

# Are Perspectival Shapes Seen or Imagined? An Experimental Approach\*

John Schwenkler (Florida State University) and Assaf Weksler (University of Haifa)

This paper proposes a novel experimental approach that would help to determine whether perspectival shapes, such as the elliptical profile of a tilted plate or coin, are part of perceptual experience. If they are part of perceptual experience, then it should be possible to identify these shapes simply by attending appropriately to them. Otherwise, in order to identify perspectival shapes they must first be constructed in the visual imagination. We propose that these accounts of perspectival identification can be tested by measuring the interference between visual and verbal working memory load, respectively, and the identification of perspectival shapes in the appearance of a 3D object.

## 1. Introduction

Philosophers and psychologists often characterize the visual experience of 3D shape as involving a duality of aspects. One of these aspects corresponds to the objective shape of visible objects as they exist in the environment at a distance from the viewer: for example, the rectangular shape of the surface of a desk, or the circular shape of a coin. The other aspect concerns what we will call the *perspectival* shape of an object, i.e. to the viewer-relative, 2D shape that would outline or occlude an object from a given point of view<sup>1</sup> For example, the surface of a desk will have a perspectivally rectangular shape if it is seen from straight-on, but its perspectival shape will be that of an irregular quadrilateral or trapezium when it is seen from a different angle; and the perspectival shape of a coin will be circular when viewed from straight-on but elliptical from a more sideways angle.<sup>2</sup>

Let us give the name *perspectivalism* to the thesis that visual experience has these two aspects. Many contemporary analytic philosophers endorse perspectivalism: for example, Brogaard (2010), Cohen (2010), Fish (2009), Harman (1990), Hellie (2006), Hill (2009), Lycan (1996), Noë (2004), Schellenberg (2008), Smith (2002), and Tye (2002). So do many experimental psychologists, including Gibson (1950), Mack (1978), Palmer (1999), and Rock (1983). By contrast, a

---

\* Both authors contributed equally to this work.

<sup>1</sup> We use “perspectival shape” instead of “apparent shape” since there is a clear sense in which the object shape is *also* apparent: e.g. the white face of the cube in Figure 1 clearly *looks* square, and the billiards rack *looks* like an equilateral triangle, etc.

<sup>2</sup> The properties of distance and slant (in 3D) are viewer-relative but do not count as perspectival properties for our purposes. That is because we consider them instead to be part of our experience of the 3D world around us. For example, the rim of a cup, when viewed obliquely, looks circular, slanted in depth, and located some distance away. What we call “perspectival properties” are (roughly) those that a painter would draw on a 2D canvas, i.e., those that 3D objects project onto a flat medium, perpendicular to the line of sight. These do not include 3D slant and distance.

comparatively smaller number of philosophers and psychologists have rejected (or at least doubted) this thesis and endorsed (or at least expressed sympathy towards) *anti-perspectivalism* instead (Briscoe 2008; Gibson 1986; Hopp 2013; Schroer 2008, 2017; Schwitzgebel 2011; Siewert 2006). According to anti-perspectivalists, we do not literally *see* perspectival shapes, nor does the ability to *recognize, draw, or describe* the 2D shape that would occlude an object in one's line of sight depend on a visual experience of any such shape.

The debate between perspectivalism and anti-perspectivalism is a complex phenomenological debate that has proven recalcitrant to the methods of philosophy. Can philosophical disputes of this sort—and this dispute in particular—be “translated” into a scientific dispute? Tradition has it that the answer is negative. After all, the dispute belongs to the realm of phenomenological, conceptual, and metaphysical analysis. Ask a scientist about it, and she will probably reply with comments about cognitive mechanisms and performance in some task, which in turn will likely strike the philosopher as orthogonal to her concerns.

Our paper models a two-stage strategy by which a philosopher can overcome this apparent impasse. First, she may show that the competing philosophical views implicate different *mechanisms* (or kinds of processes) in the accomplishment of a certain *task*. Once this is done, the dispute will “belong” to science in some minimal sense. The philosopher can now present her puzzle to the scientist in a way that makes sense to her, as a dispute within cognitive psychology. Second, she may then propose experiments that would help to reveal which of these mechanisms are operative in performing the task in question.<sup>3</sup> To anticipate, in this paper we suggest that the *task* is to identify perspectival shapes, the possible *mechanisms* for identifying it are selective attention (according to perspectivalism) and constructive imagination (according to anti-perspectivalism), and the *experiments* for testing which mechanism is operative use dual-task paradigms, pairing perspectival identification with a working memory task.

One might question the need for second stage in the strategy just outlined. Isn't the first stage enough for successful “translation” of a philosophical problem into a scientific one? We think not. The craft of experimental psychology is precisely to take a complex theoretical issue and squeeze it into the constraints of the lab--constraints which often are, to philosophical eyes, simply unbelievable. The experiments often involves measuring, in milliseconds (!), the response time of key presses, to very simple displays, themselves presented for a few hundred milliseconds or less. Furthermore, the task must be designed in such a way that the participants will be motivated to use a specific strategy to solve it – the strategy that the experimenter wants to probe. If there are shortcuts, the participants will take them, making the results useless. Also, all experiments must contrast a target group with a control group, which is identical to the target group except for the *specific* variable the experimenters need to measure. Finally, experiments often work in one direction only. That is, experimenters manipulate a variable and look for an effect. If it is found, then the experiment is successful. If it is not found, then it is not clear

---

<sup>3</sup> For earlier examples of this strategy, applied to debates about the content of perspectival experience and the thesis of attentional transparency, see Weksler (2016) and (2017), respectively.

that anything could be inferred, since any number of things can be invoked to explain this.

In short, it could be that one succeeds in translating the philosophical dispute into psychological categories, but not in a way that situates the dispute within the constraints of experimental practice. In that case, science will not be able to contribute to the debate, and we must continue conducting it from the armchair. For this reason, this paper will develop our experiment sketches in some detail in order to make a more convincing case that they could actually be run in a way that might yield significant results.

Note that despite using the word “translation,” we do not mean to suggest that philosophical distinctions, e.g., between imagination and attention, are simply identical the categories of cognitive psychology. For all we are saying, the connection between imagination and working memory load, and the connection between attention and working memory load, are both contingent. Perhaps in other possible worlds imagination and attention do not interact with working memory at all. What matters is that in the actual world they do. Without a doubt, it would be interesting to try to analyze imagination and attention, in the philosophical sense, by reference to working memory, but this project is not our aim here.

Section 2 of our paper focuses on the first stage of the strategy described above, namely mapping the philosophical dispute onto scientific categories. Section 4 will focus on the second stage, namely experimental design. In between, in Section 3, we explain why we do not think that the dispute between perspectivalism and anti-perspectivalism can be settled simply by appeal to introspection.

## 2. First Stage: Philosophical-Scientific Mapping

Consider the shapes shown in Figure 1. You will likely be able to recognize easily enough that the white face of the cube shown on the left is *objectively* a square, and the shape of the billiards rack on the right is *objectively* an equilateral triangle, though with curved corners. But likely you can also recognize, with comparable ease, that the *perspectival* appearance of each of those shapes is irregular: for example, you can see that if you had to draw the cube as it looks to you, or match its appearance with a corresponding line drawing, the corners of the top face at the top left and bottom right of the image should be more acute than those at the top right and bottom left; and that if you had to draw the billiards rack, or match its appearance with a line drawing, the line at the bottom would be longer than the other two, and the top angle wider than the ones at the bottom.<sup>4</sup> If you can make these judgments, you are capable of what we will call *perspectival identification* of

---

<sup>4</sup> Notice that we are not committed to saying that your perspectival identification judgment is accurate. Indeed, Thouless (1931a, 1931b) has famously shown that perspectival identification reports are distorted in the direction of the objective shapes. For example, while people can judge that the left perspectival angle of the billiard rack in Figure 1 is smaller than 60 degrees, they probably overestimate its size (it is actually 45 degrees). For our purposes, it is enough that people can identify the perspectival shape of an object at least roughly.

the shapes in a visual scene. And this is the task which we propose will connect the philosophical dispute over perspectivalism to the categories of experimental psychology.




Figure 1: Two images. RH image credit: The authors. LH image credit: Wikimedia Commons.

Now let us ask: What underlies our capacity for perspectival identification? It might seem that the *only* possible explanation for it is that perspectival shapes are, at least when we direct our attention appropriately, part of the content of visual perception.<sup>5</sup> This would render incredible the anti-perspectivalist position of J.J. Gibson in *The Ecological Approach to Visual Perception*:

No one ever saw the world as a flat patchwork of colors—no infant, no cataract patient, and not even Bishop Berkeley or Baron von Helmholtz, who believed firmly that cues for depth were learned. The notion of a patchwork of colors comes from the art of painting, not from any unbiased description of visual experience. (Gibson 1986, p. 286)

Similarly, Eric Schwitzgebel, in his *Perplexities of Consciousness*, writes:

What exactly is my sensory experience as I stare at a penny [obliquely]? My first and recurring inclination is to say that the penny looks just plain circular, in a three-dimensional space--not elliptical at all, in any sense or by any effort I can muster. (Schwitzgebel 2011, p. 19)

Gibson's position, and the position that Schwitzgebel is inclined to endorse, is radical. It is different from the position that Gibson had advanced in *The Perception of the Visual World*, according to which it *is* possible to perceive a visible scene "as if

---

<sup>5</sup> Thus Cohen (2010) says (in our terminology) that the fact of our capacity for perspectival identification "gives us reason" to believe that visual perception represents perspectival properties. However, he does not claim that this is the *only* way to explain perspectival identification.

it consisted of areas or patches of colored surface, divided up by contours”, as long as one adopts the appropriate perceptual attitude:

To do so, you must fixate your eyes on some prominent point and then pay attention not to that point, as is natural, but to the whole range of what you can see, keeping your eyes still fixed. The attitude you should take is that of the perspective draftsman. ... If you persist, the scene comes to approximate the appearance of a picture. You may observe that it has characteristics somewhat different from the former scene. This is what will here be called the *visual field*. It is less familiar than the visual world and it cannot be observed except with some kind of special effort ... It is strictly an introspective or analytic phenomenon. One gets it only by trying to see the visual world in perspective and to see its colors as a painter does. (Gibson 1950, pp. 26-27)

The Gibson of 1950 puts forward what we will call an *attentional* account of perspectival identification, according to which perspectival properties can be identified by adopting a distinctive attentional focus. The attentional account holds that, at least when this focus is adopted, perspectival shapes and colors are *seen*--it is an *aspect of visual experience* that we identify from within this introspective or “painterly” attitude. On a view like this, identifying perspectival properties amounts to a Stroop-like task. In a Stroop task one must classify stimuli correctly despite the influence of incongruent and dominant perceptual representations. For example, in the numerical Stroop task (Figure 2), comparing physical size is difficult because the numerical value interferes. To succeed in the task, subjects must filter out the numerical value using selective attention. And Gibson’s 1950 description of perspectival identification suggests something similar: it’s a matter of selectively *attending* to the perspectival characteristics of the visual field, and overcoming the natural habit of attending instead to the objective shapes, colors, etc. that are our usual objects of concern. If this account is correct, then perspectival identification requires selectively attending to certain visible features while suppressing attention to others.<sup>6</sup>

Congruent pair	Incongruent pair	Neutral pair – Numerical task	Neutral pair – Physical task
<p style="text-align: center;"><b>3    5</b></p>	<p style="text-align: center;"><b>5    3</b></p>	<p style="text-align: center;"><b>3    5</b></p>	<p style="text-align: center;"><b>3    3</b></p>

---

<sup>6</sup> To clarify: We are *not* putting this account forward in the spirit of interpreting Gibson’s actual views. These were quite complicated, and settling what they were falls outside the scope of this paper.

Figure 2: The numerical Stroop task (Henik and Tzelgov 1982). Participants are faster to recognize the digits or classify their sizes correctly in congruent trials (the larger number is physically larger) than incongruent ones. Image credit: Wikimedia Commons.

By contrast, anti-perspectivalists like Schwitzgebel and the later Gibson must reject the attentional account, since on their view the “patchwork of colors” that one would paint in order to depict a visible scene, or describe in identifying its perspectival appearance, is not something we can *see* at all (with or without a special exercise of visual attention). How else might they account for our capacity for perspectival identification? We are not aware of any place where Gibson or Schwitzgebel address this question directly, but a possible answer to offer on their behalf is that the “visual field” of perspectival shapes and colors is something *imagined*, rather than seen. Consequently, perspectival identification depends on a top-down process in which a two-dimensional patchwork corresponding to a visible 3D scene is constructed in the eye of the mind.<sup>7</sup> Such a *constructive* account of perspectival identification has been proposed by Robert Briscoe, who describes perspectival identification as a perceptual-cum-imaginative process of “make-perceive”:

In general, one engages in make-perceive when one projects or ‘superimposes’ a mental image on a certain region of the visually perceived world. Thus, as Rick Grush observes, one might use a mental image to decide where in egocentric space a vase should be placed in order best to obscure a picture on the desk (2004, p. 390). Similarly, when deciding how to arrange the furniture in a new home, one might begin by visualizing how things would look were one to place an armchair in a certain empty corner, or were one to hang a painting on a certain unadorned wall. It seems clear that one does not in any literal sense *see* the visualized armchair in the empty corner or the visualized painting on the wall. Perhaps the best and most familiar example of make-perceive is the experience of noticing a constellation in the night-time sky. Noticing a constellation is a hybrid visual-imaginative experience: it involves both *seeing* the stars in the constellation and *imagining* the lines that connect them at the same time. (Briscoe 2008, p. 479)

In the present case, “make-perceiving” an object’s perspectival shape could involve, as Briscoe puts it, pretending to trace an outline “on an imaginary pane of glass” (p. 480), or imagining the shape of a planar figure which, if stood upright and placed

---

<sup>7</sup> Noë (2004) holds both that sensory experiences are ‘perspectival’, and also that such ‘perspectival sensing’ plays a role, i.e. has a function, in everyday perceiving. As Noë notes, this is in keeping with long tradition in philosophy and in psychology. Our definition of the attentional and constructive accounts is supposed to be neutral on this issue. Seeing or visualizing perspectival shapes may or may not play a role in everyday perceiving. If the reader doubts that *imagined* perspectival shapes could play a role in everyday perception, we invite her to consider the phenomenon of amodal perception. Arguably, it involves mental imagery (Nanay 2010) and, nevertheless, it plays a role in everyday perception. For example, an amodally completed object can be the focus of object-based attention (Lee and Vecera 2005).

appropriately, would perfectly occlude a certain portion of the perceived scene. Each of these would be a way to construct within the visual imagination a 2D shape corresponding to the appearance of an object from a certain angle. While this construction draws on materials that are provided in visual experience, there is an important respect in which it goes beyond them. And it seems possible in principle to account in this way for the ability to draw, recognize, or describe the perspectival shapes of objects in a visible scene, even on the assumption that those shapes are not elements of visual experience itself.<sup>8</sup>

We propose, then, that this contrast between *attentional* and *constructive* accounts of perspectival identification amounts to a translation of the perspectivalism/anti-perspectivalism dispute into the categories of cognitive psychology, as attention and mental imagery are topics routinely studied by cognitive psychologists. This completes the first stage of the strategy described above.

### 3. The Limits of Introspection

Sometimes it seems easy to tell whether a given perceptual judgment rests on top-down visual construction rather than attention to a feature that is present in visual experience. For example, consider the tasks illustrated in Figure 3. It seems clear that each of (a), (b), and (c) requires an exercise of visual imagination that goes beyond what is presented immediately in visual experience. In (a) this process seems to involve tracing a path through a maze in imagination,<sup>9</sup> in (b) it involves repeatedly multiplying a percept in the imagination until a certain length is reached, and in (c) it involves imaginatively rotating one percept in order to compare it with a different one. Things are not so clear when we consider the Kanizsa illusion in (d), where the illusory triangle seems as if it is “seen” to be present simply in the original experience, even though some of its contours are no part of the visible scene itself. In each of (a), (b), and (c) the task involves a *constructive* process: it involves something more than just categorizing what is already seen. It is harder to say whether or not the same is true of (d) as well.

---

<sup>8</sup> For another constructivist account of perspectival identification, see Schroer (2017). Unfortunately this paper came to our attention too late for us to consider it in our essay.

<sup>9</sup> Or, perhaps, “filling in” the maze from the origin point until an exit is reached.

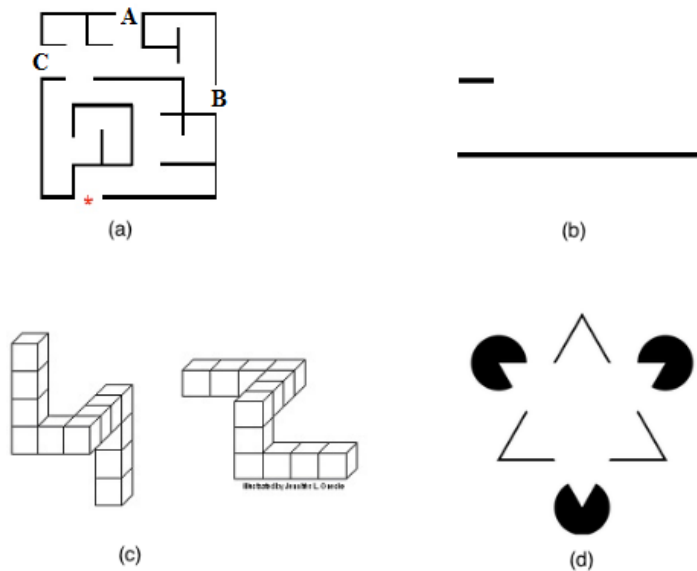


Figure 3: Examples illustrating the distinction between visual perception and visually-based “make-perceive”. In (a) the task is to say which, if any, of A, B, and C can be reached via the entrance to the maze. In (b) the task is to say how many times longer the bottom segment is than the top one. Figure (c) is a classic mental rotation task: Are the two shapes equivalent? And (d) is the Kanizsa Triangle—a ‘borderline’ case that is harder to characterize confidently just through introspection. Image credit for (a) and (b): The authors. Image credit for (c) and (d): Wikimedia Commons.

Suppose for the sake of argument that we are justified in thinking, on the basis of introspection, that the first three tasks in Figure 3 rely on constructive visualization, but that we are not so justified in being confident either way about the Kanizsa illusion.<sup>10</sup> How about the question of whether or not perspectival identification requires a process of constructive visualization? Some philosophers seem to think that this is a question that introspection can decide on its own. For example, here is Alva Noë’s response to “philosophers [who] will not scruple to acknowledge commonplaces such as this”—where *this* is the claim that e.g. a circular plate seen from an oblique angle *looks elliptical*:

It is just not true that the plate looks elliptical, [anti-perspectivalist philosophers] will say. But how can we take this seriously? Certainly it is the case that we wouldn’t likely judge the plate to be elliptical, on the basis of how it looks. Nor is it likely that we would say that it looks elliptical. But surely it does look elliptical from here! (Noë 2004, 78)

What could ground Noë’s confidence that the thesis he defends is “certainly” or “surely” correct? It seems most charitable to take this statement as a

---

<sup>10</sup> To be clear: nothing in our argument will rest on this assumption. Indeed, the experimental methods we explore below could give a way of testing it.



*phenomenological* description of the character of his visual experience, based on the reflexive awareness of experience that can be secured through introspection. This is the same kind of awareness that justifies saying that, e.g., the correct path in Figure 3a is not part of our visual experience of the maze, but does not justify a similar verdict about the white triangle in Figure 3d. Except in the face of special evidence to the contrary, we tend to trust one another's introspective reports about what our experience is like, and about how we go about solving various perceptual tasks. Is there any such reason to doubt what Noë says in the quotation above?<sup>11</sup>

Some of the possible reasons for skepticism have already been discussed in important work by Eric Schwitzgebel (2011). For example, if it really were *certainly* or *surely* the case that the surface of a coin seen from an oblique angle looks elliptical, then we would expect there to be widespread agreement about this—that is, agreement not on *which perspectival shape would correspond* to the coin's surface, but on whether that shape is a feature of visual experience itself, prior to any constructive exercise of the visual imagination. Yet as we have seen, there is no widespread agreement about this matter, even among philosophers and psychologists who have given it a lot of consideration. As Schwitzgebel (2011) points out, this disagreement is *prima facie* evidence for the unreliability of introspection in this case.

As further evidence of the difficulty of settling this question just by an appeal to introspection, consider again the images shown in Figure 1, and try to answer the following questions about them:

- (1) Are the perspectival angles at the top right and bottom left of the cube equal, or is the top one slightly wider?
- (2) Is the perspectival angle at the top of the billiards rack acute, obtuse, or 90 degrees?

As you answer these questions, do you find it *certain* or *obvious* that all you have to do is attend selectively to an aspect of your visual experience, and not construct a further shape in the imagination? When we shared an image like that of the billiards rack in an informal poll on Facebook and asked for descriptions of its perspectival angles, respondents offered a diverse range of explanations of how they approached this task, some of which sounded primarily attention-driven while others suggested processes of constructive visualization or “make-perceive”. Here is a representative sample of those replies:

- ... I imagine various angles and try to reconcile them with my perception of the image. Knowing that the object would be an equilateral triangle in actuality

---

<sup>11</sup> To clarify, this question concerns not “first-order” certainty that the perspectival shape of the plate in the world is elliptical (though perhaps we should not have such certainty, because perspectival identification is somewhat inaccurate, as discussed in fn. 4 above). The question concerns Noë's “second-order” certainty that the perspectival shape is seen (even if inaccurately) rather than imagined.

impedes my ability to imagine other angles and restrict them to the 2d image (indeed I have to ignore portions of the image to do so).

- I had to mentally rotate it to the perpendicular, otherwise the acute angle of the 3D shape tends to override the perspectival appearance. I also have to really attend to just that corner, otherwise it's difficult not to see it as just an equilateral triangle.
- Perspectively [the angle] appears to me to be obtuse. I come to this conclusion by rotating the image so that the side that appears longest is vertical. The other side is tilted above horizontal so the apparent angle is greater than 90.
- Initial vision suggested [that the angle was] acute, but I now see it as obtuse after sort of comparing its sides to the rectangular outline in which FB places the image.
- I've spent part of the morning working in Light Room recropping photos so that they're straight up and down. The program uses a grid overlay, and you tilt the image using some vertical part of the image. And I felt like I was just doing the same thing. But I think the edges of the frame are what I used.
- It appears to be an equilateral triangle on casual viewing and one has to step back and view it in a very detached way to see it as obtuse.

For the reasons explained above, we believe that each of these descriptions should be treated with considerable skepticism. But so, by the same token, should Noë's introspective judgment about the presence of an elliptical shape in his visual experience when he looks from an oblique angle at a plate. Mere introspection is not sufficient to distinguish between the perceptual, cognitive, and imaginative processes that we may bring to bear in the task of perspectival identification.

#### 4. Second Stage: Possible Experimental Approaches

We are considering the dispute between an *attentional* account of perspectival identification, on which perspectival shapes can be identified by attending appropriately to an aspect of visual experience, with the *constructive* account, on which perspectival shapes are not seen but rather visualized, in the same way as one can construct in the imagination the solution of a maze or the appearance of a 2D or 3D shape that has been rotated in space. We argued in section 3 that introspection alone cannot decide which of these accounts is correct. The present section will describe a possible *experimental* approach to the question, which turns on the possibility of operationalizing the roles of visualization and selective attention in perspectival identification.

Simplifying things *quite* a lot, the experiments we describe will test the following predictions. First, since on the constructive account perspectival identification requires top-down construction of a mental image that is "superimposed" onto one's perception of the visual world, **the constructive account predicts that there will be interference between perspectival identification and other tasks that draw on the same cognitive resources as constructive visualization.** And second, since on the attentional account perspectival identification requires *selective*

*attention* to perspectival shapes, **the attentional account predicts that tasks that interfere with selective visual attention will also interfere with perspectival identification.** Each account can be falsified by testing these predictions.

Sections 4.2 through 4.4 describe a series of experiments that could do this. First, in section 4.1 we consider a related proposal by Sean Kelly, explaining why it is inadequate to our purposes.<sup>12</sup>

#### 4. 1 Kelly's proposal

Sean Dorrance Kelly proposes a way of experimentally probing the presence of perspectival shapes in visual experience in his contribution to a 2008 symposium on Noë's *Action in Perception*. Referring to the passage from Noë's book that we quoted in section 3, Kelly writes that if Noë's claim about the experience of perspectival shapes is correct, "there ought to be some evidence of this experience" that an appropriate experiment could detect (Kelly 2008, p. 687). Specifically, Kelly proposes to investigate whether there is a *priming* effect of perspectival shapes on subsequent judgments of 2D stimuli that correspond to them. He writes:

The experiment I imagine will make use of the primed matching paradigm, and it will go like this: first present the priming stimulus, a circular object like a plate presented at an angle ... Then present as targets, among other controls, a pair of ellipses and a pair of circles. If Noë is right, then judgments about the similarity of the ellipses should be facilitated equally if not more than judgments about the similarity of the circles. I predict, however, that judgments about the similarity of the circles will be primed, but judgments about the similarity of the ellipses will not. (ibid., pp. 688-689)

Kelly also reports that the preliminary results of an experiment along these lines supported his position over Noë's: "using a picture of a penny at an angle as the prime we get a 10% facilitation effect for similarity judgments regarding circles, and no effect whatsoever—neither facilitation nor interference—for similarity judgments about ellipses" (ibid., p. 689). This, he says, puts the burden of proof on Noë to explain why, if perspectival shapes are part of ordinary visual experience, they do not give rise to a priming effect.

While we share Kelly's desire to look for experimental evidence that bears on this dispute, the logic of his proposed experiment can be challenged in both directions. First, as Kelly admits (2008, p. 687) a finding that perspectival shapes

---

<sup>12</sup> Before moving on, let us briefly explain why work on distortions in perspectival identification from the first half of the twentieth century (starting with Thouless 1931a and 1931b) seems to us to be neutral with respect to the attentional/constructive dispute. This body of work has apparently shown that perspectival identification is distorted by cues for depth or shape constancy. That could be because subjects are make-perceiving the visual field on the basis of a visual experience that represents three-dimensional, "objective" features exclusively. Or it could be because the computational processes that generate the perspectival dimension of experiences are influenced in some way by 3D information (a view like this is endorsed by Hill 2009). So the body of work in question is apparently consistent with both the attentional and the constructive views.

have a priming effect on subsequent judgments could be explained within an anti-perspectivalist framework by the hypothesis that *non-conscious* or subliminal representations of perspectival shape, e.g. representations in “early vision” that correspond to retinal input rather than distal properties of the environment, influence subsequent perceptual judgments (for an example of subliminal visual priming see Bar and Biederman 1998). Given this possibility, if Kelly’s proposed experiment were to be run again with a larger number of participants and find evidence of a priming effect by perspectival shapes, this would not count as evidence in support of perspectivalism.<sup>13</sup>

Second, Kelly’s preliminary finding that there is no such priming effect could in principle be explained by the hypothesis that the relevant priming effects are modulated by the task-relevance of the priming stimuli. And there is some evidence for this possibility: for example, using a masked priming paradigm in which a priming stimulus was shown very briefly followed by a target stimulus that had to be identified, Enns and Oriet (2007) found that identification of the target was facilitated only when the prime resembled it in respects that were explicitly relevant to the classification task. Similarly, Pohl et al. (2010) found that subconscious primes facilitated identification of a target stimulus only when the prime resembled the target in a task-relevant dimension.<sup>14</sup> Finally, reviewing a large corpus of data from social psychology, Eitam and Higgins (2010) develop the ROAR (Relevance of a Representation) framework, on which, *in general*, priming depends on relevance. Given these findings we should question whether Kelly’s preliminary finding that perspectival shapes do not prime subsequent judgments supports the conclusion that these shapes are absent from ordinary experience--as instead this lack of an effect could be due to the task-irrelevance of perspectival shapes. In Kelly’s study, as we understand it, participants were presented with tilted objects (as primes). Subsequently, they had to compare ellipses, circles, and other shapes to each other. At no stage in the experiment were participants instructed to focus on perspectival shapes, and nothing in the task required them to do so. All this gives reason to suspect that the visual system had deemed perspectival properties irrelevant in the context of Kelly’s study, and this task-irrelevance could explain why they did not give rise to any priming effect.

Finally, notice that the dispute between Kelly and Noë concerns a different question than the one we are focused on here. Kelly is a perspectivalist in the sense we have defined: he holds that “it is possible to experience [the] apparent ellipticalness” of e.g. a tilted plate (2008, p. 685), and he seems to regard this experience as

---

<sup>13</sup> Note that while falling short of establishing perspectivalism, finding the priming effect in question would establish something interesting, namely that visual representations of perspectival shapes activate shape concepts. Once activated, such concepts can facilitate and bias judgments related to them. On such a view, representations of perspectival shapes do more than just feed into modules that calculate objective shapes; they can influence our cognitive lives in a more global way.

<sup>14</sup> Note that there is a much larger body of literature about the effects of task-relevance on the phenomenon of “priming of pop-out,” in which a prime leads to faster location of a unique target resembling it, in a visual search task (for a review see for example the introduction in Michal, Lleras, and Beck 2014).

perceptual rather than imaginative. What Kelly denies is only that it is possible to experience both objective and perspectival shapes at once. Thus he writes:

Look at a plate presented at an angle and pay attention to its shape. When you are seeing it to be circular try also to see its elliptical aspect. It seems to me clear that focusing on the real shape pushes the apparent one out of experience, and focusing on the apparent shape likewise makes the real one go away ... The problem seems to me similar to the problem of trying to see the duck and the rabbit simultaneously in the duck-rabbit image. I can switch back and forth between them, but I can't—or at least I can't without extreme effort—experience both of them at the same time. (ibid., p. 685)

According to Kelly, perspectival shapes *can* be seen when we adopt the right perceptual attitude, and once this result has been achieved the new perceptual attitude *excludes* from experience the “real” or objective aspects of things that had been present in experience beforehand. By contrast, what Noë says is that perspectival and objective features are present in experience no matter what perceptual attitude we take. Kelly's proposed experiment is supposed to help decide between these views: he is offering a way to discover if perspectival shapes are present in visual experience even when we are not focused on them (which is exactly why he did not instruct participants to do this). However, the question we have posed is more fundamental than this: for if the constructive account is correct then what Kelly calls “focusing on the apparent shape” is really a way of *constructing* such a shape. In that case both his view and Noë's would be incorrect.

#### **4.2 Testing the role of visual working memory in perspectival identification**

Our challenge is to find a way of deciding between a *constructive* account of perspectival identification, according to which an exercise of visual imagination or “make-perceive” is involved in identifying the perspectival shapes corresponding to regions in a 3D display, and an *attentional* account, according to which perspectival identification involves attending selectively to perspectival shapes and directing attention away from objective ones. In the present section, we will explain how the constructivist proposal can be tested by employing a *dual-task paradigm* in which a perspectival identification task is paired with a task that is known to place high demands on visual working memory. We will argue that if perspectival identification involves a constructive visualization, then it should place demands on visual working memory as well, and thus the tasks should interfere with one another. Finding that there is no such interference between the tasks would undermine the constructivist account.




Figure 4: Hyun and Luck's experimental paradigm. The change-detection task required comparing the first display (leftmost box) to the last (rightmost box). Each display contains four colored squares. The imagery task required indicating whether the letter shape (central box) was canonical or mirror-reversed. In the dual-task condition, both had to be performed simultaneously. Credit: Hyun and Luck (2007).

As an illustration of how the dual-task paradigm works (and also a main reason to think that visualization draws on visual working memory), consider a study by Hyun and Luck (2007). They gave participants the task of determining whether a given 2D shape was a canonical or mirror-reversed letter -- a task that is supposed to require imaginatively rotating the figure in order to determine its correspondence or non-correspondence with the relevant letter shape. In the key manipulation, participants had to perform this task concurrently with a change-detection task that drew on visual working memory, as pictured in Figure 4: they were shown a visual stimulus that they had to remember, followed by a letter shape that was to be identified as canonical or mirror-reversed, followed by a further stimulus that was to be compared with the first one.

Hyun and Luck's study found interference in both directions, as shown in Figure 5: performance in the mental rotation task was impaired by the concurrent change-detection task, and the mental rotation task interfered with the change detection task as well. Moreover, the interference was *rotation-dependent*, as it increased as a function of degree of rotation: in the zero-degree condition, i.e. when the letter was presented in its canonical orientation, the concurrent memory task did not interfere at all with classifying the letter shape (Figure 5a) and the letter task interfered with visual memory only a small amount (Figure 5b) in comparison with the other two conditions. Additionally, the interference between both tasks was greater in the  $\pm 144$ -degree condition, when the letter was most significantly rotated, than the  $\pm 72$ -degree one. As Hyun and Luck write, this interference confirms the prediction that in the mental rotation task, "[s]ome sort of buffer [is] necessary to hold an intermediate representation of the object while it is being rotated" (2007, p. 154). This buffer is also drawn on in the visual change-detection task but not in the zero-degree condition of the mental rotation task, where the original sensory representation is sufficient to make the judgment.

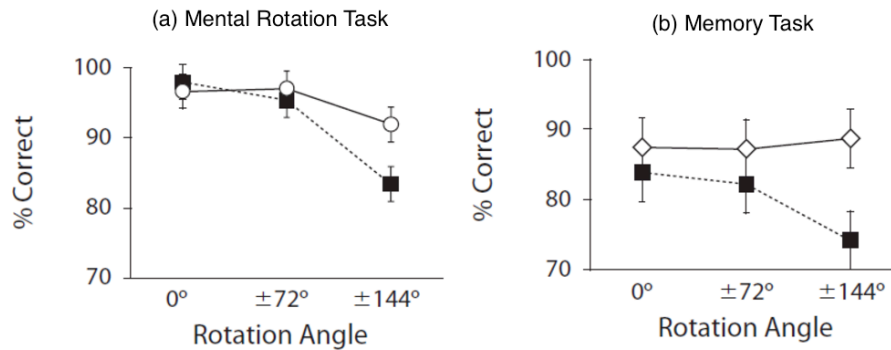


Figure 5: Dual-task interference in Hyun and Luck (2007). In (a), open circles represent performance in the mental rotation task on its own, and filled colored boxes represent performance in the mental rotation task while also performing the memory task. In (b), open diamonds represent performance in the memory task on its own, and colored boxes represent performance in the memory task while also performing the mental rotation task. Credit: Hyun and Luck (2007).

Following Hyun and Luck, we take these findings to support the hypothesis that *visual working memory* acts as a “buffer” underlying the capacity for visual imagination and visual memory. And since the constructive account of how perspectival shapes are recognized credits visual imagination with a crucial role in this process, that account yields the prediction that **the identification of perspectival shapes will interfere with concurrent tasks that place demands on visual working memory and vice versa**. This prediction can be tested by pairing a task requiring the identification of perspectival shapes with a visual change-detection task like the one used by Hyun and Luck. The display will comprise three stimuli, as shown in Figure 6A: a group of colors, followed by a photo or realistic rendering of a 3D shape<sup>15</sup> with two 2D outline shapes below it, followed by a single color stimulus. In the memory task participants will have to indicate whether the color shown at the end was included in the original configuration, and in the perspectival identification task participants will have to indicate which of the two 2D shapes match the perspectival shape of a region of the 3D image.<sup>16</sup> The *difficulty* of the perspectival identification task will depend on the congruence or incongruence between the real and perspectival shapes in the photograph: in a case where they coincide because the relevant region is displayed from straight-on, seeing the objective shape will be enough to identify the perspectival shape, and so constructive visualization should not be required (because subjects know in advance that when a shape is viewed from straight on, its perspectival shape is identical to its objective shape). However, if the viewing angle is oblique, so that the real and perspectival shapes are divergent, then there should be reliance on

<sup>15</sup> Section 4.3 will address some difficulties with using 2D images for this purpose.

<sup>16</sup> The task and stimuli illustrated assumes that retaining *color* information pertinently engages the resources underlying retention of imagined outline *shape*. We follow Hyun and Luck (2007) here. In their task (see Figure 6), participants retain color information, which interferes with retention of the imagined (rotated) letter.

visualization. As we explain just below, our claim is that the constructive account predicts that there will be a *difficulty-dependent* interference between the visual memory task and the perspectival identification task. That is, the interference will increase with the degree of non-correspondence between the perspectival shape and the corresponding objective one. Finding no such pattern of interference will count as evidence against the constructive account.

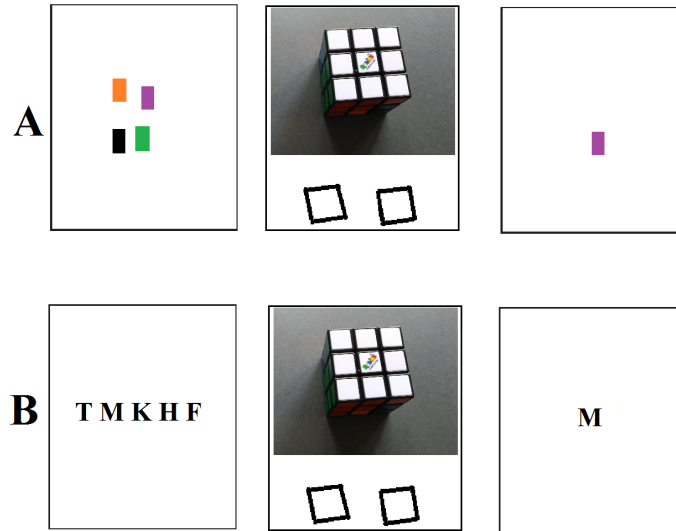


Figure 6. A sketch of the experiments we propose. Participants need to memorize colors (visual memory, A) or letters (verbal memory, B). Next, they need to indicate which of the two 2D shapes match the perspectival shape of the white face of the 3D shape. Finally, they need to indicate whether the color (letter) probe appearing in the last display matches one of the colors (letters) from the first display. In congruent trials, the 3D shape will be presented from straight on.

A point of clarification may be necessary. The prediction of “difficulty-dependent” levels of interference is needed in order to show that interference between the tasks is not due just to a *general* dual-task cost, i.e. the simple difficulty of doing two different things at once. If the two tasks draw on a common resource, i.e. the “buffer” that both holds a remembered stimulus and allows for the visual construction of a perspectival shape, then the interference between perspectival identification and visual memory should increase as the former task is made more taxing. If the level of interference were found to be independent of differences in the non-correspondence between objective and perspectival shapes, this would be inconsistent with the constructive account.

Let us sum up our proposal so far. Since, on the constructive account, perspectival identification involves an exercise of constructive visualization, the account predicts that there will be interference between perspectival identification and another concurrent task that taxes visual working memory, and that the degree of interference will increase with the difficulty of the task. We have sketched a way to test this prediction. In section 4.4 we will show how a similar series of



experiments could be used to test a prediction of the attentional account as well. First, let us consider an important methodological concern.

### 4.3 Perspectival shapes in two and three dimensions

In the experiment described in section 4.2, we proposed to use photographs or other realistic images of 3D objects as the stimuli for perspectival identification tasks. However, since images themselves are planar figures, it may seem that perspectival properties of the 3D objects that an image depicts can be seen simply by directing one's attention to the properties of the 2D display. One might worry that if this is correct, then any findings about how perspectival identification is achieved within an experimental paradigm that utilizes photographs or other 2D images might fail to generalize to other conditions. In particular, a finding that visual memory load does not interfere with perspectival identification would show only that the constructive account is incorrect concerning our perception of *pictures* of 3D objects, whereas it might be correct when 3D objects are seen "in the flesh".<sup>17</sup>




Figure 7: A photograph of a book (a) and an anamorphic transformation of it (b). In (b), the 3D shape of the book appears distorted when viewed directly, but if you look at the page from an oblique angle with the bottom of the image closer to you, the shape will appear veridical. For a full-size version, visit [tinyurl.com/yc35dwmo](http://tinyurl.com/yc35dwmo). Image credit: The authors.

One way to address this problem is to enrich the stimulus image with 3D cues, and remove 2D pictorial cues, thereby making the experience more similar to a genuine experience of a 3D object (for relevant discussion, see Vishwanath 2014). For example, because participants could gain information about the 2D shape of the stimulus by moving their head, experimenters should stabilize the head of participants. Similarly, experimenters should mask the computer edges, and/or use a dark room so that the 2D context of the stimuli will not be apparent. These two suggestions aim to remove 2D cues. It is also possible to add 3D cues by using stereoscopic goggles (for a recent study using this technology see Hibbard et al.

---

<sup>17</sup> Alva Noë makes a related point in his (2008), at pp. 694-695.

2017). These measures will probably not lead to a perfect 3D experience (for example, focus cues will be inconsistent with the experienced 3D object, see Watt et al. 2005). They are still helpful, however, because they allow experimenters to manipulate the difficulty of the perspectival identification task. For example, perspectival identification with stereoscopic goggles should be more difficult than without them. Using this manipulation, experimenters could test the prediction that there is a *difficulty-dependent* interference between perspectival identification and change detection.

We note as well that as an alternative to the strategy of manipulating 2D and 3D cues, the problem of using a flat display can also be addressed by using *inverse perspective*, or anamorphosis, to create images of 3D objects that are meant to be viewed from an angle. An example of this is shown in Figure 7. Viewed from directly ahead, the image in (a) will appear veridical and the one in (b) distorted. However, if you take the page on which this figure appears, sit down at a desk or table, and place the paper a short distance in front of you, then viewed from this angle the 3D shape of the book in the distorted image (b) will now appear veridical. However, from this angle the perspectival shape of the book's cover (see Figure 8a) is *not* the same as its shape on the page (Figure 8b) -- which makes it impossible to identify the perspectival shape just by attending to the properties of the image.




Figure 8: Perspectival shapes of the cover of the book in Figure 7, as seen from two perspectives. In (a) the shape traces the original photograph in 7a, and will correspond to the perspectival shape of the cover when the figure is viewed from the appropriately oblique angle. In (b) the shape traces the distorted image in 7b. Credit: The authors.

Our second proposal, then, is to construct the stimuli for our perspectival identification tasks using reverse perspective, and display them on a screen that is slanted away from the participant's line of sight. In this condition, participants will experience objective shapes, and will also be able to identify perspectival shapes corresponding to the appearance of the display as seen from their point of view, but the latter will not match the actual 2D shapes that are shown on the display.

#### 4.4 Disentangling the roles of working memory in perception and imagination

In section 4.2 we described a study to test the hypothesis that perspectival identification is an exercise in constructive visualization. If this constructive account is correct, then there should be interference between visual memory and perspectival identification in a dual-task condition. The dual-task experiment we described in section 4.2 was a way to test this prediction. In the present section we will describe how the same kind of experiment could be used to test one of the predictions of the attentional account, namely that **verbal memory tasks will interfere with perspectival identification more than visual memory tasks**. Before describing the experiments that could be used to test this prediction, we will explain the prediction in more detail, and also explain why we think it is yielded by the attentional account.

According to the attentional account, perspectival identification involves first *selecting* perspectival shapes as objects of visual attention, then *classifying* the properties of the shapes one attends to. The study by Hyun and Luck (2007) discussed in section 4.2 gives reason to think that the classification of a visual stimulus will not interfere significantly with visual working memory: recall that in their study there was no effect of visual working memory load on the classification of letter shapes in the zero-degree (unrotated) condition.<sup>18</sup> However, since control of selective attention requires *remembering* which stimuli are task-relevant and which are not, one might expect there to be interference between working memory tasks and tasks that require selective attention. And there is indeed evidence for this hypothesis in the experimental literature. For example, de Fockert et al. (2001) had participants perform a working memory task (remembering a series of digits) concurrently with a visual task that required selective attention (classifying a written name while ignoring a distractor face), and found that the behavioral effects of distractor interference were significantly greater in the high-load condition of the working memory task.<sup>19</sup> These findings suggest a role for working memory in the control of visual selective attention.

However, there is also compelling evidence that the *modality* of a working memory task makes a difference to its interference with selective attention -- and these differences enable us to draw out a testable prediction of the attentional account. In Baddeley and Hitch's influential (1974) account of working memory, visual and verbal working memory draw on distinct and dissociable cognitive resources specialized to carry a specific sort of information. That is, while *visual* working memory allows for the storage of the visual features of a stimulus, such as the colored box patterns used by Hyun and Luck (2007; Figure 4 above) or the dot patterns in Figure 9, *verbal* working memory allows for storage of verbal or phonological features, such as words, letters, or digits. The distinctness of these

---

<sup>18</sup> There was a small interference in the opposite direction, but Hyun and Luck attributed it to a general dual-task cost (see discussion at the end of section 4.2).

<sup>19</sup> The task in question involves names and distractor faces, which are special. But there are other studies from Lavie's lab that use ordinary stimuli (e.g., letters) and equally support the conclusion that WM load interferes with selective attention (see Lavie 2005 for a review).



forms of working memory is shown by the different ways that visual and verbal working memory, respectively, are drawn on in a given task. For example, Robbins et al. (1996) found that a concurrent visuo-spatial task has more of a disruptive effect than a concurrent verbal task on chess players' ability to remember the configuration of pieces on a chessboard and solve tactical chess positions (i.e., identify the correct move to make at a given point in a game). In Baddeley and Hitch's framework, what accounts for this difference is that visual memory relies on a "visuo-spatial sketchpad" that represents characteristics such as form, color, and movement in an internal mental image, while verbal memory is carried through an "articulatory loop" in which stimuli are rehearsed subvocally in order to guard against decay. Within this framework the explanation of the finding by Robbins et al. (1996) is thus that the visuospatial sketchpad, rather than the articulatory loop, plays a primary role in cognition in chess.

As another illustration of the distinction between the functions of visual and verbal working memory, consider a study by Sims and Hegarty (1997). Their interest was in whether spatial visualization is involved in the *mental animation* of the workings of a mechanical system. If mental animation involves spatial visualization then it should draw on the visuo-spatial sketchpad, and so there should be more interference between a primary mental animation task and a secondary task that draws on visuo-spatial memory than between the mental animation task and a secondary task that requires verbal memory instead. Sims and Hegarty's two "secondary" tasks are shown in Figure 9: the first one, memorizing a pattern of dots, draws on the resources of the visuo-spatial sketchpad, while the second, memorizing a series of letters, draws on the articulatory loop instead.<sup>20</sup> In line with the hypothesis that mental animation involves spatial visualization, Sims and Hegarty found greater interference between mental animation and memory of a dot pattern than between mental animation and verbal memory, while the opposite pattern was observed when they replaced the mental animation task with a task that required verbal reasoning: in this latter case there was more interference between the two verbal tasks than between the verbal task and the visuo-spatial one.

---

<sup>20</sup> A clarification may be necessary here. In Baddeley and Hitch's framework, what determines whether a given stimulus is stored in verbal or visual working memory is not the modality in which the stimulus is displayed, but rather the aspect of its content that is retained for later processing. For example, in the study by Sims and Hegarty (1997) the letter strings that had to be remembered in the *verbal* secondary tasks were presented as part of the *visual* display, as in Figure 9B. Nevertheless, since participants were instructed to remember the letters themselves rather than the form or color of the stimulus, they read the letters silently and thus "converted" the information into a phonological form that was then carried by the articulatory loop.

**A. Visuo-Spatial Secondary Task**



**B. Verbal Secondary Task**

True Trial:	LVYQMX	Q
False Trial:	DMKXRG	W

Figure 9: Visual (A) vs. verbal (B) memory tasks. Both tasks use visual stimuli, but task (B) is assumed to activate verbal WM (the articulatory/phonological loop), while (A) draws on visual WM (the visuo-spatial sketchpad). Image credit: Sims and Hegarty (1997).

In light of the above, return now to the study by de Fockert et al. (2001) that we used to motivate the hypothesis that increased working memory load disrupts the capacity for selective visual attention. It is important to notice that in this study the memory task that disrupted the control of selective visual attention drew on *verbal* rather than *visual* working memory: the task in question was to remember a series of digits, which though displayed visually were then stored in verbal memory by repetition in the articulatory loop. (On this point, see our clarification in fn. 20.) And other studies suggest precisely the opposite effect of *visual* memory load on the capture of attention by task-irrelevant visual distractors: for example, Konstantinou et al. (2014) found that performance in a visual classification task that required suppressing attention to an irrelevant distractor was impaired by a concurrent verbal memory task (remembering a series of symbols), but *facilitated* by a concurrent task that drew instead on visual memory (remembering a colored pattern). What seems to explain this difference is that while the concurrent verbal memory task interferes with the top-down control of attention, the concurrent visual memory task consumes the attentional resources that are needed for processing visual stimuli, as a consequence of which task-irrelevant stimuli are automatically not attended, and thus there is no need to actively suppress attention to them.<sup>21</sup>

<sup>21</sup> In making this suggestion, Konstantinou et al. (2014) rely on a central thesis from Lavie's perceptual load theory (see Lavie 2005), namely that attentional resources are allocated to task-relevant stimuli, and only if there are leftover resources, they are allocated to task-irrelevant stimuli.

The bottom line is that while there is a good deal of evidence that verbal working memory load interferes with cognitive control of distractor processing, there is *no* evidence suggesting that visual working memory load has a similar effect. On the contrary, studies testing the matter have found either facilitation of selective attention, as in the study by Konstantinou et al., or found no effect on selective attention at all (Yi, et al. 2004).<sup>22</sup> Moreover, the facilitation effect of visual memory on selective attention *has nothing to do with cognitive control of distractor processing*, but is instead a byproduct of the *visual maintenance* function of visual working memory, as explained in the previous paragraph. Apart from interference due to general dual-task cost, it appears that visual working memory load does not interfere with selective attention, or if it does, the effect is not manifested in extant dual-task experiments designed to test it. By contrast, verbal working memory load interferes significantly with selective attention (over and above a general dual-task cost), by drawing on the memory resources required to control attention in a top-down manner (i.e., resources responsible for remembering which stimuli are task-relevant).

This difference between the effects of verbal and visual memory load on selective visual attention offers a way to test the attentional account. On this account, perspectival identification is essentially a selective attention task, and so if the account is correct then we should expect to find that a concurrent *verbal* working memory task interferes more with perspectival identification than a concurrent *visual* working memory task (if indeed the latter interferes at all). That is, the attentional account predicts **that verbal working memory load will interfere with perspectival identification more than visual memory load**. Finding otherwise would be evidence against the attentional account.



Figure 10: Possible stimuli for our modified Stroop task. In (a) the size of each image is *congruent* with the size of the animal that it is an image of. In (b) the size of the

In the dual task situation in question, there are two tasks that require resources of visual attention, and consequently there are no leftover attentional resources that could be allocated to task-irrelevant stimuli.

<sup>22</sup> A study by Park et al. (2007) is an exception. The study found that visual memory load can impair selective attention, but it can do so *only when* the remembered item (face) matches the target category (face) and not the distractor category (house) in the selective attention task. We can safely avoid this complication, as long as we make sure that in our proposed experiments, it is not the case that the visual item to be remembered matches the category of the target, but not of the distractor, in the selective attention task.

image is *incongruent* with the size of the animal. Given the instruction to identify either the larger or smaller animal or the larger or smaller image, a measure of distractor suppression will be the extent to which congruent performance is superior to incongruent performance. Image credits: Wikimedia Commons.

This prediction of the attentional account could be tested in the following pair of experiments. First, the *control experiment* will compare the effects of verbal and visual memory, respectively, on a visual task that requires selective attention to a particular aspect of the display. This could be done using the Stroop task shown in Figure 10, where participants need to attend either to the size of the image or to the size of the animal that it depicts: here the expected result would be that the influence of the task-irrelevant stimulus dimension will be increased significantly by a concurrent verbal memory task (which interferes with selective attention), whereas a concurrent visual memory task will not have this effect. Second, the *crucial* experiment will pair verbal and visual memory tasks, respectively, with a perspectival identification task, as shown in Figure 6. If the attentional account is correct then the pattern of interference should resemble the one found in the control experiment. If a different pattern is observed, this would count against the attentional account.

The reason why the control experiment is so important is that there is conflicting data about whether spatial separation between target and distractor stimuli make a difference to the way that selective attention interacts with working memory load. For example, Gao, Chen, and Russell (2007) found that verbal WM load did not influence the magnitude of distractor interference in a standard Stroop task that required reading a stimulus word while ignoring its color. This suggests that the influence of verbal working memory load on distractor processing requires spatial separation between target and distractors. However, as mentioned before, de Fockert et al. (2001) found that verbal working memory load increases distractor interference in a Stroop-like task in which names are superimposed on faces. This suggests that the influence of verbal working memory load on distractor processing does *not* require spatial separation between target and distractor. Since perspectival and objective properties are not spatially separated (in the sense that when attention is directed to the location of one of them, it is apparently also directed to the location of the other), we are not entitled to assume from the outset that verbal memory load will interfere with perspectival identification even if the latter requires selective attention. The control experiment is needed in order to show that the attentional account yields this prediction.

## 5. Conclusion

This paper can be seen as a proof of concept. Is it possible to use experimental methods to help make progress in a complex phenomenological debate that has proven recalcitrant to the methods of philosophy? Our answer is: Yes, though this approach will have the same limitations and uncertainties that attend to any other experimental endeavor. First, our proposals here rest on empirical assumptions

about the roles of working memory in visualization and selective visual attention: while these assumptions are supported by experimental data they could also be overturned in the same way, or shown not to apply to the case at issue. And second, the proposals we have described here only offer a way to *falsify* the attentional and constructive accounts, and are not able to confirm that either account is the correct one. Nor, finally, have we shown that these accounts are the only way for perspectivalist and anti-perspectivalist philosophers, respectively, to develop their views. Some such limitations are always the price of taking an experimental approach to a vexing theoretical issue.<sup>23</sup>

## References

- Baddeley, A. D., & Hitch, G. (1974). Working memory. *Psychology of learning and motivation*, 8, 47-89.
- Bar, M., & Biederman, I. (1998). Subliminal visual priming. *Psychological Science*, 9(6), 464-468.
- Briscoe, R. E. (2008). Vision, Action, and Make-Perceive. *Mind & Language*, 23(4), 457-497.
- Brogaard, B. (2010). Strong Representationalism and Centered Contents. *Philosophical Studies*, 151, 373-392.
- Cohen, J. (2010). Perception and Computation. *Philosophical Issues*, 20, 96-124.
- de Fockert, J. W., Rees, G., Frith, C. D., & Lavie, N. (2001). The role of working memory in visual selective attention. *Science*, 291(5509), 1803-1806.
- Eitam, B., & Higgins, E. T. (2010). Motivation in mental accessibility: Relevance of a representation (ROAR) as a new framework. *Social and personality psychology compass*, 4(10), 951-967.
- Enns, J. T., & Oriet, C. (2007). Visual similarity in masking and priming: The critical role of task relevance. *Advances in Cognitive Psychology*, 3(1-2), 211.
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.
- Gao, Q., Chen, Z., & Russell, P. (2007). Working Memory Load and the Stroop Interference Effect. *New Zealand Journal of Psychology*, 36(3), 146.
- Gibson, J. J. (1950). *The Perception Of The Visual World*. Boston: Houghton Mifflin.

---

<sup>23</sup> The authors wish to thank Galia Avidan, Robert Briscoe, Zohar Bronfman, Randy Clarke, Paul Conway, Baruch Eitam, Jonathan Folstein, Nurit Gronau, Sean Kelly, Arnon Keren, Ariel Meirav, Lior Shmuelof, Mark Wagner, Aaron Wilbur, Wayne Wu, Yaffa Yeshurun, and two anonymous referees for helpful feedback and discussion.



- Gibson, J. J. (1986). *The Ecological Approach to Visual Perception*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27, 377-442 .
- Harman, G. (1990). The Intrinsic Quality of Experience. *Philosophical Perspectives*, 4, 31-52.
- Hellie, B. (2006). Beyond Phenomenal Naiveté. *Philosophers' Imprint*, 6(2). Retrieved from [www.philosophersimprint.org/006002/](http://www.philosophersimprint.org/006002/)
- Henik, A., & Tzelgov, J. (1982). Is three greater than five: The relation between physical and semantic size in comparison tasks. *Memory & cognition*, 10(4), 389-395.
- Hibbard, P. B., Haines, A. E., & Hornsey, R. L. (2017). Magnitude, precision, and realism of depth perception in stereoscopic vision. *Cognitive Research: Principles and Implications*, 2(1), 25.
- Hill, C. (2009). *Consciousness*. Cambridge, England: Cambridge University Press.
- Hopp, W. (2013). No such look: Problems with the dual content theory. *Phenomenology and the Cognitive Sciences*, 12(4), 813-833.
- Hyun, J. S., & Luck, S. J. (2007). Visual working memory as the substrate for mental rotation. *Psychonomic bulletin & review*, 14(1), 154-158.
- Kelly, S. D. (2008). Content and constancy: Phenomenology, psychology, and the content of perception. *Philosophy and Phenomenological Research*, 76(3), 682-690.
- Konstantinou, N., Beal, E., King, J. R., & Lavie, N. (2014). Working memory load and distraction: dissociable effects of visual maintenance and cognitive control. *Attention, Perception, & Psychophysics*, 76(7), 1985-1997.
- Lavie, N. (2005). Distracted and confused?: Selective attention under load. *Trends in cognitive sciences*, 9(2), 75-82.
- Lee, H., & Vecera, S. P. (2005). Visual cognition influences early vision: the role of visual short-term memory in amodal completion. *Psychological Science*, 16(10), 763-768.
- Lycan, W. G. (1996). *Consciousness and Experience*. Cambridge, MA: The MIT Press.
- Mack, A. (1978). Three modes of visual perception. In M. H. Pich (Ed.), *Modes of perceiving and information processing* (pp. 171-186). Hillsdale, NJ: Erlbaum
- Michal, A. L., Lleras, A., & Beck, D. M. (2014). Relative contributions of task-relevant and task-irrelevant dimensions in priming of pop-out. *Journal of vision*, 14(12), 14-14.

- Nanay, B. (2010). Perception and imagination: amodal perception as mental imagery. *Philosophical Studies*, 150(2), 239-254.
- Noë, A. (2004). *Action in Perception*. Cambridge, Mass.: MIT Press.
- Noë, A. (2008). Reply to Campbell, Martin, and Kelly. *Philosophy and Phenomenological Research*, 76(3), 691-706.
- Palmer, S. E. (1999). *Vision Science: Photons to Phenomenology*. Cambridge, Mass.: MIT Press.
- Park, S., Kim, M. S., & Chun, M. M. (2007). Concurrent working memory load can facilitate selective attention: evidence for specialized load. *Journal of Experimental Psychology: Human Perception and Performance*, 33(5), 1062.
- Pohl, C., Kiesel, A., Kunde, W., & Hoffmann, J. (2010). Early and late selection in unconscious information processing. *Journal of Experimental Psychology: Human Perception and Performance*, 36(2), 268.
- Robbins, T. W., Anderson, E. J., Barker, D. R., Bradley, A. C., Fearnlyhough, C., Henson, R., Hudson, S. R., & Baddeley, A. D. (1996). Working memory in chess. *Memory & Cognition*, 24(1), 83-93.
- Rock, I. (1983). *The Logic of Perception*. Cambridge, MA: The MIT Press.
- Schellenberg, S. (2008). The Situation-Dependency of Perception. *The Journal of Philosophy*, 105(2), 55-84.
- Schroer, R. (2008). The woman in the painting and the image in the penny: an investigation of phenomenological doubleness, seeing-in, and 'reversed seeing-in'. *Philosophical Studies*, 139(3), 329-41.
- Schroer, R. (2017). Hume's Table, Peacocke's Trees, the Tilted Penny and the Reversed Seeing-in Account. *Mind & Language* 32(2), 209-230.
- Schwitzgebel, E. (2011). *Perplexities of consciousness*. MIT press.
- Siewert, C. (2006). Is the appearance of shape protean?. *Psyche*, 12(3), 1-16.
- Sims, V. K., & Hegarty, M. (1997). Mental animation in the visuospatial sketchpad: Evidence from dual-task studies. *Memory & Cognition*, 25(3), 321-332.
- Smith, A.D. (2002). *The Problem of Perception*. Harvard University Press.
- Thouless, R.H. (1931a). Phenomenal regression to the real object. Part I. *The British Journal of Psychology* 21, 339-359.
- Thouless, R.H. (1931b). Phenomenal regression to the real object. Part II. *The British Journal of Psychology* 22, 1-30.

Tye, M. (2002). *Consciousness, Color, and Content*. Cambridge, MA: The MIT Press.

Vishwanath, D. (2014). Toward a new theory of stereopsis. *Psychological Review*, 121(2), 151-178.

Watt, S. J., Akeley, K., Ernst, M. O., & Banks, M. S. (2005). Focus cues affect perceived depth. *Journal of vision*, 5(10), 834-862.

Yi, D. J., Woodman, G. F., Widders, D., Marois, R., & Chun, M. M. (2004). Neural fate of ignored stimuli: dissociable effects of perceptual and working memory load. *Nature neuroscience*, 7(9), 992.

Weksler, A. (2016). Retinal images and object files: towards empirically evaluating philosophical accounts of visual perspective. *Review of Philosophy and Psychology*, 7(1), 91-103.

Weksler, A. (2017). Attention to mental paint and change detection. *Philosophical Studies*, 174(8), 1991-2007.