HOW BIG DO THINGS LOOK?

Ron McCLAMROCK

ABSTRACT: The idea that we have direct and infallible knowledge of appearances is still deeply entrenched; and even scholars who reject this idea often still presume that our normal awareness of the shape and size of objects includes awareness of something like the shape and size of the image it projects onto the retina. I show here how these ideas are undermined by some new empirical evidence regarding these features, as well as by some observations concerning the phenomenology of size, the familiar moon illusion, and the persistence of illusions more generally. These considerations further suggest a path for dealing with the phenomenology of appearance more broadly.

KEYWORDS: perception, consciousness, phenomenology, vision, mind

1. Considering Apparent Size

There are few philosophical dichotomies as deeply entrenched as that between how things *appear* and how they *are*. But over 100 years into both modern psychology and phenomenology, talk about how things *appear* to us or *look* to us in the case of vision is still a bit of a mess. I'm going to try to sort out a tiny bit of that mess here.

But before I do that, I'm going to ask you to start with a little personal phenomenology of perception. Now, we all know that the Moon is huge, and really far away. But how big does it *look*? One obvious way to try to make that question more concrete would be to transform it into a question about how big a disk you'd need hold up at arm's length to subtend the same angle or project the same "size in the image" as the disk of the Moon. Or, what you might take to be the same question, *how big a disk would you need hold up at arm's length for it to* look *the same size as the disk of the Moon*? Go ahead, consider it; perhaps take note of your answer before reading on.

So, here's the point about how things look—that is, about our visual experience of them, or if you like, the phenomenology of vision—that I'd like to make: The visual angle subtended in our view of an object (or what we might call the *size in the image*) is often unavailable to us, in the sense that we don't know even what should be basic and obvious facts about relative *size in the image*. That makes such size in the image a terrible candidate for something that's part of the appearance or visual experience of the object. It's not just that there's often more to the appearance than merely size in the image (although that's surely true as well);

It's that size in the image is – under normal circumstances – just not part of the character of visual experience at all.

As noted, I'm interested in how things *look*; how they *appear*; or what the *character of the visual experience of them* is. The substantial body of literature in psychology on the perception of size informs but does not resolve questions about *apparent* size. Even if we had a well-worked-out story of our eventual perceptual judgments about the size of objects, that wouldn't tell us how big or small the objects *appear*, unless of course we are to assume that how big something *appears* to us and how large we *take it to be* on the basis of our perception must be the same thing. But almost no one wants to do this. Most obviously, this conflicts with the phenomenon of the persistence of illusion—the stick in the glass of water looks bent even though I don't believe it is, and the lines of the Müller-Lyer illusion look to be different lengths even though we know they aren't.

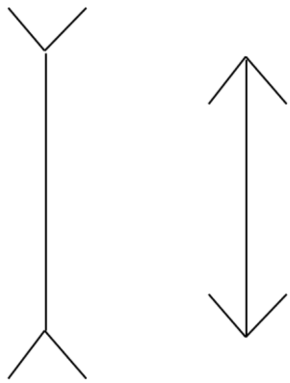

[Figure 1: The Müller-Lyer illusion]

Traditionally this has been pulled apart by separating our *judgments* about distal three-dimensional objects from what's directly *given* in the image available to us—something like sensory data or visible geometry. I take the core idea of such data is that it's something that is displayed in the two-dimensional image available from the point of view of the perceiver (what Alva Noë has nicely captured in characterizing what he calls the "snapshot" view of vision or of visual experience (Noë 2004, chapter 2)). In the case of visual size, I take this to be something like the angle-subtended "size in the image" noted above. So, a 6-inch saucer held 2 feet from my eye subtends a visual angle of about 14°, as does a 12-inch dinner plate if held about 4 feet from my eye. A snapshot taken from the position of my eye of each of these will trace out the same-sized circle on the photo, and so they should have the same size in the image.

Let me be clear: I neither mean to assume that such images must be involved in a direct way in vision nor that such snapshot metaphors are at all helpful (as explained below, I think they're actually somewhat harmful). But I do think that the use of the expression is clear enough here and evokes something about a way of thinking about vision that's commonplace, so I'll go ahead and use it.

Noë himself uses a closely related idea of "Perspectival properties" (or "P-properties"), which are taken as an ineliminable (but not, for him, exhaustive) component of visual appearance and experience. As he puts it, "*how things look with respect to size from here* ("size in the visual field")... corresponds to the size of the patch that one must fill in on a given plane perpendicular to the line of sight in order to perfectly occlude an object from view." These "perspectival sizes" are what he calls *P-properties*; and "they are themselves *objects of sight*, that is, things that we see. They are visible." So you *see* the "P-shape" projected by a plate; and with respect to trees of the same size but at different distances from you, "you can see the difference in P-size of the trees even though you also see that they are the same in size." (Noë 2004, 83)

So, as noted, one entrenched and simple way of dealing with the character of appearance is to take *appearance* as what shows up in visual geometry or P-properties, and *judgment* to be those things that we infer and make judgments about beyond what's given in the P-properties. On this view, what's given in the image for a sphere or a disc in perpendicular orientation to the line of sight is round and its size is crudely speaking the angle that subtends—or if you like, its P-size. The actual distal size that we *judge* an object to have is taken as inferred from its appearance in the image plus what we know about distances, relative sizes of objects, and perhaps all sorts of other information. Or, put slightly differently, the size that actually *appears* to us is the angle subtended, while the difference in size between objects that subtend the same angle is taken to be *judged* rather than given in appearance, using information beyond what's given in appearance.

The idea that what *appears* to us (as opposed to being *judged* or *believed* by us) at least *includes* (and maybe just *is*) what shows up in the projection or the image is not only commonplace, but persists even among those whose views on perception make significant room for other possibilities. As noted above, Noë holds on to the idea that P-properties are at least an ineliminable part of appearance. And psychologist Richard Gregory—whose views of visual perception emphasize its top-down nature—still says that "when we see a bicycle wheel from an oblique angle it *has the appearance of* an ellipse;" and that "an engine driver *sees* the rails as converging into the distance." (Gregory 2009, 135-6, my emphases) A snapshot of a bicycle wheel from an angle will certainly trace a two-dimensional form that is elliptical; but to say then that it *looks* elliptical is simply to make the assumption that I'm calling into question here—i.e. that the sameness of *projection* implies sameness of at least some salient aspect of *appearance*.

2. How Big Does the Moon Look? Some Data

I asked you earlier to consider how big the moon looks; or how big a disk you'd need to hold up at arm's length to match its projected size. Let me return to that now.

The correct answer is that a disc of about 1/4" in diameter at arm's length will match the angle subtended by the full moon. That's the size of one of those little pieces of paper you have left after you've used the hole punch. The moon subtends approximately a half of a degree, and one degree takes up about a half an inch at arm's length (depending on arms, of course, but we're just ball-parking it here). Next time you see the Moon, check it out; you might notice that that the tip of your pinky at arm's length subtends a far bigger angle than the disk of the Moon. Perhaps you are surprised by this; perhaps not. Preliminary data I've gotten suggest that for many people, it's the former.

Here are two bits of small-scale data collection on intuitions about the apparent size of the moon. Both were done in college classrooms. The first was an open-ended survey with a small classroom group (N=14); the second was a forced multiple-choice survey in a larger one (N=108). In both cases, students in philosophy classes were asked the following question: *How big a disk would you need hold up at arm's length to "appear" the same size as the full moon (that is, to line up with its outline, or project the same size visual image, or subtend the same visual angle)?*

Open-ended survey: In the open-ended survey, students were asked to answer the question in whatever way seemed best to them – e.g., they could give a standard-sized object (like a U.S. quarter, or a baseball), or they could give an actual quantitative measurement if they preferred, in inches (or centimeters, or whatever). They wrote their answers down without discussion.

Results: None of the answers to the question within a factor of 2 of being correct; all were significantly larger. The smallest size given was .7" (U.S. dime), and only 7% of answers were under 1." The median size given was about 2", and over 20% gave answers larger than 5."

Forced-choice survey: Students in a large class were asked the same question, and encouraged to formulate their answer (without discussion). They were then given five options of familiar objects with their sizes (in inches) and asked to pick the one *closest* in size to what they'd thought of. The answers were collected by an electronic classroom response system (the *iClicker*).

Results: The options given and responses are summarized below:

A CD or DVD (4.75")	17%
The top of soda can (2.13")	17%
A U.S. quarter (.96")	51%
A U.S. dime (.7")	12%
A scrap from a paper-hole-punch (.25")	3%

Discussion: What should we make of these results? Clearly almost nobody gets it right, even in the forced-choice case. A quite small number place it even within a factor of three, and a quite significant number are off by a factor of eight or more. Overall, the answers under both conditions show the subjects to be quite poor at judging the relative sizes of the angles subtended or P-properties exhibited by the moon and the various objects in this situation.

A small aside on the two experiments: Although the gross results are similar in the two cases, the median answer is clearly lower in the forced-choice case (around 1" as opposed to around 2" in the open-ended case). I suspect this is an anchoring phenomenon, as the 1" answer was the middle answer, and the highest-end possibility (4.75") was significantly smaller than the largest self-generated sizes.

In any case, the results overall suggest that at least in this particular kind of situation, people are kind of awful at judging the relative sizes of the angles subtended or P-properties. By itself, it's a counterexample to the claim that apparent size overall is both fixed by size in the image (or P-properties) and consciously available for introspective knowledge. We at least sometimes to get "size in the image" or angle subtended or maybe even P-properties quite wrong—that is, we judge that things that share P-size (like the Moon and the punch-hole paper at arm's length) *don't* share them, and that things that don't share them (like the Moon and the quarter at arm's length) *do*.

How general is this? Access to P-properties is clearly not a requirement of our normal visual experience; but whether the degree of mismatch on apparent size seen here will show up across lots of other situations is still an open question. My own suspicion is that this is quite general. I can't tell you when the nose of someone close to me subtends a bigger or smaller angle in projection from here than the whole face who is farther away, or whether a figure of a person across the room is taller or shorter in the image than my thumb at arm's length (at least, not without lining them up and maybe squinting.) If I were directly aware of "image sizes in appearance," I should be able to do that, at least far better than I can. Size "in the image" just seems like it is often not cognitively available to us.

3. Persisting Illusions More Generally

Before saying why this matters in the bigger picture for our views on perceptual consciousness, it's worth mentioning some ways in which this shouldn't be at all surprising. After all, the same phenomenon considered here shows up in general in the persistence of illusions.

So, for example, in the Müller-Lyer illusion (above), the double-forked line *appears* longer than the arrow-ended line. But of course their P-sizes are exactly the same. Even though we (being familiar with the illusion) *know* that, and so *judge* them to be the same length, appearances (or maybe our visual systems) just won't listen. Knowing it's an illusion doesn't change its appearance.

Or, take the illusion most closely related to our current case: The traditional Moon illusion. That the Moon appears larger at the horizon than it does higher in the sky has been noted since at least Aristotle. The recognition that this is fundamentally a phenomenon of perceptual psychology was well-established in the scientific community by the 17th century (largely by noting that the angle subtended by the Moon was constant across various heights in the sky); while the details of how this works have been debated over the last 100 years, perhaps now coalescing around our contemporary conventional wisdom that this results from using a variety of cues about distance that come from visible intermediate objects (and perhaps other sources) (see Egan 1998; Ross & Plug 2002).

Still, the earlier idea that the illusion is generated by a kind of atmospheric refraction is still quite commonplace in the population at large. In fact, in the forced-choice experiment discussed above, the students were asked the follow-up question "Would you have to use a different-sized object to match/line up with the moon when it's higher or lower in the sky?," 42% of the students picked "Yes, a smaller one when the moon is higher," 37% picked "No, the same," and 15% picked "Yes, a bigger one when the moon is higher." That's a lot of endorsement of the view that the angle subtended by the Moon is affected by its height in the sky, whether by "atmospheric refraction" or some other means.

Whatever folk explanation is offered, the fact is that the Moon subtends the same angle and presents the same P-size at the horizon as it does higher in the sky, but for at least many people appears larger at the horizon. But the larger-Moon-at-the-horizon persists as an illusion even though both the P-size and the judged size (we don't believe the Moon itself is *actually* bigger at the horizon) remain the same. And in this way, it's like other cases of persistent size illusions.

4. So What?

The concrete results here about the case of the Moon, the phenomenon of persistence of illusion, and the more impressionistic phenomenological reports about the relative size in the image of thumbs and figures and noses and faces would seem to suggest that there's a pretty widespread lack of knowledge of P-size. How might this bear on the idea that P-size is a feature of appearances?

I suppose one *could* try to dig in and conclude that we're terrible at knowing how big things look, since obviously the Moon and the quarter inch disk at arm's length *do* look the same size in the end. We could try in that way to hold on to the idea that the size that things appear (or at least some given aspect of that) is a straightforward function of the visual angle they subtend—that is, of their P-size. And one can find the occasional philosopher who will do this (see, for example, Schwitzgebel 2013).

But for many, this seems like a hard line to take. How things appear has been traditionally the kind of phenomenon about which we've take ourselves to have (nearly) infallible knowledge (as opposed to our quite fallible knowledge of the distal objects of perception). Appearance has been taken as what's in some sense *given* to us; and whatever mistakes we might make about the objects, the idea that we make rampant or even nearly universal mistakes about how things appear should seem pretty problematic. And all of the examples of how things look to us that I've discussed so far (including the Moon illusion, the Müller-Lyer, and the like) would have to be written off as cases where we are just wrong about how things look – wrong about the forked line looking longer than the arrow-headed one, for example.

Barring this, it seems like we need to reject the idea that how big something looks—or even an experientially available part or aspect of how it looks—is simply a matter of the angle it subtends (or its P-size). This fits nicely with the cases at hand, and overall is a natural part of a generally anti-snapshot view of visual experience.

This shouldn't really be seen as surprising or puzzling as a fact about perception. Our visual system is really good at integrating lots of subtle cues about a scene into perceptual information about the sizes and distances of objects. Binocular stereopsis, defocus blur, motion, and experience with the objects from various perspectives all work together to solve the problem of figuring out the three-dimensional scene (or at least the aspects useful to us in it) that we face. But it's an open question which pieces of the overall information that the visual system has access to and uses in this process are passed along in a way that shows up directly in the character of our visual experience.

Visual processing may not be an entirely informationally encapsulated system, but it clearly uses far more information than it passes along to the character of visual experience in an explicit way. So, for example. information about binocular disparity is clearly used in perceiving depth and distance, but our unified gaze normally has access to the depth but not to the binocular disparity itself. Taking seriously the *task* analysis of vision (in the sense of Marr 1982), we can see vision as working to provide us a characterization of the three-dimensional environment we encounter. It makes use of lots of information, some of which we might think of as visible geometry, features "in-the-image," or P-properties. But whether it passes along to our visual *experience* any particular features it might detect and use is an open and empirical question. In principle, it needn't tell us much about its internal doings.

Ignorance about the P-sizes in the case of the Moon (and others mentioned above) suggests that the information passed along to experience may not include these. And if it's not available to our judgment at all, including our judgments about appearances, it seems like a particularly bad candidate for something that is to constitute a central piece of the character of visual experience in the way we wanted from the notion of appearance.

Although I won't argue the point here, I think this will turn out to be the case for some other features of stimuli that have often been taken as features of appearance, such as color and lightness; and perhaps in other modalities of perception, similar patterns will be seen. The features of appearance are, I suspect, far more entrenched in our engagement with the world more broadly than with the explicit character of the proximal stimuli we pick up.

But for now, I'll leave it at this: At least sometimes, P-properties are not visible at all, but invisible, rather hidden behind a veil of perception – not one separating our consciousness from the distal world, but one separating it from our own proximal stimuli.[1]

References:

Egan, Frances. 1998. "The Moon Illusion." *Philosophy of Science* 65: 604-623. Chicago: University of Chicago Press.

Gregory, Richard. 2009. *Seeing Through Illusions*. Oxford: Oxford University Press.

Marr, David. 1982. *Vision: A Computational Approach*. Cambridge, Massachusetts: M.I.T. Press.

Noë, Alva. 2004. *Action in Perception*. Cambridge, Massachusetts: M.I.T. Press.

[1] Thanks to P. D. Magnus and Nick Jacobson for their feedback on earlier drafts of this paper.

Schwitzgebel, Eric. 2013. *Perplexities of Consciousness*. Cambridge, Massachusetts: M.I.T. Press.

Ross, Helen. and Plug, Cornelius. 2002. *The Mystery of the Moon Illusion: Exploring Size Perception*. Oxford: Oxford University Press.