

Perceptual Intentionality

John R. Searle

Perception looms large in our philosophical tradition. There are a number of reasons for this, but two stand out. First, for the three centuries after Descartes epistemology, especially skeptical epistemology, was at the center of Western philosophy; and of course there is really no way to face these issues without advancing a theory of perception. A second reason is that philosophy is generally concerned with the relation of human beings and reality, and the primary conscious medium by which humans relate to the real world is the medium of perceptual experience.

Unfortunately, the entire philosophical tradition rests on a disastrous mistake. This mistake in turn rests on a failure to understand the intentionality of perception. So the article will have two aims: I want to expose and correct what I will call the Bad Argument, which afflicts just about all of the major philosophers after Descartes, in fact all known to me; and second I want to provide at least the beginnings of an account of the intentionality of perceptual experience.

1 The Bad Argument

The account of perception that I will present is a form of *direct realism*, according to which, in a typical perceptual experience, we are directly aware of objects and states of affairs around us. "Directly" means that the relationship is unmediated. Looking at the desk in front of me is not like seeing the desk on television or reflected in a mirror. In those cases I do not see the desk directly; right now I see it directly. Traditionally direct realism is opposed to representative realism, according to which there is a real world out there; but we cannot perceive it directly. We only perceive representations of it in the form of our impressions, ideas, or sense data. Famous representative realists are Descartes and Locke.

Because direct realism is so obviously consistent with both our experience and with common sense, Why has it been so frequently denied? Here is the amazing fact. None of the Great Philosophers in the

tradition that I am familiar with – and by Great Philosophers, I mean Descartes, Locke, Berkeley, Hume, Leibniz, Spinoza, Kant, and I would have to add Mill and Hegel – were direct realists. They all, in one way or another, accepted an argument that goes back at least to the 17th century, and for all I know it may even go all the way back to the Greeks. The standard argument against direct realism is called the Argument from Illusion. And here is how it goes: I said that I see the desk in front of me, but suppose I were having a hallucination. I could be having an experience exactly like this, completely indiscriminable from this one with exactly the same content as this one, and be having a total hallucination. You could have the same experience without there being a desk on the other end of the experience. But now comes a crucial step. The character of the experience in the hallucination case and the character in the veridical case, the bad case and the good case respectively, is the same; so any analysis of one has to be applied to the other. But in the bad case, though one is not aware of a desk, one is certainly *aware of* something. One is conscious of something, and, at least, in some sense, one “sees” something. That something cannot be a material object, because there is no material object there in the case of the hallucination. Yet I am aware of something. Let us give a name for such somethings; they were called “ideas” by Berkeley, Descartes, and Locke, “impressions” by Hume, and came to be called “sense data” in the 20th century. So I will stick with “sense data” and say we are aware of a sense datum. A sense datum is an entirely mind dependent, ontologically subjective entity. Now, by the principle that both the good case and the bad case should receive the same analysis, it follows that in the good case I am not aware of a material object but only of sense data. But now it seems to follow that in all experiences I am aware only of sense data, not of mind independent material objects. And the question then arises, What is the relationships between the sense data that I do see and the material object that apparently I do not see? This argument in various forms survives right to the present day. What is wrong with the argument? On the surface, at least, it rests on a simple fallacy of ambiguity. The expression “aware of”, “conscious of”, and even “sees” in this argument, are ambiguous. The ambiguity can be illustrated by using a very simple and unproblematic example. If I push my hand very hard against the top of the desk, I am aware of the surface of the desk. This is the intentionality sense of “aware of” that has the desk as the intentional object. But it is obvious that I am also aware of a painful sensation in my hand

assuming I push hard enough. So which am I aware of: the painful sensation in my hand or the top of the desk? After all there is only one object there and one experience there. Which one is the genuine case of awareness? The answer is that the expression of "aware of" is used in two different senses. In the intentionality sense the desk is the object of the awareness, and I am aware of the desk. But there is another sense being exhibited here, where I am "aware of" a painful sensation; and that case is not one of intentionality, because the awareness and the painful sensation are identical. This is a crucial point. Where intentionality is concerned the sensation is not identical with the object, but there is another constitutive sense in which the awareness and the thing one is aware of are identical. The proof then that there two different senses of "aware of" are being used here is that the semantics are different. In the intentionality sense, the subject S has an awareness A of object O implies A is not identical with O; but in the constitutive sense where subject S has an awareness of entity O, A is identical with O, the painful sensation and the awareness are identical. Now let us apply this to the famous Argument from Illusion that we considered earlier. In the sense in which I am aware of an object when I look at the desk, the intentionality sense, in that sense when I have a hallucination I am not aware of anything. There is nothing there; hence I could not be aware of anything. Nonetheless, I am having a conscious visual experience and it is tempting, given the way our language works, to erect a noun phrase to stand for that awareness and make it into the object of the verbs of perception. So "aware of", "conscious of", are used in two different senses. We feel immediately hesitant to say that one "sees" anything in the hallucination case, so we are tempted to put sneer quotes around "sees". But what is going on, I hope, is obvious and clear. In every case there is an ambiguity in the crucial phrases "aware of" or "conscious of"; because in the intentionality sense in which I am aware of something when I see it, in the case of the hallucination I am not aware of anything. I have a conscious experience, but that conscious experience is not itself the object of the experience; it is identical with the experience.

Once pointed out that this is such an obvious fallacy it is hard to see how anybody could have made it; but nonetheless there it is and it produced the idea that is common to the Great Philosophers that one does not perceive the world or does not perceive it directly. One perceives only the contents of one's own mind, one's own sense data. Ironically, the argument is repeated by people who think that they are defending

naïve or direct realism. They are called “disjunctivists”, and they accept the validity of the argument but wish to reject its conclusions on the ground that they think the first premise is false. The first premise says the hallucination case and the veridical case have exactly the same content. The experience is indistinguishable for the obvious reason that it has exactly the same content and the same form of awareness. Disjunctivists think that in order to protect naïve realism you have to reject that premise. Thus in order to save naïve realism they reject an obviously true premise. And what I am suggesting here is that you can have naïve realism and still avoid making the fallacy that led to its rejection. To repeat this point: the disjunctivists correctly see that the conclusion of the argument – the denial of direct realism – is false; but they think in order to save naïve realism they have to reject the true premise that the content of the good and the bad cases can be exactly the same. They accept the validity of the argument but think it is unsound, because the first premise is false. On the other hand, my argument is that the first premise is entirely true and the argument is *invalid*, because it rests on a fallacy of ambiguity.



Prague Lecture on *The Intentionality of Visual Perception*, Institute of Philosophy, May 2011

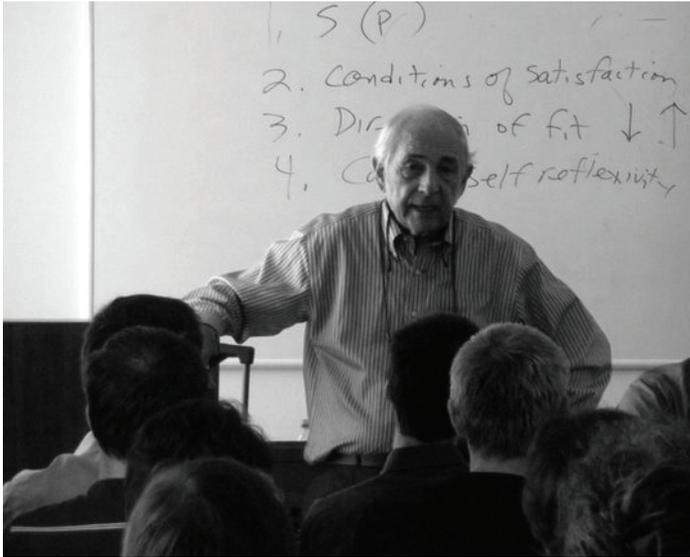
The same fallacy, by the way, afflicts the so-called Argument from Science that says science has shown we can never see objects and states of affairs in the world, but only our own experiences of those objects.

The story goes that the neurobiological events that lead from the sensory receptors to the conscious experience determine that all that we can ever be aware of is the conscious experience. I hope it is obvious that this is the same fallacy: the fact that one can give a causal account of how the conscious experience occurs does not show that one does not see the objects and states of affairs on the other end of the conscious experience. To suppose that is to suppose that the experience itself is the object of perception. And that is the Bad Argument all over again. I am not going to go through the entire history of philosophy to show how the Bad Argument keeps creeping up in Descartes, Locke, Berkeley, Hume, etc. I do not know that Kant ever uses the argument in this form, but he certainly accepted the conclusion.

2 The Intentionality of Perceptual Experience

I am going to assume that the Bad Argument has been refuted and that we can accept direct realism as the accurate account of veridical perception. How does it work? Before answering the question I want to situate the discussion within contemporary and recent philosophy. Traditionally, analytic philosophy is obsessed with truth conditions. To analyze a concept or a sentence is to give its truth conditions. Thus, to take two famous examples, Frege makes his distinction between *Sinn* and *Bedeutung* in his analysis of the truth conditions of identity statements, such as “The Evening Star is identical with the Morning Star”. Again, Russell in his Theory of Descriptions tried to analyze the truth conditions of sentences, such as “the King of France is bald”, which contain an apparent reference to an entity that does not exist. In my book *Intentionality* (1983) I extended this project from sentences to mental states, and I attempt to analyze not just truth conditions but what I call *conditions of satisfaction*, generally; so that I would be able to give the form of the conditions of satisfaction not just of sentences but of beliefs, desires, hopes, intentions, and perceptual experiences. Typically, philosophers do not worry about how the sentence or expression gets its conditions of satisfaction; they just want to know what are the conditions of satisfaction. And the answer, in general, is that in the case of linguistic elements the conditions of satisfaction are imposed by convention, and the form of that imposition is called assigning a meaning to the sentence. The meaning of the sentence is that which determines the conditions of satisfaction, and the sentence has its meaning by the

conventions of the language of which it is a part. Our present task is much more difficult than that. Perceptual experiences with their raw phenomenology do not in general have conventions that fix their conditions of satisfaction. All the same they do determine conditions of satisfaction. Our question is, How does it work?



Prague Lecture on The Intentionality of Visual Perception, Institute of Philosophy, May 2011

The key to understanding perception is to recognize both that it has intentionality in the philosopher's sense, according to which intentionality is simply the directedness or the aboutness of the mind, but it has a special form of intentionality that I call *presentational intentionality*. Perceptual experiences, visual or otherwise, are in the intentionality sense, directed. But the intentionality of visual experience differs from the philosopher's favorites of belief and desire in a number of important respects. Beliefs and desires are typically representations of objects and states of affairs in the world. But when I see something I do not just have a representation; I have a direct presentation of the object. My visual experience, so to speak, reaches right up to the object.

What does it mean exactly to say that the perceptual experience is a presentation and not just a representation? There are a number of features of perceptual presentations. I do not have the space to go into all of them, but I will mention some of the most salient. The most imme-

mediate and obvious feature of perceptual presentations is that they are directly caused by the object itself. The present features of the object I am seeing are the conditions of satisfaction – what makes the visual experience “veridical” – but they are experienced immediately, and they are experienced as causing the very perceptual experience that has the object and the features as the rest of the conditions of satisfaction. So there is a causal self-reflexivity to perceptual experience in that it is only satisfied if the object or state of affair seen causes the experience itself that has those conditions of satisfaction. Memories and prior intentions also have this causally self-reflexive feature, but they are not direct presentations. You are not immediately aware of the object you *remember* in way that you are if you actually *see* it.



Prague Lecture on The Intentionality of Visual Perception, Institute of Philosophy, May 2011

A second feature of the presentational intentionality of visual experience is that it is always of the here and now. Everything you see, touch, hear, you are hearing as existing right here and now. This, as a feature of the phenomenology, remains true even in cases where we know the object is not right here and now. If I look at the star that I know ceased to exist millions of years ago, all the same, I am seeing it as if it were visible right here and now. The third feature of perceptual experiences that differs from features of other more commonly discussed forms of intentionality, such as beliefs, is that they are experienced as non-detachable from the conditions of satisfaction. If I close my eyes and think about objects, I experience certain representations of these objects and I can shuffle them around at will. But if I open my eyes and

look directly at objects in front of me, I cannot shuffle my experiences around at will. They are experienced as directly connected to the object that I am seeing, and this is a feature of the phenomenology that holds even for hallucinations. Even if it is a hallucination and even if I know it is a hallucination, all the same, I do experience it as non-detachable from the object that it is as if I were seeing.

So presentational intentionality differs from such things as beliefs and desires in at least these three crucial respects. It is causal and it is experienced as directly causal, that is to say, it is experienced as directly caused by the object or the state of affairs perceived. Secondly it is always indexical; its conditions of satisfaction are always immediately tied indexically to the occurrence of the experience itself. Every time I see something I see something as existing right here and now, and every time I see an event I see that event as occurring right here and now. Third, it is because of these two features that it is not detachable from the conditions of satisfaction. I cannot shuffle it around as I can with other sorts of mental representations.

But these reflections so far do not answer our question, How does it work? And I now need to explore that question in some detail.

3 The Subjective Visual Field and the Objective Visual Field

If you close your eyes and put your hand over your eyes, you will have an experience which is something like seeing a black area with yellow patches in it. You can try this for yourself. Now, strictly speaking you do not see anything, because your eyes are closed; but you do have conscious experiences which you would naturally describe in the vocabulary that I just used. You do not see anything, because you cannot see anything with your eyes closed. But, nonetheless, you do continue to have conscious visual experiences. This area of your conscious experiences I propose to call the *subjective visual field*. Now open your eyes and the subjective visual field is suddenly filled with everything that constitutes your conscious awareness of the objects and states of affairs in your vicinity. To have a name for what you can perceive I want introduce another technical term, the *objective visual field*, which consists in all of the objects and states of affairs that are visible from your point of view and given your physical condition at any particular time. The objective field is ontologically objective in that it is perceivable by any similarly endowed and similarly situated person. The chair that I see

is part of the objective visual field and it is ontologically objective, because it is accessible to anyone equally. But, in addition, we have found it necessary to postulate the subjective visual field, which consists entirely of the conscious experiences going on in my head. Much of the rest of this article will be about the relations between the subjective visual field and the objective visual field.

The most important thing, I can tell you immediately, is that in the objective visual field everything is seen or at least can be seen; in the subjective visual field nothing is seen nor can be seen. This is not because the entities in the subjective field are invisible, but because they are the seeing of anything. The idea that the entities in the subjective visual field are themselves seen is the basis of the Bad Argument. The awareness itself is erected into the object of awareness; and this denies the intentionality of perception, because the experience is not the intentional object of the experience, it is the experience itself.

The question I want to address now is, How does the raw phenomenology of the visual experience, the actual qualitatively subjective entities in my head, set the conditions of satisfaction – that is how is it that they have the content that they do? Another way to put this question is to ask, What fact about my subjective experiences makes it the case that I am seeing or at least seem to be seeing specific types of objects and states of affairs in the objective visual field?



Prague Lecture on The Intentionality of Visual Perception, Institute of Philosophy, May 2011

This is not a trivial question and I will not fully answer it in the scope of this article, but I can at least suggest some general principles which have to govern any answer.

Two traditional answers to this question both fail. They are resemblance and causation. According to the resemblance theory the visual experience presents red objects because it is itself red. There is a resemblance between the visual experience and the object perceived. Everything is wrong with this answer. To begin with, the visual experience could not possibly itself be colored; and it is not itself perceivable. Causation will not do as an answer either. According to this answer the experience of red has red as its conditions of satisfaction, because it is caused by red things. This will not do, because causation by itself does not carry any intentional content. Suppose that every time I see a red object it causes in me a pain. This does not have the consequence that "red" means pain. Neither resemblance nor causation by themselves are the answer to our question. What is the answer to the question?

4 The Hierarchy of Perception and Basic Perceptual Features

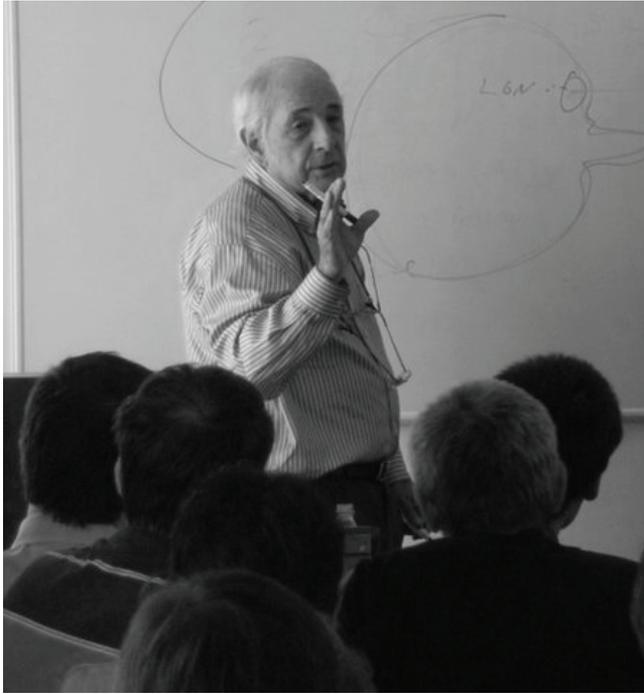
Perception is hierarchical structured in that the top level depends on our ability to see the bottom level until finally you get down to the basic level of things that you can see without seeing anything else by which you can see them. So, take an example, I do not just see colors and shapes, but I see cars, trees, houses; and indeed I do not just see any cars, trees, and houses, I see *my car* in the parking lot, or *my house*. The point about the hierarchical structure is this: in order to see that it is *my car*, I have to see that it is a car of a certain make and age; but in order to see that it is a car of a certain make and age, I have to see that it is a car at all; and in order to see that it is a car at all, I have to see that it has certain color and shape. Color and shape are basic relative to the other features, so being *my car* is less basic than being a car of a certain type; being a car of a certain type is less basic than being a car; and you go down until you finally reach basic perceptual features, such as color, shape, movement, etc. The intuitive idea is that a basic perceptual feature is one you can see without seeing something else by way of which you see it. This cannot be quite the right way to put it, because on this criterion we cannot distinguish color and shape; and maybe the right way to describe the basic feature is to say "colored shapes" or "shaped colors". The intuitive idea is clear: perception is hierarchically structured and hierarchy bottoms out in basic perceptual features; and corresponding to the hierarchy in the objective ontology is a hierarchy

in the subjective visual field. The experience of the higher level features – such as being my car – requires the experience of the bottom level features – such as color and shape.

So far then we have both the notion of basic perceptual features, which are ontologically objective features of the world that are perceivable, and the notion of basic perceptual experiences that have the basic perceptual features as their objects, as their conditions of satisfaction. And we can now narrow our question so that it is much more manageable. How does the raw phenomenology of the basic perceptual experience fix the basic perceptual features as the objects of perception, as their conditions of satisfaction?

A condition of adequacy on our answer to this question is that the connection between the raw phenomenology and the conditions of satisfaction must be internal or essential. That is, for example, it could not be this very experience if it was not an experience of seeing, or at least seeming to see, something red. Now what fact about the experience gives it that intentional content? And remember in answering that neither resemblance nor causation by themselves are going to be enough.

To answer the question I propose to emphasize the fact that our perceptual relations to the external world are causal and experienced as causal throughout. The whole Humean tradition about causation makes it difficult to see that causation is everywhere and is experienced as everywhere. Hume taught us we cannot experience necessary connection. I think we experience it pretty much all day long. Wherever we consciously perceive or engage in action we experience objects causing perceptual experiences in us and ourselves causing bodily movements and other sorts of changes in the world. Causation, to repeat, is everywhere. We live in a sea of causal relations and we constantly experience causal connections. The reason that this experience has red as its conditions of satisfaction is because it is the essence of red, it is part of the very definition of “red”, that it consists, at least in part, in the ability to cause this sort of experience. It is tempting to put this point conceptually, but that would be misleading if it gave us the idea that in order for an animal to be able to perceive red it has to have some conceptual skills. The point rather is this: it is of the very essence of something being red, and hence of the concept of red, that red consists, at least in part, in the ability to cause this sort of experience. What I have said about “red” is, of course, true of all other colors. Their essential feature is this ability to cause these sorts of experiences.



Prague Lecture on The Intentionality of Visual Perception, Institute of Philosophy, May 2011

The analysis is a bit trickier for the so-called primary qualities, but I think it works for them – or at least for the two dimensional ones – as well. Lines and shapes are in part defined in terms of their ability to cause such and such experiences. So a straight line causes this sort of experience. Notice it will not do to say that a straight line is one that “looks like this”, or red is a color that “looks red” under these conditions; because “looks like” and “looks red” have a meaning deriving from “is”. To look red means to look *to be* red, looks *as if it is* red; and that is what we are trying to explain. The fact about it that makes it look as if it is red is what needs explaining, and the explanation I am offering is that red, by its essence, consists in the ability to cause this sort of experience.

Well, this does not give us an analysis of a very rich array of perceptual features. We are down just to the bare bones of basic features. What about the others? I do not have the space to go into detail here.

I am writing a whole book about this. But, just to mention a few of the points: depth or three-dimensionality is not a basic perceptual feature, rather we experience depth because our experience of the basic perceptual features impacts on our nervous system in a way that our Background mastery of the principles of perspective enables us to treat them as carrying three-dimensionality as part of the conditions of satisfaction. As part of the basic experience I see two parallel lines that are getting closer toward the top, but of course I do not just see two parallel lines; I see two railroad tracks receding away from me in the distance. How do I get from the perception of the basic perceptual features to the perception of the distance as the railway tracks move away from me in the distance? The answer is that my Background skills contain a mastery of the principles of perspective, and this enables me to treat the basic two-dimensional features three-dimensionally.

Notice I am not saying we make an inference. We do not need to. We just see distance immediately. I can literally see that the desk is further away from me than the chair. But in order to do that the nervous system must be able to operate on the basic perceptual experiences in a way that imposes the three-dimensional understanding on a visual array in my subjective visual field that can be produced by a two-dimensional stimulus.

What about more complex cases? Suppose I see that this is a California Coastal Redwood Tree, or that the car in the parking lot is my car. Some complex visual phenomena can be defined in terms of the basic features. To learn to identify the Coastal Redwood Tree as a Coastal Redwood is to be able to identify it in terms of its basic features: color, shape, texture, etc. Being a California Coastal Redwood Tree is defined as having a certain combination of basic perceptual features – at least as far as visual experience is concerned. Of course scientists will give a fancy scientific definition in terms of the DNA, but what we are talking now is what I can literally see. The phenomenology of the complex is made out of the phenomenology of the simple.

We also need to explain recognition of previously experienced objects. So, for example, my perception of my car carries much more than just complex perceptual features. The facts about it that make it the case that I perceive it as mine involve the notion of *recognition*. And how does recognition get into the phenomenology? In the case of seeing it as my car, I do not just see it as a car having such and such features; but I see it as identical with a car that I have experienced earlier on count-

less occasions of having such and such features. That is, as far as the phenomenology is concerned, this perceptual experience is simply the latest in a sequence of perceptual experiences all of which are caused by the perception of my car.

5 Conclusion

This article is only really the beginning of an account of how perceptual experiences fix conditions of satisfaction, but I hope the direction is clear. There are three points that I wish to emphasize. One, direct realism is preserved once we abandon the Bad Argument and resist any temptation to resurrect it. Two, perception is obviously intentional; but it has a special type of intentionality that I call presentational, rather than just representational. And third, we can begin to explore how the raw phenomenology of the perceptual experience sets the condition of satisfaction that it does. For an animal to perceive a basic property as F is for it to have a perceptual experience that it experiences as caused by something whose essential character is the ability to cause F. I am not, of course, saying that animals think all of this, nor do humans except for rare philosophers; but this is what is going on in the phenomenology of their experience, and indeed if you watch an animal engaged in complex conscious behavior – a dog digging for a bone, for example, or chasing a cat – you see precisely this combination of the conscious experience of causation and the identification of intentional objects.

*Department of Philosophy
University of California
314 Moses Hall 2390
Berkeley, CA 94720-2390
USA
searle@berkeley.edu*

References

SEARLE, J. R. (1983): *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.