range of kinds of causal interactions between X and Z that such mediations make possible.

But there is one additional causal element which can even further increase the flexibility and range of potential causal interactions between X and Z, and hence which should also be included in a reflexive theory of perception. Sometimes there would be significant additional evolutionary advantages to having *deferred* rather than *immediate* reflexive perceptual causal loops between X and Z, so that, rather than X causing sense organ z_i to cause Z to immediately cause some change in X, instead it only causes Z to acquire a *disposition* to cause some change in X. For example, if a predator Z saw a prey animal X on the far horizon, then it would be unable to immediately eat X because of the distance between them. But by Z being capable of sensorily acquiring an X-caused *disposition* to eat X, which would be behaviorally manifested when Z was close enough to X, members of Z's species would maximize their opportunities to eat members of X's species.

Thus an X-related disposition is a disposition whose manifestation would involve some causal interaction with X. However, causal interaction need not involve actual physical contact between X and Z, because, for example, an X-caused disposition for Z to flee from predator X could, when manifested, also have a causal effect on X, in that the fleeing prevents X from eating Z. Depriving a predator X of food has a legitimate causal effect on X, just as depriving someone of oxygen can genuinely cause them to die.

Also, there is no reason why some such reflexive, X-related dispositions should not be immediately manifested when caused by X, so that dispositional causality is flexible enough to permit both immediate, and deferred, reflexive causality between X and Z. In addition, as with any dispositional form of causality, if for some reason the appropriate manifestation conditions for a disposition are not actualized, then no

actual X-related behavior would be produced by Z. Thus a reflexive dispositional theory of perception is not wedded to a simplistic behavioristic model of interactive causality (Dilworth, 2005a).

One further issue should be mentioned. As with any causal theory of perception, the present theory needn't be committed to any particular view of the nature of causality, or of its relata. So when X causes Z to acquire X-related dispositions, it might strictly be properties F of X that cause Z to acquire F-related dispositions toward X, or X-related events that cause corresponding event-related dispositions in Z toward X, and so on.

With all of those preliminaries taken care of, a reasonably precise statement of the reflexive theory of perception could be given as follows. An organism Z perceives an object, property, event etc. X just in case X causes a sense-organ z_i of Z to cause Z to acquire an X-related disposition. Or, assuming transitivity of causality for sensorily mediated perception, Z perceives X just in case X causes Z to acquire an X-related disposition.

As for the issue of perceptual intentionality or aboutness, the reflexive theory automatically supports at least one very basic kind of intentionality. As one might expect, on the reflexive theory a perception *refers to*, or *is about*, the worldly item X that both causes the perceptual state, and to which item X a reflexive disposition is acquired. This point will also be used in Section 'Reflexive versus pure input views of perception' to show that the theory has a significant advantage over purely input-based causal theories, in that standard indeterminacy problems due to input causal chains can be disambiguated via dispositional output causality.

In the rest of the paper this same reflexive theory will be applied to the analysis of purely mechanistic, non-biological perception.

Applying reflexive perceptual concepts to non-biological mechanisms

Enough groundwork has now been laid to permit an initial exploratory application of perceptual concepts to non-biological mechanisms. Initially a straightforward application, with examples, of the preceding reflexive theory of perception to mechanisms will be provided, followed by discussion in succeeding sections of possible objections etc. A mechanical device or system capable of perception will indifferently be described as a perceptual system, as one that perceives the relevant worldly items, and similar cognates.

A primary distinction to be made is that between mechanical devices Z which include modules having a sensory function—but which devices Z are not themselves perceptual systems—from those devices that in addition have genuinely perceptual capacities. For example, closed-circuit television (CCTV) systems include video cameras, plus various kinds of additional equipment, such as recording equipment for passive storage of the video images, or optical character recognition (OCR) equipment for recognizing license plate numbers on cars, and so on. Clearly the cameras in such integrated systems Z qualify as sensors or sensory devices S, which function as transducers for visual information, which is then used for various additional purposes. Our question concerns what else, if anything, is required for the relevant integrated systems Z to also count as perceptual systems, or as systems that can perceive the relevant worldly items.

According to the biologically-derived reflexive theory of perception, such a system *Z* perceives the item *X*, which causes *Z*'s sensor *S* to cause certain effects in *Z* itself, just in case those effects in *Z* include one or more *X*-related dispositions. Thus a passive CCTV system that records video information for later human analysis would not count as a perceptual system, because it itself acquires no dispositions toward *X*. Nevertheless, its human operator could use such a passive CCTV system as a sensory extension of her own perceptual capacities, in that an extended reflexive chain of causality exists in which *X* causes the sensor *S* to record the information on videotape *V*, which recorded information *V* causes the human, via her own sense-organs, to acquire some disposition toward *X*. Hence in this case, sensorily assisted human perception of *X* likely occurred, but no purely mechanistic perception.

On the other hand, in the case of a CCTV system *Z* that includes some optical character recognition (OCR) equipment, which is used to read the license plate on speeding vehicles on a stretch of highway, such a system might become a bona fide perceptual system, if some further method is provided by which *Z* can look up the address of the driver *X* having the relevant license plate, and then automatically generate a mailing, which sends a notice of the relevant fine for speeding to the owner *X* of the vehicle. In the case of such an integrated system, *the system itself*, and not just its human operators or creators, acquires an (immediately manifested) disposition toward the owner *X* of the relevant speeding vehicle, and hence it counts as having *perceived* the owner's speeding vehicle, in addition to having acquired some sensory information about it. Such cases of mechanistic, genuinely perceptual systems involving fully automated serving of legal summons to persons guilty of civil infractions already exist, and are becoming increasingly common.

To be sure, such relatively impersonal, bureaucratic systems do not perhaps provide the best intuitive examples of genuine cases of perception. An intuitively better example might be an art gallery which includes a security system, which both sensorily detects the presence of an intruder *X* at night, and which then imprisons *X* inside the gallery with a system of steel shutters until the police can arrive to remove him. The physical immediacy of the security system's sensorily caused reflexive disposition toward the intruder *X* is more obviously a genuinely perceptual case than the more bureaucratic examples, and one which significantly contrasts with a purely passive sensory security system, which just records the intruder on videotape without making any attempt to apprehend him.

Here is another kind of example that equally sharply distinguishes sensory devices from perceptual systems. In the near future it is likely that at least partially robotic cars or vehicles will be produced, which can carry out some of their maintenance tasks without any human intervention being required. Now current vehicles already include simple sensors or gauges for levels of oil, gas or brake fluid, but these typically rely on their human users to take appropriate action, and so intuitively do not themselves count as being genuine perceptual systems. However, if a robotic car was designed so that when it sensed a low oil level, it automatically drove into a garage and refilled itself with sufficient oil, without any human intervention, it would be hard to deny that this was a case of genuine perception by the car of the low oil level, since it clearly had acquired a disposition toward the low oil level which it manifested by refilling the oil reservoir, which reflexive pattern of causation satisfies the reflexive theory of perception as applied to such cases.

On the other hand, one can also envisage such matters being automatically handled by a division of labor, as follows. Cars would include appropriate arrays of

sensors of internal conditions such as low oil or low gas levels, but the vehicle itself would not read or use them. Instead, the garage or filling station would automatically check the standardized array of sensors in the car when the vehicle enters the garage, and carry out whatever corrective actions would be required, such as topping up the gas tank or refilling the oil reservoir. (Many garages already have diagnostic equipment for reading onboard computers that store sensory information about the vehicle, so that including appropriate corrective actions in such systems would not be difficult). In such cases, it would be the *automated garage equipment*, rather than the car itself, which would genuinely perceive the sensory states of the car, in that it is the automated equipment rather than the car that acquires immediately manifested dispositions toward the cause of the relevant sensory states of the car.

To conclude this series of examples of mechanistic perception with a relatively simple and low-tech case, many thermostatically controlled devices such as furnaces and refrigerators would also qualify as devices capable of perception on the reflexive view, in that, for instance, the thermostat in a furnace could sense when the ambient temperature is too low, and activate a temperature-related disposition by the furnace to raise that temperature by turning itself on. In such cases the relevant causality of the sensors and the reflexive causal loop is transparently simple, but on the reflexive theory such cases are just as much genuine cases of perception as more complex or higher level kinds of perception. Indeed, this is exactly what one should expect, given the present reflexive account of perception as a primary, evolutionary fitness-enhancing mechanism, since the relevant kinds of sensorily mediated reflexive causal interactions must be applicable to even the simplest and most low level kinds of interactions of an organism with its environment.

To sum up this section, it has been shown that, in spite of the initially biological and evolutionary origins and credentials of the reflexive theory of perception, it can also provide a very straightforward, purely naturalistic account—which is often intuitively convincing as well—of purely mechanistic cases of both sensory but non-perceptual mechanical states on the one hand, and both sensory and genuinely perceptual mechanical states on the other hand. As in evolutionary cases, the critical difference between merely acquiring sensory data, and genuine perception, is that perception requires some appropriate interaction—or at least a dispositional readiness to interact—with the cause of the perception, in conformity with the reflexive theory of perception.

Objections, and complementary versus competing theories of perception

Some potential objections and competing theories will be considered here and in the following sections. To begin, probably the most common or widespread conception of perception is of a process in which *information* about worldly items is causally received via one's sense-organs, with no additional behavioral effects being required. Also, this informational conception probably retains its intuitive force, even though attempts to make such a view naturalistically precise—as in nomic covariation or informational semantics views such as those of Dretske (1983)—have not been entirely successful. According to such purely input causality conceptions, perception as such need not involve any behavioral or action-related aspects at all. Hence the most characteristic feature of the reflexive theory—its claim that reflexive dispositions are involved in all

genuine perceptual cases—is also perhaps the most suspect feature of the theory, from the point of view of pure input causality theorists.

However, before we assume that the two kinds of theories must be incompatible, so that we have no choice but to set up a theoretical fight to the death, the following interesting compatibilist theoretical possibility should be investigated. It might be possible to argue that *the proper domain of each theory is disjoint from that of the other*, as follows. Clearly both commonsense conceptions of perception, and more precise nomic covariation accounts, are primarily oriented to the explanation of relatively high level kinds of perception, that is, to cases in which relatively sophisticated perceivers, such as human beings, receive semantically or conceptually structured information about the world. Call this kind of perception '*informational perception*'.

However, equally clearly the evolutionarily based, reflexive theory of perception as described here is primarily concerned to explain a quite different set of phenomena, namely the ubiquitous causal interactions of organisms of all kinds through millions of years of the evolutionary process, where almost all of the relevant organisms are cognitively or conceptually completely incapable of processing any high level semantic information whatsoever. Also, the accounts given here of purely mechanistic perception, in terms of relatively low level causal interactions, may also be completely explicable without any consideration of higher level kinds of informational perception.

Thus, at least initially, a rapprochement may be possible between the two approaches, in which sense-organs are viewed as having two quite distinctive roles in two distinct kinds of perception. First, conceptually or semantically sophisticated perceivers are able to use their sense organs to receive semantic information about the world. Such perceivers then rationally decide how to act on the basis of the information perceived. So it might be argued that sophisticated perceivers, engaged in high level kinds of perception, are able to conceptually separate informational reception from rational decision-making and action, so that a purely input causality concept of perception is the right one for such informational perceptual cases.

But at the same time, a quite distinct low level kind of use of sense organs might also be possible, in which the senses act purely as raw causal transducers, with no informational transfer whatsoever. In such cases, the low level cognitive utility of such sensory processing, since it lacks any informational content, must instead be provided in some other way. Here, the argument continues, the reflexive theory could become relevant, by describing a simple interactive causal mechanism through which even very simple organisms could, via evolutionary processes, potentially achieve evolutionary success without needing any abilities to process semantic or cognitive information. Thus the overall picture is one in which low level perception is devoid of information, and causally interactive to compensate for this lack, while high level informational perception is a more articulated structure in which informational processing is separated from rational decision making.

Indeed, this bifurcated, compatibilist perceptual picture is such an initially attractive one that I think that it deserves to be treated as the default or fallback view of the current overall perceptual situation, at this stage of our theoretical understanding of perception and intentionality. Or in other words, it is the view which, at least initially, seems to best cover all of the available evidence, and which could only be dislodged or superseded by strenuous attempts to show e.g. that high level informational kinds of perception are themselves best explained in reflexive

perceptual terms, or that pure input causality perceptual views are theoretically completely unworkable. (As it happens, elsewhere I have argued for both of these more strenuous views, in Dilworth (2004, 2005a, b, c), but this does not diminish the attractions of the compatibilist view for those who remain unconvinced).

This ecumenical, compatibilist view is particularly attractive as applied to the current issues of mechanistic perception. The commonsense intuition is that, strictly speaking, no purely mechanical device could genuinely *perceive* anything. This intuition could now at least partly be explained as the intuition that no machines—or at least, no machines of the currently available levels of sophistication—are capable of *informational* perception, nor of comparable high level abilities such as that of rational decision-making, so that their perceptual abilities are confined to low level, informationally devoid kinds of reflexive perception. But I have not attempted to argue that machines are capable of any kind of informational perception, but only of reflexive perception, so I can, at least for initial purposes, happily accept that my arguments pertain only to such non-informational, reflexive kinds of mechanistic perception.

Such a compatibilist position presumably would be particularly congenial to those, such as John Searle (1983, 1992, 1998) who also wish to closely associate intentionality and meaningful perceptual content with consciousness and conscious perceptual experience. Clearly most biological organisms, such as amoebae and earthworms, are too simple to be capable of consciousness, while simple mechanical devices such as thermostatically controlled furnaces are not conscious either, so something like the above compatibilist view of perception, offering a different analysis for simpler non-conscious forms of perception, would presumably be unavoidable and hence attractive to such thinkers.

Enactive and sensorimotor views of perception

Another kind of high level view of perception is offered by those, such as Noë (2004) and Hurley (1998), who view perception as being closely associated with action. On such 'enactive' views, roughly speaking, conscious perception of worldly items is closely associated with a mastery of the main *sensorimotor contingencies* concerning that item, including knowing how to interact with it in various practical ways. On such a view, perceiving an item X, including conscious perception of it, is not a distinct cognitive act, but instead it is simply a part or aspect of skilled interactions with the world.

Perhaps it is clear enough that such 'enactive' views of perception, far from being competitors of the reflexive view, are potentially closely allied with it. Indeed, the reflexive theory could be regarded as providing some previously missing low-level causal structures, which would be an essential part of a complete enactive theory of perception. Enactive theorists have up to now largely concentrated on high level conscious perception, but again, an evolutionary account of perception must also be able to explain very low level perceptual activities of organisms, which organisms may completely lack anything we would recognize as practical knowledge, or as knowing how to interact with the world.

Issues concerning mechanistic perception throw such currently missing features of enactive theories into high relief, in that it is obvious that a thermostatically

controlled furnace completely lacks any practical knowledge of sensorimotor contingencies associated with ambient temperatures. Equally, an art gallery security system has no practical knowledge of how best to deal with intruders. But both systems are nevertheless capable of carrying out or 'enacting' very low level kinds of perceptual skills. Arguably, the reflexive theory of perception has the best available account of how simple mechanical systems are able to do this, in spite of their extremely limited causal structures. Thus in the case of enactive theories of perception, a compatibilist view of their relations to the reflexive theory is likely not only a useful first approximation, but probably also part of the ultimate theoretical truth about the nature of perception generally.

Reflexive versus pure input views of perception

Recall from Section 'Objections, and complementary versus competing theories of perception' that probably the most common or widespread conception of perception is of a process in which information about worldly items is causally received via one's sense-organs, with no additional behavioral effects being required, which amounts to a pure causal input view of perception. The informational semantics of Dretske (1981) etc. provides one well-known naturalistic attempt to make that conception more precise. Also recall that the most characteristic feature of the reflexive theory, namely its claim that reflexive dispositions are involved in all genuine perceptual cases, seems to be in direct conflict with such pure input views.

However, as an initial point, since the reflexive view claims that perception involves behavioral *dispositions* rather than actual behavior, it is compatible at least with the folk psychological idea that perception can occur without any overt action having to immediately ensue. One could acquire a behavioral disposition whose manifestation conditions are not currently actual, so that reflexive perception could occur without any immediate behavioral manifestation. Also, since this could be the case in an arbitrarily high percentage of actual perceptual cases, mere anecdotal observation of perceivers and their immediate behavior would equally support either a pure input view or a reflexive view of perception. For the acquisition of a dispositional state of readiness during reflexive perception need have no immediate behavioral consequences or symptoms at all.

A more specific concern is the following. Couldn't someone be in a state of paralysis, whether short-term or permanent, so that they are unable to do anything, but nevertheless be able to perceive things?¹ My answer is yes, they could, because this situation is compatible with their acquiring genuine dispositions during their paralysis. To be sure, as long as they are paralyzed, there could not be any manifestation of those acquired dispositions, but all dispositions are such that they are only manifested under appropriate circumstances. For example, a sample of salt has the dispositional property of being soluble in water, even if it happens to be located on a waterless planet, which situation precludes it from ever actually dissolving.

But there is also a significant advantage to the reflexive theory's dispositional approach, as compared to that of pure input views, as follows. A fundamental problem that has bedeviled causal theories of perception is that of the perceptual status of items in *causal chains* (Coates, 2000; Dilworth, 2005b). For example, when

¹ My thanks to an anonymous referee for a query of this kind.

one sees the red surface of an object, there is an unbroken causal chain of nomic dependencies from the red surface of the object to the eyes of the perceiver. Hence the red light that is immediately incident upon the perceiver's retinas nomically covaries with the resulting information in the optic nerves just as much as it covaries with the surface properties of the red object. But what, then, makes it the case that the person sees the red surface of the object, rather than e.g. instead seeing the light that is incident upon their retinas? Or in other words, how can a pure input view of perception explain how one particular item in such a causal chain has a privileged status as the actual item that is perceived, when the nomic covariance regularities are exactly the same for all of the items in the causal chain?

Such issues have led Dretske and other informational semanticists to conclude that indication of, i.e., nomic covariance with, a worldly property *F* by itself only provides a necessary, but not a sufficient, condition of perception or perceptual representation of *F*. Dretske now holds that perceptual representation of a property *F* occurs only when a system has the teleological *function* of indicating or providing information about that property, which function it possesses because of its evolutionary history (Dretske, 1995; see also Millikan, 1984; Price, 1998, 2001).

However, such an account is completely useless for explaining machine or non-biological perception, because machines do not have evolutionary histories. Fortunately, the reflexive theory has a straightforward, *non-teleological* solution to such causal chain problems. Though input causation via sense-organs cannot itself uniquely specify a particular item in a causal chain, the dispositions acquired by the perceiver can. Thus for instance, it is because a red object causes a person to acquire dispositions *toward that red object itself*, rather than toward light incident upon the retina, that the relevant perception counts as perception of the red object rather than of incident light (Dilworth, 2005b).

This point can readily be applied to cases of mechanistic perception as well. For example, the art gallery intruder alert system presumably involves a chain of causation in which the movement, sound, or appearance of an intruder triggers a relay or other device to close a circuit, which then closes the shutters in the art gallery. But as far as the detection devices themselves are concerned—the sensors for sound, movement, or visual appearance—they equally detect the intruder, the light or sound incident upon the visual or auditory sensor, and so on. Thus it is only the relevant behavioral perceptual dispositions acquired by such systems that determine what the system counts as perceiving.

For example, if the triggering of the auditory sensor solely had the effect of causing a noise-cancelling device to operate, which emitted compensating sound waves to cancel out the detected sounds, then clearly the system is not perceiving intruders as such, but only ambient sounds, however caused, and the system would be using its perceptual skills to minimize the noise level in the art gallery rather than to detect and trap intruders. But those very significant differences in item perceived plus overall causal functioning are entirely the result of dispositional differences in output causality, since the input, sensor-related causality could be identical in each case.

Misperception issues for machines

Another standard problem with nomic covariation or 'indication' pure input explanations of perception is that they are unable to adequately account for the

possibility of *misperception* or perceptual misrepresentation. For if perception of a property F is explained solely in terms of nomic covariation of a sensory state with the presence of property F, all genuine perception of F would be automatically true or correct. In Dretske's case, this is another reason as to why he has now moved to an account of perception as indication plus its being the *proper or natural function* of a perceptual system to indicate a property. For a system can have a proper function because of its evolutionary history even if, on a particular occasion, it fails to function adequately and hence produces misperception (Dretske, 1995).

But again, such an account is useless to explain machine misperception, since machines have no evolutionary histories and hence no proper or natural functions in Dretske's sense. Yet arguably mechanistic perception must, if it is to be genuine, also allow for the possibility of misperception. Here are some *prima facie* examples of machine misperception, using adapted forms of the previous examples of machine perception from Section 'Objections, and complementary versus competing theories of perception'. First, the system that uses CCTV and OCR equipment to read license plates of speeders, and then sends them notices of their fines, starts sending the fines notices to the owners of all cars whose license plates it photographs, whether or not they were speeding. Second, the art gallery security system develops a hypersensitivity to any internal sensory disturbance, so that even the accidental overnight entry of a fly to the gallery causes the steel shutters to close and the police to be informed of the presence of an intruder. Third, the robotic car that monitors its own gas and oil levels, and automatically drives into a garage and replenishes them when needed, starts to replenish them in this manner even when they are already nearly full. Fourth, the thermostatically controlled furnace, which is programmed to turn on the heat at 70 degrees and off at 72 degrees, begins to let the temperature drop to 65 degrees before it turns on the heat.

Now as an initial point about these examples, in spite of their great variety it is possible that none of them involve any purely input sensory change or malfunction from cases in which the relevant systems correctly perceived the relevant items. The misperceptions might all be due to changes in the subsequent processing of the unchanged output characteristics of the relevant sensors. For example, in the art gallery case a sensitive sound sensor might have been used in the basic design of the system, whose transduced output would include electrical signals corresponding to any sounds, from the loudest to the quietest. When correctly perceiving human intruders, the system triggers shutter-closing etc. only when the mix of output signals includes some corresponding to relatively moderate to loud sounds in the art gallery. But in the fly misperception case—in which the system misperceives a fly as a human intruder—the system ends up also allowing some relatively low ambient sound levels in the gallery, such as the buzzing of a fly, to trigger its shutter-closing etc. But these are changes in the *output* causality of the system, not changes in its purely sensor-related input causality. And similar analyses could be provided of the other examples as well.

This point concerning the possible unchanged input causality, but changed output causality in such examples is important, in that it extends the conclusion of the previous section that it is *dispositional output causality*, not input causality, which is the predominant factor in determining which item X is perceived by a system that includes some sensory input. The current point is that output causality also is the predominant factor in determining whether *misperception* rather than correct perception of an item X occurs, given that it is X that is being perceived by the system.

Thus output causality not only determines what item X is perceived by a mechanical system, but also whether the system correctly or incorrectly perceives it.

Conclusion

First, a reminder of the limited scope of the argument in this paper. It has not been argued that all perception, including the most sophisticated human kinds, can best be explained by the reflexive theory of perception and its integral use of dispositional output causality factors, but only that the theory provides the best available account of simple mechanistic cases of perception, as discussed here. Thus the compatibilist option, as discussed in Section ‘Objections, and complementary versus competing theories of perception’, of distinguishing two quite distinct concepts of perception—one being a sophisticated, high level and pure informational input kind, and the other a low level, causally interactive theory along the lines of the reflexive theory—still remains an open option at this stage, and one that is bound to be attractive to many. So in conclusion, it will be useful to again address the concern that simple mechanistic cases of perception must lack critical factors that are present in normal human perception.

My basic strategy in dealing with such concerns has not been confined to acknowledging the compatibilist option, though that is indeed one legitimate possibility. More broadly, I have argued that there is a basic kind of perception that is indispensable in understanding the role of perception as a primary mechanism underlying evolution (see Section ‘The fitness-enhancing evolutionary role of perception via sensory organs’). This kind of perception is indeed a fundamentally low level activity, and it should sharply be distinguished from any higher level kinds of cognitive processing that have, via millions of years of evolutionary adaptation, also become prevalent in humans and other higher animal species. Thus for instance, an art gallery security system can, as argued here, genuinely perceive or misperceive intruders. But this claim, I would argue, is a very low level claim that does not go beyond the relatively simple interactive causal structures covered by the reflexive theory of perception.

Here are some of the higher cognitive factors lacking in such a system. The system, when perceiving an intruder, does not have any introspective access to its perception—it does not know, or believe, that it has detected an intruder. Nor is it clear that it processes any intruder-related information. Nor does it, in any significant higher sense, *interpret* its input sensory data as indicating that an intruder is present. Nor does it believe that it is legitimate for it to temporarily imprison an intruder because its owner has been officially permitted by the local police force so to confine illegal intruders. Nor, having perceived an intruder or misperceived a fly as an intruder, can it deliberate about what is the best course of action for it to take, all things considered, as in traditional models of rationality. And so on. Instead of all of these, it simply acquires some immediately manifested behavioral dispositions toward the item that, via sensory mediation, it thereby perceives.

Hence I would argue that, once a complete catalog of this kind is assembled—of all of the higher features of cognition that are completely lacking in such a simple causal system—we would have also assembled a complete catalog of all of the reasons for wishing to deny that such a system genuinely perceives in the full human sense. Or in other words, it is not basic perception as such that such intuitive

concerns address, but instead the whole integrated package of sophisticated human abilities that *includes* perceptual abilities. My strategy in response is basically a standard cognitive science response of cognitively separating out some relatively minimal, purely low level perceptual factors in this mix from the higher cognitive factors, and then agreeing with those who are concerned that art gallery security systems completely lack any such higher cognitive factors or abilities. Thus this kind of perception is indeed a modular and low level cognitive activity, as has been argued by Fodor (1983) for any kinds of perception.

The compatibilist option discussed earlier would concern whether there is, in addition to the purely low level reflexive concept of perception just discussed, also a distinctive kind of purely input-based, mid-range perception, applicable to humans but not to simple mechanistic systems, and involving such issues as perceptual information-processing, distinctions of conceptual versus non-conceptual perceptual content, kinds of perceptual intentionality not analyzable in reflexive terms, and so on. A significant value of the study of mechanistic perception, as well as the related study of the evolutionary role of perception, is that they enable us to clearly separate out, perhaps for the first time, these two distinctive constellations of issues about perception. Then we can investigate interesting philosophical strategies for either explaining all such issues in terms of low level perception, or for more perspicuously separating out their distinctive cognitive roles.

Acknowledgments The author thanks the Editor, and an anonymous referee, for very helpful comments.

References

- Cantoni, V. (Ed.) (2005). *Human and machine perception: Communication, interaction, and integration*. New Jersey: World Scientific Publishing Co.
- Coates, P. (2000). Deviant causal chains and hallucinations. *Philosophical Quarterly*, 50, 320–331.
- Cummins, R. (1996). *Representations, targets and attitudes*. Cambridge, MA: MIT Press.
- Dilworth, J. (2004). Naturalized perception without information. *The Journal of Mind and Behavior*, 25, 349–368.
- Dilworth, J. (2005a). The reflexive theory of perception. *Behavior and Philosophy*, 33, 17–40.
- Dilworth, J. (2005b). Perceptual causality problems reflexively resolved. *Acta Analytica*, 20, 11–31.
- Dilworth, J. (2005c). A naturalistic, reflexive dispositional approach to perception. *The Southern Journal of Philosophy*, 43, 583–601.
- Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.
- Dretske, F. (1995). *Naturalizing the mind*. Cambridge, MA: MIT Press.
- Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Fodor, J. (1990). *A theory of content and other essays*. Cambridge, MA: MIT Press.
- Hurley, S. (1998). *Consciousness in action*. Harvard, MA: Harvard University Press.
- Millikan, R. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.
- Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.
- Price, C. (1998). Function, perception and normal causal chains. *Philosophical Studies*, 89, 31–51.
- Price, C. (2001). *Functions in mind*. New York: Oxford University Press.
- Searle, J. (1983). *Intentionality*. New York: Cambridge University Press.
- Searle, J. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press.
- Searle, J. (1998). *Mind, language and society*. Philadelphia, PA: Basic Books.