

5

Hearing As

William G. Lycan

The expression of a change of aspect is the expression of a *new* perception and at the same time of the perception's being unchanged.

—Wittgenstein, *Philosophical Investigations*

Ambiguity—rabbit or duck?—is clearly the key to the whole problem of image reading. . . . [S]uch interpretation involves a tentative projection, a trial shot which transforms the image if it turns out to be a hit.

—E. H. Gombrich, *Art and Illusion*

It continues to puzzle me, and I'm sure others, that so little is written by philosophers on aspect perception, a phenomenon both ubiquitous and fascinating. It doesn't surprise me at all that every bit of what has been written is on vision.

One of my purposes in this chapter is to review some views and lessons about "seeing as," and then to see which of those may carry over to aspect perception in hearing. The latter may actually be of more general interest, because, while visual aspect perception is fun, it is best known through gimmicks such as ambiguous figures, while hearing-as plays at least two more central roles in human life.

But first a little groundwork is needed. (1) What is perceived-as is perceived, and more specifically an object of perception; to see X as such and such one must see X. (2) There is a contrast between merely seeing and seeing-as, in that two people may see X and see X qua X, but be seeing X under different aspects. (3) Aspect perception is commonly contrasted with representation; seeing-as is thought to outrun what is visually represented. (4) It may be held that aspect perception is necessarily or at least normally a

conscious phenomenon (in whichever sense of the "c"-word). So: What are the objects of hearing? What are the typical contrasts between merely hearing X and hearing X as such and such, and between auditorily representing X and hearing X under an aspect? And what is the relation between hearing as and consciousness?

5.1. Background Assumptions

1. I shall simply assume that perceiving is a matter of representation.¹ I shall further assume that we are talking of representations hosted and employed by whole persons, not of the merely subpersonal representational contents assigned to elements of our perceptual modules or other specialized brain agencies.

And we must ask, what does hearing characteristically or proprietarily represent? Here I have little of my own to contribute. There is already a prominent contender: O'Callaghan's (2007) carefully defended view that (a) hearing represents *sounds*, which are worldly events rather than objects or properties; and (b) a particular sound is an "event . . . of oscillating or interacting bodies disturbing or setting a surrounding medium into wave motion" (60).

Nowadays nearly every philosopher agrees that sounds are particulars,² and that audible qualities like volume, pitch, and timbre are properties of sounds. There are two immediate contrasts with vision and its objects. First, though it is generally agreed that vision represents properties or properly instances such as color and shape, it is controversial whether vision represents *individual* things or events in the world. Second, as Kania (2010) observes, it is nearly tautologous that what we hear are sounds and sounds are what we hear; but there is no parallel tautology for seeing, except (I note) the slightly strange "What we see are sights."³ One might take this to suggest that sounds

¹ With one small qualification: There is a puzzle about what taste (proper, unaided by smell) represents at the whole-person level (Lycan 2018); one possible solution to the puzzle is to fall back and say that taste represents only in the debased, primitive way set aside by Ramsey (2007, chap. 4).

² A literature that burgeoned in the 2000s rejected the whole idea that sense modalities represent, e.g.: Campbell 2002; Travis 2004; Fish 2009. But for critique, see Siegel 2010a and Schellenberg 2011.

³ An exception was Parnau (1999), who took sounds to be qualities. "Sights" in the ordinary sense are special things, as in "sightseeing," "seeing the sights," but no where more special than at Macbeth's banquet: "You make me strange / Even to the disposition that I owe, / When now I think you can behold such sights, / And keep the natural ruby of your cheeks, / When mine is blanched with fear." The dumbfounded Ross naturally asks, "What sights, my lord?"

are in some way more subjective or more closely mind-related than are the objects of seeing, but if so, that does not show in the current array of theories, for each of them takes sounds to be real phenomena in the environment that affect us physically. (Which is hardly tautologous, though I take it to be fairly obvious.)⁴

Some of the competing ontologies: Casati and Dokic (1994) had contended that sounds are vibration events, with no mention of a medium (notice the implication that there would be sounds even in a perfect vacuum). Before that, Perkins (1983) had identified sounds with “trains of airwaves stretching from the listener’s ear back to the sound-making object throughout the duration of the sound” (172)⁵; Roy Sorensen (2008) also defends a wave theory, as do O’Shaughnessy (2009) and Nudds (2009). Roger Scruton (1997) puts forward a less forthrightly physicalist view: that sounds are “pure” (subjectless) events and, even though they are real and external to perceivers, they are not reducible; their intrinsic properties are limited to ways in which they appear or would appear. (Scruton thinks this is important for the understanding of music, because, he says, it is important to hear music as independent of its source.)

These theorists duke it out with each other on grounds that include the perceived durations and locations of sounds, and intuitions regarding identity and individuation. I shall not take sides, except to assume physicalism about sounds. But I shall insist, following Sorensen (2009) and Farennikova (2013), that *absences* can be perceived as such,⁶ for in listening to music it is sometimes important to hear the silence of a dramatic rest (think of that last rest in the “Hallelujah Chorus”), or the absence of the third from an unexpected open fifth.

⁴ I would, being temperamentally an arch-realist about everything. But I gather the traditional Berkeleyan question about the tree in the forest persists among the folk, even though its counterpart about the tree in the quad does not. (No one contributes to popular culture by asking, “If a tree falls in a forest and no one is there to see it, is it really on the ground?”) So perhaps, *without accepting Berkeleyan idealism generally*, some speakers are tempted to withhold the word “sound” if there were no sentient creature available to hear it. In 1884 (April 5 issue), *Scientific American* declared as fact that “sound is vibration, transmitted to our senses through the mechanism of the ear, and recognized as sound only at our nerve centers. The falling of the tree or any other disturbance will produce vibration of the air. If there be no ears to hear, there will be no sound” (218, quoted in Wikipedia).

⁵ It is astonishing that neither O’Callaghan nor any other recent writer on sounds cites Perkins. *Sensing the World* won the American Philosophical Association’s Johnstone Prize in 1983. It was perhaps the most significant contribution since Reid’s to the dethroning of vision. (“Vision” comes only as chapter 6, following “Hearing,” which follows “Feeling Heat and Cold.”)

⁶ This creates the obvious problem for perceptual psychosemantics, for perceiving an absence cannot be a state caused by the object perceived, nor is it a case of perceiving an object that is not really there. The most straightforward approach here is to hold that absences can themselves be causes, though that will constrain theories of causation.

2. We hear sound-events, then, and we do so in virtue of our auditory systems’ representing them. We hear some of their properties, ditto. Very likely, as part of their total inputs and detailed computational processing, our auditory systems also represent sounds or bits of sound and properties that whole persons do not hear; following usual practice I will continue to call those representations “subpersonal.”

There is an important question about the level of conceptual sophistication at or under which events and properties are perceived. A conservative view would be that we strictly perceive only the environmental properties that our sensory systems are built to detect. For vision, this might be joint colors, shapes, and locations, or a bit less conservatively, also depth, motion, and change as such. As noted previously, it would be slightly more liberal (and already controversial) to contend that we visually represent individuals, or objects such as blocks. Still more liberally, it may be argued, as does Siegel (2010b), that we represent objects under Rosch’s “basic categories” (Rosch et al. 1976) such as “dog,” “shoe,” “chair,” “red,” and “lady,” and even members of specific natural kinds such as pine trees. At a Kuhnian extreme (Hanson 1958), which preceded and in part inspired *The Structure of Scientific Revolutions* (1962),⁷ but especially Churchland (1988), we directly see things described in terms of any theory you like: a woman suffering from Bright’s disease, the advancement of global warming, an economic recession, or a relation in Riemannian space.⁸ For audition, a conservative-to-liberal spectrum of examples might be a sound of such-and-such volume in a certain timbre range; a tone at this or that pitch; a repeated banging of two objects together; splashing liquid; a waterfall; an explosion; a bell ringing; a melody; a melody played on an oboe; a human voice; a voice saying words; a voice beginning to recite “Tintern Abbey”; the sound of a Rolls-Royce’s door closing; the opening instrumental bars of Handel’s “Lift Up Your Heads”; Martin Luther King concluding his “I Have a Dream” speech; the Chicago Symphony Chorus under Margaret Hillis; the bowel sounds of one suffering from Baird’s syndrome; the launch of a Zircon missile; the end of the world.

Elsewhere (Lycan 2014) I have contended that in the case of vision, the liberalism issue is hard if not impossible to settle. It is not an entirely empirical

⁷ I wonder if you have to be as old as I am to remember Hanson’s extraordinary life and work: https://en.wikipedia.org/wiki/Norwood_Russell_Hanson.

⁸ Though Churchland denies that perception is propositional representation, so he is not saddled with the claim that the visual system “has” concepts like “economic recession.”

question, but no ordinary appeal to common sense or to introspection will do, and psychosemantic ideas cannot help. Austen Clark (2000, 2004) puts forward a conservative view, but does not rule out more liberal possibilities. I know of only one actual detailed argument for a robust liberalism: Siegel's aforementioned, based on her "method of phenomenal contrast," but I maintain (Lycan 2014) that it breaks down on each of several points.⁹ Most pertinently for our purposes here, the argument fails to distinguish actual visual representation from merely "seeing as"; does the subject see, by visually representing, a pine tree *per se*, or does she merely see a tree-shaped object and see it *as* a pine? As we shall learn, this distinction is even more vexed in the case of hearing.

3. On questions of consciousness: To be blunt, I think they have nothing to do with the present issue. What sort of "consciousness" might be thought to be required for aspect perception? The idea may be that in a paradigm case, that of deliberate *aspect-flipping* as with the duck-rabbit, one cannot do that unawares, and so it may be held that aspect perception is necessarily conscious in that the perceiver must be aware of the relevant mental aspect over and above what is more strictly seen. But that is not so. To stay with the example, it's widely observed that were you to see what is in fact a duck-rabbit figure amid a whole scene involving ducks, you would automatically and quite unconsciously see it as a duck, but were you to see the very same figure in a rabbit-filled picture, you would just as automatically and unthinkingly see it as a rabbit. Aspect perception is ubiquitous—or so I shall argue, contra Wittgenstein himself—and we are rarely more aware of doing it than we are of perceiving at all. (More generally, we are seldom aware of our own mental state of perceiving, though of course we can be and sometimes are. To be introspectively aware of it, we must at least to some small degree attend to it rather than to the external objects and properties it represents [Lycan 1996, 2008]. But for a mental state to be a "conscious" one in this sense, as opposed to, say, Dretske's [1993] or Block's [1995] quite different respective senses, is a very superficial phenomenon of little ontological significance.)¹⁰

⁹ And see the debate between Siegel and Alex Byrne in Siegel and Byrne 2016.

¹⁰ I here ignore the view that goes under the heading of "phenomenal intentionality," according to which intentionality and on some versions all genuine representation is a phenomenological, hence conscious, phenomenon (Horgan and Tienson 2002); I have rebutted that idea at length in Lycan 2008.

4. I've been assuming that there is a distinction between what is strictly seen, which I identify with being represented by vision alone, and what is represented only by dint of "cognitive penetration" or other influence of background information on the visual system's output that would result in perceptual belief. But the idea of "vision alone" depends on a modular picture of a sense modality (Fodor 1983), or at least a clear distinction between perception and cognition, that may or may not be supported by the most current neuroscience.

I shall continue to assume that wide-open Kuhnian-Churchlandian relativism is false, that perception is not "theory-laden" through and through, and that there is a distinction between what is more or less strictly perceived and what appears a certain way but can do so only because of the subject's knowledge or beliefs.¹¹ The distinction may be weakened in one way or another; Lycan (2014) argued that for vision, it comes in grades or stages corresponding to layers of processing. And recent literature on "cognitive penetrability" shows that even for vision, which neuroscientifically is fairly well understood, the influence of empirical beliefs remains contentious; indeed, there are different kinds of "penetration" that have only just begun to be distinguished.¹² Unfortunately this muddies our question of what the difference is between merely seeing X by seeing X's visible properties and seeing X *as* having this property or that.

Finally, we cannot assume that a single perceptual representation or perceptual experience has just one intentional object or content. Peacocke (1992), Lycan (1996), Noë (2004), and Schellenberg (2008) have each argued that a representation can have more than one content, systematically related to each other by some asymmetrical priority relation. Principally, the perceptual state may represent one object or property *by* representing a more primitive or less ambitious one. For example, Schellenberg argues that we perceive "situation-dependent" properties of external objects, and *thereby* the higher-level properties of the same objects. I shall call views of this kind "layering" views. My own model for layering of this sort is that of deferred linguistic referring: uncontroversially, we refer to a thing by referring to a distinct but in some way more accessible thing that is saliently associated with it.¹³

¹¹ For resistance to the radical view on the basis of vision science, see Gilman 1992, 1994.

¹² See the essays collected in Zeimbekis and Raftopoulos 2015.

¹³ A number by a numeral, a novel by a copy of it or by the author's name, a color by a sample, a military unit by its commander, hospital patient by her complaint ("That pancreas was discharged this morning and went home"), or a restaurant patron by what he ordered ("The ham sandwich wants his check"). *Multiple layering*: In the same gesture, we can refer to a numeral, thereby referring to a number, thereby referring to an office in the building, thereby referring to its occupant.

Layering may be of some aid to liberals such as Siegel, since we might resent something like a pine tree by representing something more primitive. And there are obvious auditory candidates (I say only “candidates” because I do not want to beg any questions about what higher-level properties may be represented in audition): by hearing a sound of such-and-such timbre and duration, you hear a bang, and by hearing the bang you hear a collision; you hear a major triad by hearing three tones sounded simultaneously or in close succession, and perhaps you hear a D₄ minor triad by hearing the triad; by hearing noises as of a human voice, you hear words, and (some contend) by hearing those you hear meanings.

5.2. Aspect Perception

5. It is more than time to set out some facts and standard claims about aspect perception. Naturally, they have always been put in terms of vision.

It is even more unfortunate that for the public and for most professionals, our entrée to aspect perception has been through ambiguous figures, paradigmatically the duck-rabbit (figure 5.1).

We say the figure can be seen as a duck or as a rabbit, but that is inaccurate. The figure is a *picture*, not an animal or other environmental object. Except in unusual circumstances such as a very rare *trompe l'oeil*, we cannot see the picture either as a duck or as a rabbit, but only as picturing a duck or a rabbit. Or to put it better, following Wollheim (1996), we can distinguish between seeing-as and “seeing *in*.” It is less that we see the duck-rabbit figure as




Figure 5.1 Duck/rabbit picture

a picture of a duck (though of course we can do that) than that we see a duck and alternately a rabbit *in* the picture, so to speak not noticing the picture. Since aspect perception is of crucial importance in understanding visual art (Gombrich 1960), seeing-in has been more important to aesthetics and art criticism than has everyday seeing-as.

There are auditory pictures too, in the sense of real-world sounds being deliberately used as representations of events actual or nonactual. A raconteur may interpolate oral sound effects in his funny story. A radio play without words (the opposite of a silent movie) might depict an occurrence in nature, or for that matter a sequence of human interactions. Occasionally this is done in music, typically when there is a battle theme. Isaac wrote “A La Battaglia” in 1487. Apparently Byrd wrote a keyboard suite called “The Battell” (ca. 1588); I don’t know that piece. Janequin produced a madrigal, “La Bataille,” in which choral voices imitate battle sounds (much as he portrayed flocks of birds in the much better known “Chant des Oiseaux”). Famously, the “1812 Overture” is actually scored for, and with sufficient funds is played using, real cannons rather than percussion instruments to represent the cannon fire that had actually occurred sixty-odd years before.¹⁴

In this chapter I should try to focus on perceiving-as, and set aside perceiving-in; so, a bit ironically, the duck-rabbit should not be a paradigm. For the same reason, neither should the Necker cube (figure 5.2).

There is no actual cube, but only the representation of one, and the aspect-flipping it invites is as between the view from the top (with front facing leftward) and the view from the bottom (front rightward), in three dimensions. And likewise for most of the ambiguous figures displayed on pages or seen in paintings. A seeing-as proper in visual art will ordinarily be of a piece of sculpture. For example, Henry Moore’s *Nuclear Energy* can in part be seen as a mushroom cloud, or as a skull—and yet, isn’t that phenomenology still more like seeing-in? That’s why I said “in part.” Of course, Moore could have constructed a purer case, a sculpture that itself really could be aspect-flipped as between a mushroom cloud and a skull. For that matter, someone could make a real three-dimensional but transparent Necker cube and hang it

¹⁴ But cp. Peter Schickel’s hilarious parody, in which the cannon shots are represented by a triangle going “ting.”

Apparently Tchaikovsky had been preceded in this use of actual cannon by Beethoven, in a short 1813 piece called “Wellington’s Victory.”

What should we say of leitmotifs? Do they represent, strictly speaking, or is there merely an intended association? Think of a version of “Peter and the Wolf” that is souped up in this regard.




Figure 5.2 Necker cube

somewhere. For all I know someone has, and ditto the duck-rabbit; in fact, I'm sure that somewhere there is a three-dimensional duck-rabbit figure.

And the preceding point exhibits a third difference between seeing-as and hearing-as: I doubt anyone could ever have introduced the whole topic of aspect perception by playing examples of representational music. For one thing, in drawing and in painting, ordinary three-dimensional objects are represented by (not counting *impasto*) two-dimensional ones, while in representational music events are represented by sounds, which are already real-world events. Also, the parallel question would be, not whether the musical sound itself could be aspect-flipped, but whether the represented sound could be, and I find it hard to think of a good example of doing that.¹⁵ At least historically, the paradigm cases of "seeing as" are really seeings-in; but hearings-in are rare at best, and the paradigm of "hearing as" is hearing-as proper.

6. Truisms coming from Wittgenstein (1953):

- (1) "The expression of a change of aspect is the expression of a *new* perception and at the same time of the perception's being unchanged" (196). Indeed, we are tempted to say that the object seen has "changed," even

¹⁵ What about sounds heard *on recordings*? A musical note or chord heard on a CD can certainly be heard-as and aspect-flipped; and a muffled sound heard on a cassette tape can be heard as a cough or a distant pistol shot. Perhaps these are straightforward cases of hearing-in. Note that there is a close analogue for vision: people and their actions seen on television. But my own view is that those are not cases of pictorial representation. Rather, they are cases of technologically enhanced hearing and seeing *per se*. Thanks to audio and video recordings, you hear the orchestra itself, and you see the president. It lacks only the cachet and the prestige of being there in person. Granted, this is a large subject.

though we "see that it has not changed" (193) and know that it in itself has not.

- (2) We can speak both of the "continuous seeing" of aspects and of the "dawning" of aspects (194). Such a dawning may be very dramatic. Think of those pointillist screens that appear to be just *mélanges* of various colored dots; handed one, we are told to try to see a human figure or rudimentary scene in it. The task is hard, and when or if we succeed, there is a rush of revelation.
- (3) "The flashing of an aspect on us seems half visual experience, half thought" (197). But:
- (4) We "describe the alteration like a perception; quite as if the object had altered before my eyes" (195). Accordingly:
- (5) The perception/cognition distinction does not apply straightforwardly. In particular, it is inaccurate to divide the experience into what we strictly see and how we interpret it in thought. Seeing-as is "like seeing and again not like" (197).
- (6) Seeing-as is part of visual phenomenology (though Wittgenstein would never have put it that way).
- (7) Seeing-as can be relational. E.g., we can "see a likeness" between two faces (193); this seems to amount to more than just seeing the first as like the second and the second as like the first).
- (8) Perceiving-as is profoundly influenced by context. Two perfectly congruent duck-rabbits will be seen as being quite dissimilar if one is surrounded by ducks and the other by rabbits.
- (9) In the striking cases that called our attention to the topic, the alternate aspects compete. We cannot see both the duck and the rabbit simultaneously, but must flip between them.
- (10) Aspects can be consciously and deliberately flipped, at will. (Doing that is fun.)¹⁶
- (11) Perceiving under an aspect is irreducibly conceptual. Even if strictly visual content is nonconceptual, perceiving-as is or at least requires the application of a concept. (That is strongly suggested by the grammar alone: To see X as . . . is to see X as F or as bearing R to Y.)

¹⁶ For most anyone, not just for philosophers and psychologists. Why?

7. At least two more theoretical claims have since been made, and each is plausible for the case of seeing:

EXPECTATION: Perceiving X under a particular aspect is at least in part a matter of expectation, of how one would tentatively expect X to behave or its parts to be perceived should circumstances alter. Gombrich speaks of how we “project” possibilities from what we perceive.

ATTENTION: Aspect perception is an attentional phenomenon. Empirical studies by Ricci and Blundo (1990) and Kleinschmidt et al. (1998) suggest that the neural correlates of ordinary attention are activated during their subjects’ aspect-flipping activities. Earlier, Chastain and Burnham (1975) had shown that the aspect seen in an ambiguous figure depends on which part of the figure the subject focuses on.

ATTENTION in particular is a very important idea. For one thing, it opens the aspect-perception issue to empirical findings from a different and burgeoning area of research. More importantly, it would reduce Wittgenstein’s mystery to a more tractable question about attending: granted that visual attending in particular is a visual phenomenon, in that *at least* it affects visual phenomenology, but what is the sense in which what is seen does not change? Here the latter question has a nonrhetorical answer: By every current psychological theory,¹⁷ attending is a matter of “heightening” or “strengthening” or “enhancing” (or something) an existing representation, as opposed to creating a new representation. What has not changed when an aspect dawns is the content of the existing representation—which is a good thing to mean by “What is strictly seen does not change.”

Orlandi (2011, 2012) uses **ATTENTION** to argue both that (11) is false, i.e., that aspect perception does not require the application of concepts, and that aspect perception does not require interpretation of any sort, if “interpretation” means an inference-like subpersonal process.¹⁸

¹⁷ See the references in Mole 2010.

¹⁸ Gauker (2017) defends an even stronger anti-concept view of aspect perception. He holds that all perceptual reference is merely “marking” in similarity spaces, and argues that aspect perception is just a special case of that.

Even if we are unconvinced by Orlandi and Gauker and insist on the truth of (11), there is a serious question of *which* concepts are mobilized in aspect-perceiving. DeBellis (1995) points out that although music listening and appreciation are shot through with aspect perception and the details are given extensive analysis by music theorists both traditional and psychology-informed, even




Figure 5.3 Dot array

8. If we do restrict consideration to aspect perception in propria persona as opposed to perceiving-in, it still comes in several subcases. First, before we get to ambiguous pictures where it seems obvious that everyday concepts are involved, there are simpler organizational examples such as dot arrays (figure 5.3), nonpictorial figures in the sense that there is no issue of a three-dimensional interpretation.

The array in figure 5.3 can be seen either as four horizontal rows of three dots each or as three columns of four each.

Figure 5.4 can be seen as a diamond shape standing on one of its points or as a square tipped up onto one of its corners. (Notice that this case doubles as an illusion, for seen as a diamond it looks bigger.)¹⁹

And Nickel (2007) offers a richer example (figure 5.5).

Nickel says we can see an arbitrarily chosen set of constituent squares “as prominent.” For example, in figure 5.5 we can see the top row as prominent,

fairly sophisticated listeners do not often have the precise technical concepts applied by the music theorists (“second inversion,” “hemiola,” “Neapolitan sixth”). DeBellis tends to identify the conceptual with the doxastic, so is disinclined to grant even that the listeners hear things under concepts corresponding crudely to those technical terms, but we could insist on the latter without attributing particular beliefs to the listeners.

Notice that a similar issue infects even the duck-rabbit paradigm. Even sophisticated philosophers and psychologists do not aspect-see ducks or rabbits in the strict zoological sense of those terms. I suppose there are ordinary concepts corresponding to Putnamian stereotypes as opposed to Putnamian natural kinds. But a person need not have heard of ducks or rabbits even in those ordinary senses in order to aspect-flip on the duck-rabbit figure. It may be that the person must have some rudimentary idea of “ear” or “beak.”

¹⁹ This is for the obvious reason that in diamond mode, both its top-to-bottom measurement and its side-to-side measurement are longer by $\sqrt{2}$ times than the square’s sides.




Figure 5.4 Square / regular diamond




Figure 5.5 Nickel's array

alternately the left column, or alternating squares starting with the top left, the complement of that—or most any smaller or larger combinations of squares—while representing just the same figure and its elements.²⁰

There are auditory analogues. Think of a regular succession of beeps at half-second intervals. You can hear them as grouped in twos, or in threes, or in fours. You can hear them in different meters—3/4 versus 4/4, for example. If you're a musician you may be able to hear the difference between 3/4 and 6/8. Or imagine that the "piece" is notated in 3/4, but it is coming to a big cadence and the penultimate two measures are a hemiola. A classic hemiola is precisely a dramatic aspect-switch to duple meter within triple.²¹

Moreover, there are Nickelian differences of prominence that can be flipped at will. Imagine, or actually listen to, a major triad that is being sustained indefinitely. Hear the tonic as being merely harmonized by the third and fifth. Now switch to hearing the third as the melody note, with the tonic as the bass and the fifth as harmony. Now hear the fifth as the melody and imagine that the third and tonic are about to drop a tone each to make V in first inversion; better, imagine that the tonic's drop to the leading tone will be slightly delayed, making a suspension; that will reinforce hearing the fifth as melody.

For these kinds of cases, "organizational" as we may call them, ATTENTION looks pretty good. In the visual ones it does seem that we are doing small shifts of attention, even though Nickel says that to get his effect we need not change the focus of our vision. ATTENTION is less obvious in the musical cases, but still perhaps tenable.

Second subtype: The standard cases of seeing ordinary objects under aspects. My favorite example (please pardon my using it yet again): in chapter 17 of Booth Tarkington's *Penrod and Sam*, Penrod is sitting reluctantly in church and comes out of a doze to find himself looking at a coconut, sitting on an inverted dinner plate "miraculously balanced" on the back of a pew a few rows in front of him. He pulls himself together and realizes that the "coconut" is Georgie Bassett's head. Had he wanted to, he could then have relieved the continuing tedium by aspect-flipping between Georgie, coconut, and perhaps other kinds of object. Less amusingly, we can see a dog as a wolf or vice versa, a spot on the wall as a bug, a rock formation as a human face, a convex surface as concave.

(Wittgenstein himself held, on the basis of ordinary-language considerations since discredited by Gricean pragmatics, that except in abnormal circumstances, an ordinary object is not perceived as the kind of object it is, or rather that to say it is commits a linguistic error. That's just wrong; concave

Peacocke (1983), Macpherson (2006), and Nickel (2007) mobilize these and other examples in acting to the representation theory of sensory qualities; Block (2010) uses Gabor patches similarly. But that issue is not our concern here.

²¹ And how wonderful that is. Your whole body will move to emphasize the effect, if you don't stop yourself. My choral conductor wife reminds me that Handel's "And the Glory of the Lord" is a joyous riot of hemiolas.

surfaces can be and normally are seen as concave, a spot on the wall is seen as a spot on the wall, and in the novel, Georgië's head was nearly always seen as a head if not as Georgië's.)

There are auditory analogues for these ordinary cases too. A gunshot can be heard as a popping like that of popcorn, or as a gunshot. Wind whistling through an aperture can be heard as a human voice, as can a human voice. A voice can be heard as speaking English, or just as making noises. A tune can be heard as being played on a wind instrument, or as on a more specific type of wind instrument.

Third subcase: Perceiving ordinary objects under *very high-level* aspects. A cloud can be seen as a human head or as a particular type of animal. A display of colors and dots on a map can be seen as a Republican triumph. Naked peaks in New Zealand can be seen as the advancement of global warming.²²

Here too there are auditory cases; merely consider the more high-level possibilities considered previously as candidates for representation on a very liberal view: a distant rustling as Margaret Hillis conducting, or a distant rumble as the launch of a Zircon missile.

Is aspect perception itself a kind of representation? Some authors assume that it is,²³ but, supposing that vision "itself"/"alone" does not represent the cloud we're looking at as a Tibetan antelope having stomach trouble, is there really a visual state of mine that does represent a Tibetan antelope having stomach trouble? (Obviously some mental state does, viz., the thought that occurs to the viewer, but that's not the present question.) I suppose this just reproduces Wittgenstein's original mystery.

5.3. Music

9. In beginning this chapter I said that hearing-as plays at least two central roles in human life. I have in mind its function in listening to music and, even more central, its function in understanding speech.

²² I here ignore a fourth and still more diffuse sort of case mentioned in Lycan 2014: The nonce use of everyday objects to represent the movements of ships, or of military units in a battle. The phenomenon is visual, because through the isomorphism intended by the players, a viewer can learn how the real units did move in relation to each other, and possibly "see" when one destroyed another. But it seems a stretch to say that a particular spoon or coffee cup can be seen-as a brigade of infantry.

²³ And some just assume that it isn't. Hence the use of seeing-as examples as objections to the representation theory of sensory qualities, noted in note 20. (But if seeing-as is a specifically visual and not merely cognitive phenomenon, then how is seeing a human head as a coconut not, at least in part, visually representing the head as a coconut?)

But care is needed here, because if hearing-as per se is to have a function, or even play a distinctive role, in those activities, it must somehow contrasted with representation, or (not to beg the question raised a moment ago) at least with some more basic kind of representation. To start, we must insist that aspect perception is a whole-person-level achievement and contrast it as before with the subpersonal representations that figure in the mechanics of visual, musical, or language processing. Interestingly, there are a number of person-level musical illusions.²⁴


But if the content of a perceptual state is multiple and layered, as I believe, matters are complicated further. Layered perceptual representation is not per se aspect perception, but it is whole-person rather than merely subpersonal, and its more distal intentional contents may coincide with aspects. Suppose Siegel is right and we can visually represent pine trees as such. I would insist that we do that *by* representing objects having shapes. And suppose I do see a pine tree by visually representing a pine tree. I may be seeing the tree *as* a pine. But I might not be; I might even know it is a pine but for some reason be seeing it as some other variety or as something else entirely. So, under our suppositions, there are those two incompatible aspect-perception possibilities. (And if Siegel is wrong, the third possibility is that I cannot visually represent a pine tree as such, but can see what I do represent as a pine tree.)

Is there an auditory analogue? Suppose I can represent a high-level musical entity, such as a closing cadence in E^b major. Normally I would hear it as such, but might I know it to be one and still hear it as something else? That is hard to imagine—what could a cadential formula be but a cadence?—though we cannot rule out the possibility a priori. But, obviously, if hearing itself cannot represent anything so abstract, I can still hear the sound as a closing cadence in E^b major.

Music commentators and theorists use the "as" locution all the time. A selection culled from Mark DeBellis's (1995) book on conceptualization of music:

Harmonically, one can hear m. 3 as revolving around F, and m. 4 around G. In that case, the entire progression of mm. 1–5 can be heard as an expansion of the I–IV–V–I of the opening measure.

Example 5.1 The three diminished 7th chords



It is . . . possible to hear the establishment of the D-flat triad in m. 49 as the arrival of tonic harmony.

This rhythmic arrival is heard also as the beginning of a new hypermeasure.²⁵

But such usage leaves it open whether a trained listener's auditory system actually represents things like harmonic expansions and hypermeasures. In their monumental mingling of traditional music theory with cognitive and perceptual psychology, Lerdahl and Jackendoff (1983) are well on the optimistic side of that question, thereby diminishing our scope for musical as-auditory perception that outruns representation.

Perhaps we should approach the matter by considering ambiguity (unambiguous ambiguity, so to speak). Tonality in music results from unequal division of the octave, in Western music into half steps versus whole steps; *scale*, regardless of key or absolute pitch as notated by letter name, is precisely a division of the octave into half steps and whole steps, and different scales will sound quite different in character even to one who knows nothing about music. For ambiguity, then, we look to equal division of the octave. We can divide it into tritones (on a keyboard, diminished fifths or augmented fourths),²⁶ into minor thirds, into whole tones, or into half tones. The classic case of tonal ambiguity is the ("full") diminished 7th chord, of which there will be three in any key, along with the diminished triads that are their proper parts. Example 5.1 shows them in C.

(You can get a remarkable effect by rippling up and down these chords for a while without resolving any of them.)²⁷ In tonal music a full diminished 7th can be resolved directly to a 6th chord by lowering two of its notes

²⁵ See DeBellis's references (1995, 10). The writers are music theorists, not philosophers or psychologists.

²⁶ On a keyboard instrument, different tritones will not be truly the same interval, because of equal temperament. And a true tritone can be heard *either* as an augmented fourth or as a diminished fifth. Mediaeval music theorists, who had a Platonic/Christian way of parsing the cosmos in musical terms, did not like the tritone, and reportedly called it "diabolus in musica." I'm pretty sure some of them would have meant that literally.

²⁷ My composer friend Arthur Weink once wrote a lied that had such rippling as its piano accompaniment.

Example 5.2 From Mozart's Sonata No. 6 in D Major, K. 284

a semitone each (if a 6th is considered sufficiently restful in the context), but more commonly it will resolve to a straight 7th chord, which can serve as V⁷ of a new key. The breadth of the ambiguity is appreciated when we see that lowering any one of its four tones will produce a different 7th. For example, in the first of the three chords in example 5.1, we can get either B⁷, D⁷, F⁷, or A^{b7}.

Just listening to a diminished 7th being sounded out of context, it takes real sophistication to hear that four-way ambiguity, but it is there to hear.²⁸ And that experience should count as hearing-as, for what is represented does not determine any resolution at all.

Contrast example 5.2: a whopping cadential formula, especially when it's tricked out with a conventional device such as the trill on the supertonic.

Not much ambiguity there! Now consider it minus its final I: If you were to play all but the last chord and stop abruptly, even hearers who know nothing about music would yelp with frustration, or at least feel frustrated. (But there is still a potential ambiguity: that cadence, however final-sounding, *could* be the closing cadence of a movement within a longer work, where the movement itself is in V relative to the key of the whole. So a true disciple of Schenker adept at "structural hearing" would try to hear the cadence as, rather, a half-cadence, and might succeed.)

²⁸ Perhaps this is a musical counterpart of an ambiguous figure. Casey O'Callaghan has suggested to me that Deutsch's surprising "tritone paradox" is one also: <http://deutsch.ucsd.edu/psychology/pages.php?i=206>.

Example 5.3 From Mozart's Sonata No. 6 in D Major, K. 284

Example 5.4 From Schubert's "Der Wegweiser," in "Winterreise"

But the really ubiquitous case of deliberate and indispensable ambiguity is the pivot chord in a modulation,²⁹ ubiquitous because modulation (temporary change of key) is an important feature of any very interesting piece. Another Mozart example, very mild, is example 5.3, where the pivot chord is the v^6 in measure 5 that doubles as ii^6 in the new key (A), setting up the new dominant (E^7) in measure 6. In analysis, pivot chords are sometimes notated with a hyphen: v^6-ii^6 . (I called this example "mild" because the new key is closely related to the old, and the pivot chord is not highly ambiguous. The use of a diminished chord as pivot is of course less conservative because of its degrees of freedom aforementioned, and may introduce a more striking modulation.)

A more daring one, from a minor key to the minor one whole step higher, is shown in example 5.4.

²⁹ I'm speaking of a classic modulation that does work by pivot chord or "common chord"; not all do, and some modulations, for one reason or another, are entirely unprepared. (Some are semi-conventional, as when in pop music the last verse of a song is moved up a semitone.)

Here the pivot occurs by an intervening step of chromatic alteration, melodically lowering the third of old (Fm) V to make new (Gm) iv .³⁰

I've suggested that hearing-as is crucial because modulation of this type requires tonal ambiguity in the pivot. But it is tricky to say *when* the pivot chord aspect-dawns. It does not dawn the instant the pivot is sounded, because at that point the progression is still within the old key and could simply continue in that key. It does not seem to dawn at any time until the next, alien chord sounds; so far as the pivot chord was being aspect-heard in its split second of existence, it was being heard as in the old key. But by the time that next chord sounds, the pivot chord is past and no longer being heard at all. The best we can say is that the alien chord makes us hear the pivot ambiguously in retrospect; it is partly a memory phenomenon. (What if the pivot had been the opening though pickup chord of a piece, and the next chord established the piece's key? But that merely reveals the fact that the opening chord of any musical work is ambiguous and cannot by itself establish any key at all; indeed, even the first and second chords together do not suffice.) We might think this is a *reductio* of the idea that aspect switch is doing the work in modulation. But (a) that the auditory *ambiguity* is doing the work is undeniable, (b) "ambiguity" in this sense is defined in terms of hearing-as, and (c) the advent of the new key certainly presents as the dawning of an aspect. Another puzzle. Perhaps I am looking for too sudden a dawn; what's certain is only that by the time a couple of measures have gone by, I am now hearing certain chords as I that I had been hearing as V, and the like.

³⁰ I will not resist footnoting a favorite of mine from Lennon (example 5.5).

This qualifies as a modulation because of the suddenness and prominence of the F ; major chord in what was C. Actually E is just V of the relative minor (Am) and was set up by Am as pivot, but the piece could have continued in Am or for that matter in A (major). As is, re-pivoted by F ; major that is VI of Am, it modulates beautifully and restfully back to C.

Example 5.5 From the Beatles' "I Want to Hold Your Hand"

What of the two theoretical claims formulated in section 7? EXPECTATION is nearly a byword in music theory. Indeed, much more so than Gombrich would have imagined in thinking about visual art, because music is a temporal and a performance art, and we are literally expecting events that have not yet occurred. The great Leonard Meyer (1956) went so far as to explicate the value and interest of a melody or of a harmonic progression entirely in terms of expectation, formalized in turn in terms of Shannon-Weaver information theory. More to the point, the aspect under which we hear an ambiguous tone or chord is indicated by what we expect to follow it—in the case of a dissonance, how we expect it to resolve.

But now for an unexpected moral from our attention to aspect perception in music: hearing-as does not seem to have much to do with attention; and so it casts doubt on ATTENTION. The psychologists cited previously were thinking of our ways of attending to pictures: focusing on one or more parts of a picture, or scanning the picture in this order or that. But we do not do that with chords, not even when a chord is sustained for an appreciable period of time (as in the ordinary case of modulation, it is not). This is sad news, because ATTENTION offered real hope for the understanding of aspect perception generally.

10. Music also affords an entirely different category of hearing-as: the emotional. We hear a passage, or a melody, or a scale or (some say) even a key as having affective qualities, the paradigm being happiness or sadness and specifications of those. I personally think the understanding of music in terms of emotion is absurdly overemphasized, but there's no denying that many go in for it, or that it has a solid psychological basis. If there were no natural association between musical properties and affect, there would be no movie soundtracks, to say nothing of opera.

The literature on the expressing (in some sense) of affect by music is huge, and I cannot enter into any of those debates; I shall only comment on emotional properties as they may be aspect-heard. It is sometimes said that a piece or passage "is sad," "expresses sadness," or "is heard as sad." Given the all too obvious fact that a musical sound-event is not itself a sentient creature, what could such statements mean?—but now merely as concerning aspect perception.

I shall just follow Peter Kivy (2002) in supposing that when we say such things and we are talking about the music itself, they do not mean things about real emotional effects the music will have, or would have, or tends to have, and so on, upon listeners. Listening to music obviously can affect

listeners emotionally, for any of a variety of reasons, but there is no evidence for the claim made by "arousal" theorists that merely listening to well-performed "sad" music has per se any tendency to make a concertgoer actually feel sad; indeed the better the performance, the more exhilarated the audience members are likely to be.³¹ The "sadness" is a property of the music, even though a sound-event cannot be sad in the same sense a person can be. I sympathize with the "contour" theory defended by Kivy (1980) (though I acknowledge the objections he himself went on to make in Kivy 2002, I believe they can be overcome): a piece of music or passage "is sad" in that its literal or near-literal properties are analogous to salient properties of a sad person, much as (Kivy's central example) a St. Bernard's face looks like a caricature of a sad human being.

Sad music being in slow, halting tempo, subdued in dynamics, with drooping, faltering melodies, and sad people walking in slow, halting gait, with drooping bodies, speaking in subdued, halting voice, cannot be altogether coincidental....

Melancholy people's voices tend to "sink," and tend to remain in the low vocal register....

[T]he "contour" of music, its sonic "shape," bears a structural analogy to the heard and seen manifestations of human emotive expression.

(2002, 37–40)

How, then, can such music be heard as expressing sadness? Simple anthropomorphism. For obvious evolutionary reasons, we have at least a slight propensity for seeing things as animate, and (again following Kivy) I suggest this carries over to hearing as well. At least, we can hear a sound-event *as* expressing sadness (even when we know it is actually doing no such thing) in much the way we see the St. Bernard's face, or a natural configuration, *as* doing so.

And this is a clear case of hearing-as that is not musical representation. So far as it is representation at all, it defies all existing psychosemantics. And if it qualifies as cognitive penetration, it seems to be of a special kind.

³¹ A personal example: I once attended a performance of Lyell Cresswell's orchestral piece "Vaur," whose title is Gaelic for "sorrow" or "deep grief" and which according to the New Zealand website SOUNZ (<http://soundz.org.nz/>) "is concerned with different aspects of sorrow." I loved the piece, especially for its novel orchestration effects. At a party following the concert I exclaimed to Cresswell himself that I got nothing from it about sorrow or grief but had expressed my listening pleasure in excited mental exclamations of "that is so cool!" I think he was disappointed; perhaps he holds an unusual view.

5.4. Speech

11. Moore famously noted the huge phenomenological difference between hearing someone utter a sentence in a language we do not understand and hearing someone utter a sentence in our native language. O'Callaghan (2011) points out that this difference can be made fairly precise and controlled for, in experiments making use of "sinewave speech."³² The difference offers a prime *candidate* for hearing-as, because the sense in which what is "strictly" heard is obvious to common sense.

But open to correction by science. Psychologists speak of the "language module,"³³ and it is held by some that we do not first hear sounds that are speech sounds and then use background information to infer or construct a linguistic entity; rather, we are built to just hear meaningful utterances as such. According to a strong version of this view, we simply hear *meanings*. According to an even stronger version, we hear strong implicatures, and in such cases do not even hear literal meanings.³⁴ If some such modularist views are correct, then we actually represent semantic sound events, and are not merely hearing-as. But O'Callaghan (2011, 2015) argues that the empirical evidence for the modularist claims is scanty at best. So, as we did in the case of music perception, let's turn to known ambiguities.

Of course there are oronyms and mondegreens, phonological ambiguities that depend in part on dialect.³⁵ My old colleague Stan Munsat tells a story in which "I buried it" is misheard as "A bear ate it" or vice versa. My personal favorite is the possibility of hearing the Beatles' line "The girl with kaleidoscope eyes" (from "Lucy in the Sky with Diamonds") as "The girl with colitis goes by." And my wife remembers from 60 years ago an entire genre of pun jokes that started with "She was only . . ." an instance of which is "She was only a stablehand's daughter, but all the horsemen knew her" ("horse manure"). But I am more interested in ambiguities of meaning.

³² Remez et al. 1981. To hear examples: http://www.lifesci.sussex.ac.uk/home/Chris_Darwin/SWS/. (Thanks to Casey O'Callaghan for the reference.)

³³ See, e.g., Pinker 1994.

³⁴ Actually most of those who claim we hear meanings are philosophers, e.g., Strawson (1994). I know of no one who has defended the implicature claim in print, but I have heard more than one linguist say it. It is supported both by phenomenology and by appeal to memory, but those considerations are not strong.

³⁵ The term "mondegreen" is attributed to the columnist Sylvia Wright (<http://wordinfo.info/unit/3347>), who reportedly misheard the Scottish folk song line "They hae slain the Earl of Murray, and laid him on the Green" as "They hae slain the Earl of Murray and Lady Mondegreen."

As Paul Ziff used to emphasize, nearly every English sentence has permissible readings other than the obvious one(s), readings that would never have occurred to anyone. How these are ruled out in normal disambiguation is entirely unknown; even to speak of "disambiguation" is to suggest a process of some sort, but Gricean reasoning or even relevance-theoretic adjustment would be far too cumbersome. Still, we may suppose for the sake of argument that the intended reading is linguistically represented by the hearer. But sometimes the sentence uttered is more conventionally and obviously ambiguous, and the hearer genuinely can't tell from intonation which reading is intended. Moreover, suppose the utterance is recorded, and we can play it over and over. Then we could deliberately aspect-flip it.³⁶

For syntactic ambiguities perhaps ATTENTION has some purchase; when we aspect-flip Groucho Marx's "This morning I shot an elephant in my pajamas," we may be relying on tacit emphasis and focusing via slightly different-sized gaps between words. But semantic ambiguity is less amenable; I don't see how attentional differences would apply to uniformly pronounced lexemes.

12. There is a more interesting and I think indisputable sort of aspect perception in listening to speech: interpreting illocutionary force. As hearer of a particular utterance I may be genuinely undecided as between a mere statement of fact, a warning, an actual threat, or possibly even a reassurance. As children say, "Is that a threat or a promise?" Assuming there is no ambiguity in the sentence uttered, syntactic and semantic representations will be fixed, and my uncertainty will concern only the speaker's intentions. Or, alternately, the context may leave no uncertainty even when the speaker's tone is flat and does not contribute. I hear the threat very definitely as a threat.

An even more striking example is the hearing-as we do when attending a play. The illocutionary forces conveyed by the actors are only simulated; the human beings on stage are not really making statements, asking questions, or issuing requests or commands, but only simulating those actions. But in anything like the normal case of watching a play, we vividly hear their utterances

³⁶ It is now forgotten, and certainly was by me until I happened to see Davies 2011, that Wittgenstein explicitly applied his discussion of aspect perception to linguistic understanding. He held that sometimes we "experience" the meanings of words (210). He even considers a hypothetical condition of "aspect-blindness" applying to pictures (213–214) in which the sufferer could see only one aspect and not the other(s). That would translate, he goes on to say, into inability to hear alternate meanings of ambiguous words.

And Wittgenstein actually mentions both the appreciation of puns and the phenomenon of "secondary" lexical meanings that depend asymmetrically on normal ones.

as having those forces, and react emotionally, even when we know that the utterances do not actually have those forces.

Even in theater there can be uncertainty and genuine ambiguity on the latter point. Some avant-garde drama is interactive, well past “A-effect”; audience members may contribute genuine speech acts, to which an actor may respond but while still in character. Or it may be deliberately left open which part of the venue is the stage and which human beings are actors, stagehands, or audience members.³⁷ Or—never mind the avant-garde—the real world might intrude: The stage manager might step out, distraught, to announce that (backstage) one of the actors has been shot.³⁸

Thus, illocutionary aspect ambiguity can run riot. Here again it is not explained by linguistic representation. And here again, although differences in attending may play a role, ATTENTION cannot be the whole story or even much of it.

5.5. Epilogue

13. Aspect perception is not literally paradoxical, but it remains mysterious. My tentative diagnosis: To sort out exactly what aspect perception is and how it works, we will have to gain a much fuller empirical story than we have now, not just about vision or just about vision and audition but about all the exteroceptive sense modalities. But even complete empirical information will not settle the matter, because the mystery involves philosophically contentious questions such as the nature of representation and how to understand matters of cognitive penetration generally. Moreover, it will have to be decided how to lay folk categories and concepts over the neuropsychological facts (and for that matter how much of the relevant folk psychology must be respected and preserved). Some of that decision may turn out to be purely verbal, but some will not.

I’m quite sure that none of that will happen in my lifetime, and probably not in yours. It may happen before the heat death of the universe.

Acknowledgments

I am grateful to Casey O’Callaghan, Nico Orlandi, Jenny Judge, and Elizabeth Skoe for bibliographic help, and to John Brackett for technical discussion and diminished chords and others.

References

- Block, N. J. 1995. On a confusion about a function of consciousness. *Behavioral and Brain Sciences* 18: 227–287.
- Block, N. J. 2010. Attention and mental paint. In E. Sosa and E. Villanueva, eds. *Philosophical Issues, 20: Philosophy of Mind*. Oxford: Wiley-Blackwell. 23–61.
- Campbell, J. 2002. *Reference and Consciousness*. Oxford: Oxford University Press.
- Casati, R., and J. Dokic. 1994. *La philosophie du son*. Nîmes: Chambon.
- Chastain, G., and C. A. Burnham. 1975. The first glimpse determines the perception of ambiguous figure. *Perception and Psychophysics* 17: 221–224.
- Churchland, P. 1988. Perceptual plasticity and theoretical neutrality: A reply to John Fodor. *Philosophy of Science* 55: 167–187.
- Clark, A. 2000. *A Theory of Sentience*. Oxford: Oxford University Press.
- Clark, A. 2004. Feature-placing and proto-objects. *Philosophical Psychology* 17: 471–484.
- Davies, L. 2011. Aspect perception and understanding the meanings of words: Wittgenstein’s *Philosophical Investigations*. *Sophia* 12.
- DeBellis, M. 1995. *Music and Conceptualization*. Cambridge: Cambridge University Press.
- Dretske, F. 1993. Conscious experience. *Mind* 102: 263–283.
- Farennikova, A. 2013. Seeing absence. *Philosophical Studies* 166: 429–454.
- Fish, W. 2009. *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.
- Fodor, J. A. 1983. *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Gauker, C. 2017. Three kinds of nonconceptual “seeing-as.” *Review of Philosophy and Psychology* 1–17. DOI: 10.1007/s13164-017-0339-2.
- Gilman, D. 1992. A new perspective on pictorial representation. *Australasian Journal of Philosophy* 70: 174–186.
- Gilman, D. 1994. Pictures in cognition. *Erkenntnis* 41: 87–102.
- Gombrich, E. H. 1960. *Art and Illusion*. London: Phaidon.
- Hanson, N. R. 1958. *Patterns of Discovery*. Cambridge: Cambridge University Press.
- Horgan, T., and J. Tienson. 2002. The intentionality of phenomenology and the nomenclology of intentionality. In D. Chalmers, ed., *Philosophy of Mind: Central Contemporary Readings*. Oxford: Oxford University Press. 520–533.
- Kania, A. 2010. Review of Nudds and O’Callaghan 2009. *Natur Dialekt Philosophical Reviews*, August 17.
- Kivy, P. 1980. *The Corded Shell: Reflections on Musical Expression*. Princeton, NJ: Princeton University Press.
- Kivy, P. 2002. *Introduction to a Philosophy of Music*. Oxford: Oxford University Press.
- Kleinschmidt, A., C. Buchel, S. Zeki, and R. S. Frackowiak. 1998. Human brain activity during spontaneously reversing perception of ambiguous figures. *Proceedings of Royal Society of London, Series B: Biological Sciences* 265: 2427–2433.

³⁷ Bruce Wilshire (1982, ix–xiv) vividly and appreciatively describes attending a performance of Robert Whitman’s *Light Touch*, in which this conceit is employed.

³⁸ I owe this sort of example to Mayank Bora. An actor friend of mine once appeared in a performance during which one of the other actors died, suddenly, actually, on the stage.

- Kuhn, T. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Lerdahl, F., and R. Jackendoff. 1983. *A Generative Theory of Tonal Music*. Cambridge, MA: MIT Press.
- Lycan, W. G. 1971/1985. Gombrich, Wittgenstein and the duck-rabbit. *Journal of Aesthetics and Art Criticism* 30: 229–237. Reprinted in J. V. Canfield, ed., *The Philosophy of Wittgenstein: Aesthetics, Ethics and Religion* (New York: Garland Publishing).
- Lycan, W. G. 2008. Phenomenal intentionalities. *American Philosophical Quarterly* 45: 233–252.
- Lycan, W. G. 2014. What does vision represent? In B. Brogaard, ed., *Does Perception Have Content?* Oxford: Oxford University Press. 311–328.
- Lycan, W. G. 2018. What does taste represent? *Australasian Journal of Philosophy* 96(1): 28–37. <http://www.tandfonline.com/epint/P7RszBA2tMZfhF3WqtKA/full>.
- Macpherson, F. 2006. Ambiguous figures and the content of experience. *Notis* 40: 82–117.
- Meyer, L. 1956. *Emotion and Meaning in Music*. Chicago: University of Chicago Press.
- Moore, C. 2010. *Attention Is Cognitive Unison: An Essay in Philosophical Psychology*. New York: Oxford University Press.
- Nickel, B. 2007. Against intentionalism. *Philosophical Studies* 136: 279–304.
- Loë, A. 2004. *Action in Perception*. Cambridge, MA: MIT Press.
- Nudds, M. 2009. Sounds and space. In M. Nudds and C. O'Callaghan, eds., *Sounds and Perception: New Philosophical Essays*. Oxford: Oxford University Press. 69–96.
- Nudds, M., and C. O'Callaghan, eds. 2009. *Sounds and Perception: New Philosophical Essays*. Oxford: Oxford University Press.
- O'Callaghan, C. 2007. *Sounds: A Philosophical Theory*. Oxford: Oxford University Press.
- O'Callaghan, C. 2011. Against hearing meanings. *Philosophical Quarterly* 61: 783–807.
- O'Callaghan, C. 2015. Speech perception. In M. Matthen, ed., *The Oxford Handbook of the Philosophy of Perception*. Oxford: Oxford University Press. 475–494.
- O'Shaughnessy, B. 2009. The location of a perceived sound. In M. Nudds and C. O'Callaghan, eds., *Sounds and Perception: New Philosophical Essays*. Oxford: Oxford University Press. 111–125.
- Orlandi, N. 2011. The innocent eye: Seeing-as without concepts. *American Philosophical Quarterly* 48: 17–32.
- Orlandi, N. 2012. Embedded seeing-as: Multi-stable visual perception without interpretation. *Philosophical Psychology* 25: 1–19.
- Ornstein, R. 1999. What is sound? *Philosophical Quarterly* 49: 309–324.
- Palazzo, C. 1983. *Sense and Content*. Oxford: Oxford University Press.
- Palazzo, C. 1992. *A Study of Concepts*. Cambridge, MA: MIT Press.
- Perkins, M. 1983. *Sensing the World*. Indianapolis: Hackett.
- Putnam, S. 1994. *The Language Instinct: How the Mind Creates Language*. New York: W. Morrow.
- Reid, W. 2007. *Representation Reconsidered*. Cambridge: Cambridge University Press.
- Remez, R. E., P. E. Ruben, D. B. Pisoni, and T. D. Carell. 1981. Speech perception without traditional speech cues. *Science* 212: 947–950.
- Rodriguez, C., and C. Blundo. 1990. Perception of ambiguous figures after focal brain lesions. *Neuropsychologia* 28: 1163–1173.
- Rosen, E., C. Mervis, W. Gray, D. Johnson, and P. Boyes-Braem. 1976. Basic objects in natural categories. *Cognitive Psychology* 8: 382–439.
- Schellenberg, S. 2008. The situation-dependency of perception. *Journal of Philosophy* 105: 55–85.
- Schellenberg, S. 2011. Perceptual content defended. *Notis* 45: 1–37.
- Scruton, R. 1997. *The Aesthetics of Music*. Oxford: Clarendon Press.
- Siegel, S. 2010a. Do visual experiences have contents? In B. Nanay, ed., *Perceiving the World*. Oxford: Oxford University Press. 333–368.
- Siegel, S. 2010b. *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Siegel, S., and A. Byrne. 2016. Rich or thin? Debate in B. Nanay, ed., *Current Controversies in Philosophy of Perception*. London: Routledge. 59–80.
- Sorensen, R. 2008. *Seeing dark things: The philosophy of shadows*. Oxford: Oxford University Press.
- Sorensen, R. 2009. *Hearing silence: The perception and introspection of absences*. In M. Nudds and C. O'Callaghan, eds., *Sounds and Perception: New Philosophical Essays*. Oxford: Oxford University Press. 126–145.
- Strawson, G. 1994. *Mental Reality*. Cambridge, MA: MIT Press.
- Travis, C. 2004. The silence of the senses. *Mind* 113: 57–94.
- Wilshire, B. 1982. *Role Playing and Identity: The Limits of Theatre as Metaphor*. Bloomington: Indiana University Press.
- Wittgenstein, L. 1953. *Philosophical Investigations*. Trans. G. E. M. Anscombe. Oxford: Basil Blackwell.
- Wollheim, R. 1996. Seeing-as, seeing-in, and pictorial representation. In *Art and Its Objects*. Cambridge: Cambridge University Press.
- Zeimbekis, J., and A. Raftopoulos, eds. 2015. *The Cognitive Penetrability of Perception: New Philosophical Perspectives*. Oxford: Oxford University Press.