

Attention

Philosophical and Psychological Essays

EDITED BY

Christopher Mole, Declan Smithies,
and Wayne Wu

OXFORD
UNIVERSITY PRESS

9

Attention and Iconic Memory

Ian Phillips

1. THE SPAN OF APPREHENSION

The extent of our “span of our apprehension” is one of the oldest questions in philosophical psychology, with roots going back to antiquity.¹ As experimental psychology flourished toward the end of the 19th century, investigators sought to gain experimental purchase on the question. Thus, in 1886, Cattell conducted a pioneering investigation which he took to settle the matter (under the relevant experimental conditions). Cattell (1886) employed a “gravity chronometer” to present subjects with a series of letters for about 1 ms. After presentation, observers were able to recall at most three to five letters from such displays, irrespective of how many characters were presented, a finding later confirmed by Sperling (1960).

Unfortunately, Cattell’s experiment did not settle the matter.² His experiment speaks directly only to the following question: (1) How much of a briefly presented scene can subjects *recall* (or, *report*)? Yet Cattell explicitly conceives his work as an experiment “on the limits of consciousness” (1886, p. 310), not memory. Thus, his question is not (1), but (2) How much of a briefly presented scene are subjects consciously aware of?

Insofar as Cattell took himself to have determined the limits of consciousness, he must have implicitly presumed that, in this context, the answer to (1) settles the answer to (2). Cattell also writes that his question concerns “the number of impressions consciousness can at one time attend to” (p. 311). This introduces a third question, *viz.*, (3) How much of a briefly presented scene can subjects *attend to*? Once again, it is natural to read Cattell as implicitly presuming that the answer to (1) here settles the answer to (3).

The same implicit presumptions appear in other work from this period. Thus, in his textbook discussion, Whipple (1914) moves freely between talk of “attention,” “consciousness,” and cognitive “grasp.”

In a single “pulse” of attention only a small number of impressions can stand out clearly: the area or span of consciousness is definitely limited. In

the sphere of vision, we find that if we give but a single glance at any heterogeneous collection of objects, such as the goods displayed in a store-window, or the jumble of odds and ends in an old tool-chest, we are able to grasp and enumerate only a very few, perhaps four or five, of these objects. (p. 263)

Yet Whipple is not at sea conceptually. Discussing visual-apprehension tasks a few pages later, he notes how experiments “of this type have been variously designated as tests of ‘quick perception,’ of ‘observation,’ of ‘degree of attention,’ or even as ‘memory tests’” (pp. 278–79). To this, he quickly concedes that, strictly speaking, “it is impossible to draw any hard and fast lines between tests of apprehension, tests of memory and tests of fidelity of report” (p. 279). Does this concern Whipple? No. Evidently, he doubts the practical relevance of such distinctions.

Not everyone was so sanguine.³ Reviewing the literature in 1926, Gill and Dallenbach (1926) argue vociferously that “the term ‘range of attention’ has been mistakenly applied to [experiments to date]. These experiments give us merely the range of cognition” (p. 248). They then go on to argue, rather unpersuasively, that the true extent of attention ranges “from approximately 17–42 stimulus-objects” (p. 256).⁴

Fast forward a century, and a remarkably similar dispute over the relations among report, attention, and consciousness remains at the forefront of endeavors to understand our mental lives. Many philosophers and psychologists follow Cattell and Whipple in employing subjective report as their definitive criterion of consciousness. Naccache (2006) is explicit: “Consciousness is univocally probed in humans through the subject’s report of his or her own mental states. . . . The ability to report one’s own mental state is the fundamental property of consciousness” (p. 1396).⁵ Similarly, a related, though not entirely overlapping, group argues that there can be no consciousness without attention.⁶

However, a growing body of opinion takes the view that questions (1)–(3) are not only conceptually distinct but also experimentally separable. Thus, Lamme (2003) argues that “attentional selection is inherently independent of either awareness or memory, but determines whether we go from phenomenal to access awareness or from iconic to working memory” (p. 14). In other words, Lamme thinks that we are conscious of more than we attend to, and that attention then gates between awareness and short-term (working) memory on which report is based.⁷ As Lamme’s invocation of iconic memory attests, the growth of this latter body of opinion is in large part due to Sperling’s (1960) pioneering partial report paradigm. What Sperling’s paradigm shows, according to many, is that in a single glance subjects can *see* a great deal more than they can subsequently *report*. The difference is typically explained by a limit on short-term memory, gated by an attentional bottleneck.

In the next section, I set out how Sperling’s work is used to argue for the claim that conscious experience in this sense “overflows” attention, memory, and report. In section 3, I then propose an alternative *postdictive*

interpretation of Sperling's findings that fails to support this conclusion. Attention is often taken to play a fundamental role in Sperling's task. In section 4, I argue that this poses a serious problem for standard accounts, since they fail to give an adequate account of the *objects* of attention in partial reporting. In section 5, I show how a story about attention can be provided given a postdictive interpretation of Sperling. Consequently, we have positive reason to prefer a postdictive interpretation over its orthodox rivals. If this is right, Sperling's work fails to establish that conscious awareness overflows attention or cognitive access.

2. SPERLING'S PARADIGM

Presented with a stimulus display, as shown in figure 9.1, for 15–500 ms, and asked to report as many of the letters as possible, subjects consistently report an average of 4.3 letters in their correct positions. However, if *partial* reports are elicited from subjects by playing at random a high, medium, or low cue-tone (indicating which row is to be reported) immediately following grid offset, then trained subjects consistently report an average of 3.04 letters in the cued row (Sperling, 1960; Averbach & Sperling, 1961). This finding—that subjects are, on average, able to report more letters with respect to a row that has been cued immediately after display offset than with respect to a randomly chosen row in an uncued trial—is known as *partial report (PR) superiority*. Sperling found that such PR “superiority decreased until, at cue delays greater than about 300 ms, there was no superiority” (Coltheart, 1980, p. 185).⁸

How should we explain PR superiority? Sperling's idea was that there must be a kind of “sensory memory store”—a construct Neisser (1967, p. 20) termed “iconic memory”—encoding information about at least nine of the letters. Sperling proposed that this store endures for around 300 ms—enough time for the cue to determine selectively which letters are transferred to a smaller capacity but more durable form of memory underlying reports. I do not want to question the

U	I	V	F
X	L	G	Q
B	S	W	K

Figure 9.1 Typical stimulus display used in Sperling's partial report task (Sperling, 1960; Averbach and Sperling, 1961).

correctness of this claim with respect to subpersonal/nonconscious *informational persistence*. The persistence of some form of information regarding nine-plus letters is essential to explain PR superiority (see the seminal discussion in Coltheart, 1980). What remains to be established is that this information is *conscious*.

For his part, Sperling took his results to show that subjects do consciously perceive at least nine letters, even if they can remember only around four. His opening question is whether, “*more is seen than can be remembered*” (1960, p. 1) and, in the light of his data, he gives an affirmative answer. Likewise, Neisser (1967, pp. 19–20) understands iconic memory in terms of a certain kind of visible persistence, in part defined “introspectively.” Sperling’s view has been hugely influential among philosophers of mind who hold “that the Sperling experiment directly shows the existence of phenomenal states that are not cognitively accessible” (Block, 2007, p. 489), or, as Block also likes to put it, that conscious experience *overflows* cognitive access.⁹

Two assumptions are needed to argue from PR superiority to the overflow of conscious experience. The first assumption is that if a subject correctly reports three specific letters in a given row, then that is strong presumptive evidence that the subject enjoyed conscious experience of those three letters as the specific letters reported. The second assumption is that any aspect of experience present in a PR condition *would* have been present even if some other PR had been cued. This assumption lies behind Sperling’s remark that “The most reasonable procedure is to treat the partial report as a random sample of the letters which the S has available” (1960, p. 7).

Putting the two assumptions together, it is easy to see why Sperling’s paradigm is thought to have striking consequences concerning experiential richness. By the first assumption, the letters reported in a given PR condition (say, three top-row letters reported after the playing of a high tone) are consciously experienced. By the second assumption, these letters would also have been experienced if a different tone had been played. But if a medium tone had been played, different letters from the middle row would have been reported. These letters would also have to be regarded as aspects of the experience by our first assumption. Repeating the argument with the bottom row and a low tone, we are obliged to conclude that at least nine letters are experienced, despite our capacity for report in any single case being limited to half that many. In other words, the two assumptions legitimate summing PRs to establish the true extent of awareness.

Note that the argument from Sperling’s findings to overflow goes beyond a simple appeal to introspection. Sperling’s work was explicitly intended to afford experimental purchase on the fact that in such tasks subjects “enigmatically insist that they have seen more than they can . . . report afterwards” (1960, p. 1). It would be to abandon this aspiration to fall back on subjects’ judgments (*pace* Block, 2007; Tye, 2006).

Moreover, such a fallback would have to rule out a number of alternative explanations of subjects' insistences (see my, forthcoming, for further discussion; also Byrne, Hilbert, & Siegel, 2007; and Papineau, 2007).

3. AN ALTERNATIVE INTERPRETATION

I have no wish to challenge the first assumption above. In Sperling's experiments, we have no reason to doubt that subjects' reports are grounded in conscious experience. Even if the absence of subjective report does not establish the absence of conscious experience, it would be a radical move, indeed, to claim that the kind of subjective reports in question, typically treated as paradigmatic criteria for awareness, do not evidence awareness here. In contrast, we do possess independent reason for rejecting the second assumption. The reason it is universally assumed that any aspect of experience present in a PR condition would have been present even if some other PR had been cued is that PRs are elicited by a tone played just *after* stimulus offset (e.g., Tye, 2006, p. 511; Dretske, 2006, p. 17). The fact that the cue comes after display offset is supposed to show that our experience of the display is *cue independent*. Yet this form of reasoning is widely rejected in interpreting a large number of experimental paradigms where the perception of an initial (target) stimulus is affected—sometimes dramatically—by a second (modulator) stimulus presented a short time after target offset. These findings evidence perceptual *postdiction*. As Choi and Scholl (2006) distill it, the basic idea is that “our conscious perception of the world is not an instantaneous moment-by-moment construction, but rather is formed by integrating information presented within short temporal windows, so that new information which is obtained can influence the immediate past in our conscious awareness” (p. 385).

I discuss a number of standard examples in my “Perception and Iconic Memory” (Phillips, forthcoming): backwards masking (Alpern, 1953; Enns & Di Lollo, 2000; Weisstein & Wong, 1986); sound-induced visual bounce (Dufour, Touzalin, Moessinger, Brochard, & Desprès, 2008; Sekuler, Sekuler, & Lau, 1997; Watanabe & Shimojo, 2001); apparent motion (e.g., Eagleman & Sejnowski, 2003; Steinmann, Pizlo, & Pizlo, 2000; Wertheimer, 1912); and the flash-lag illusion (Eagleman & Sejnowski, 2000; Mackay, 1958). Here, let me just mention one striking and rather different case.

If a moderately strong, 1000 Hz sine-wave frequency is to be heard as a tone, the burst of stimulation must be at least 8–10 ms long in normal subjects. Below that limit, the stimulus is heard as a click (Doughty & Garner, 1947; Stevens & Davis, 1938, pp. 100–102). However, in studies on patients with vertebral-basilar artery insufficiency (i.e., restricted blood flow to auditory nerve centers), Creel, Boomslinger, and Powers

(1970) discovered that this limit is increased to the order of hundreds of milliseconds. This raises the question of how subjects experience the first few hundred milliseconds of such a sine-wave (i.e., that period that on its own would be heard merely as a buzz). The answer is rather remarkable:

No patient has ever reported the signal to change during its time span. A given patient will hear a 400-msec. burst as noise. He will hear a 500-msec. burst as tone. One might suppose that he would experience the first 400 msec. of the stimulus as a buzz and then shift to tonal sensation for the final 100 msec., that his auditory neural processes would report noise until the point at which the nervous system cracked the code. Not so. He experiences the decided sensation right back to the start of the signal. (1970, pp. 537–38)

In other words, in such a subject, whether a stimulus lasting, say, 300 ms is heard as tone or noise depends essentially on whether the stimulus continues on or ceases.

There are many such paradigms, and the precise effects and time-courses in each are complex. However, in some, it is clear that visual features can be modulated even by modulating stimuli presented (a) several hundred milliseconds after initial target offset, and (b) in a different (e.g., auditory) sensory modality. Such cases immediately raise the prospect of a postdictive interpretation of Sperling's data. According to such an interpretation, Sperling's auditory cue directly modulates subjects' experiences of the letter display, despite its arrival up to 300 ms after display offset. On such a postdictive interpretation, experience is not cue independent. Just the opposite. Consequently, it cannot be assumed that experienced letters would have been experienced had a different cue been sounded. PRs cannot be summed and the overflow argument collapses.

In outlining this postdictive account of Sperling's task, I have not mentioned attention. Yet, Sperling and his commentators treat it as obvious that attention plays a critical role in the task. Thus, it might well be objected that the kind of postdictive modulation of experience appealed to above can only account for the data at the cost of denying a proper role for attention. In other words, it might be objected that recognizing Sperling's task as an attentional task shows why it cannot be given a postdictive interpretation. This challenge is a serious one. But behind it lies a deeper concern. For, as I now argue, *no* orthodox account of PR superiority provides an adequate role for attention. Moreover, this failure is irremediable. In contrast, it is possible to understand how attention operates given a postdictive account, even if its operation is not quite as we might first have thought (see section 5). This constitutes a clear reason to favor a postdictive interpretation.

4. THE PROBLEM OF ATTENTION IN SPERLING'S PARADIGM

There is a trivial sense in which attention is implicated in partial reporting. In reporting that, say, the top row contains a *U*, an *I*, and an *F*, subjects

judge and so *think* that the top row contains these letters. Yet, as Martin (1997) puts it, “whatever we are prepared to call an object of thought—be it the things thought about, what one thinks about them, or the proposition one thinks in thinking these things—we can also take to be an object of attention” (p. 77).

There is also a *technical* sense in which attention is implicated in partial reporting. For example, according to Coltheart (1983, pp. 53–54), partial reporting may be best understood in terms of Kahnemanian object files. The basic idea is that on seeing a grid such as in figure 9.1, a dozen object files are quickly created, one for each grid location. Cueing is then understood in relation to these files. As Coltheart and Coltheart (2010) summarize the idea:

The subject’s task is to select only those files that match the cue and to transfer information about each of these files . . . to durable storage. . . . What is the role of visual attention here? . . . It might be that checking each object file to determine whether it contains the cued property requires visual attention to be directed to that file. . . . On the other hand, it might be that, after an object file has been selected as matching the cue. . . , visual attention needs to be directed only to each selected file. (p.355)

Though prevalent within the psychological literature, this notion of attention is not conscious visual attention as we ordinarily understand it. Though connected to our ordinary notion, the concept is fundamentally a technical one, to be understood in relation to an information-processing account of visual cognition (cf. Martin, 1997). As such, it does not pose the problem this chapter is primarily concerned with (though see section 5, where I appeal to aspects of subpersonal processing).

Our problem is with the role that *conscious visual* attention plays in Sperling’s task. It is this kind of visual attention that is being invoked in the following remarks:

The subjective response to the high signal tone is “looking up.” Since eye movements cannot occur in time to change the retinal image with any of the presentations used (Diefendorf & Dodge, 1908) a successful looking-up must be described in terms of a shift in “attention.” . . . The reaction time for the attentional response, like the reaction time for more observable responses, is greater than zero. Therefore, if the S is given an instruction before the presentation, he can prepare for, or sensitize himself to, the correct row of the stimulus even though there is not time enough for a useful eye movement. The response to an instruction which is given 0.05 sec. before the stimulus is probably the same as the response to a similar instruction that is given 0.1 sec. later, immediately after the exposure. The short time difference, 0.1 sec., accounts for the similar accuracy of responding in these two conditions. Once his attention is directed to the appropriate row, the S still has to read the letters. This, too, takes time. (Sperling 1960, p. 24)

In standard cueing paradigms, a *prior* cue directs our visual attention to some region (or object), and then (assuming that the target is presented

quickly enough to avoid inhibition of return) our capacity to report a subsequently presented target in the region (or change to the object) is enhanced in terms of speed or detail of response. Sperling suggests viewing his results along similar lines, as a matter of a *posterior* cue directing our attention to a given row which we can then read and report letters from.¹⁰ This immediately raises a puzzle, however. Given that the cue directs attention only after the grid has been offset, how can it be *the grid* to which the cue directs visual attention?

Illusion Accounts

Tye (2006) defends the natural interpretation of Sperling's view when he proposes that "In [Sperling's] experiment, the tone has the effect of focusing the subjects' attention on one particular part of the array that apparently is still before them" (p. 517). Call this view the *illusion* account. According to this account, "after the array has been extinguished, it *appears* still to be displayed" (p. 511). Because it still appears to be displayed, subjects can selectively attend to it.

Dretske (2006) argues that this illusion view cannot be right. Instead he recommends what I call an *imagistic* account:

[Partial reporting in Sperling's paradigm] is not a case of attention being drawn to the queried row since the signal for which row to report occurs after removal of the stimulus. There is no longer anything out there (where the stimulus was) for their attention to be drawn to. Rather, subjects extract this information from what they describe as a conscious but rapidly fading image ("icon") that persists for a short time after removal of the stimulus. They use the information embodied in this conscious experience to identify letters in a stimulus that is no longer physically present. (p. 175)

Dretske here denies that subjects attend to anything apparently *in their environment* when the cue occurs. Rather, he suggests, they have available a visual image which they attend to in reporting the relevant row's letters.

On the face of it, Dretske's objection appears to rehearse a slip he made forty years ago in a different context:

Being *directly aware* of something is a state of affairs which implies that the element of which one is directly aware *must* [on the next page Dretske adds: "logically must"] exist at the time one is directly aware of it. Since science has shown us that coffee pots need not exist at the time when, as we ordinarily say, we see them, we must conclude that we are never directly aware of coffee pots—nor anything else the perception of which involves a causal sequence involving a temporal interval. (1969, p. 72)

Pitcher (1971) gives the now standard reply to this traditional time-lag argument:

the finite speed of light does not entail that we do not directly see things and states of affairs in the "external world," but only that we must see them

as they were some time ago. We see real physical things, properties, and events, all right, but we see them late, that is all. . . . It is a mere prejudice of common sense—and one on which the time-lag argument trades—that the events, and the states of objects, that we see must be simultaneous with our (act of) seeing them. (p. 48; cf. Ayer, 1956, pp. 93–95)

Illusion theorists can say much the same in response to Dretske's objection concerning visual attention. The fact that attention is deployed after the display has been offset does not show that we do not attend to the display itself (as opposed to an image), but merely that we attend to the display *as it was some time ago*. There is no more reason to think that, because the display is offset, we must attend to an image of the display than there is reason to think that, because light takes eight minutes to travel to sun, we cannot attend to the sun but only a solar image.

When we think about the role illusion accounts assign to attention, however, it is clear that something has gone wrong. For, although it is acceptable to think that a cue could direct attention toward a row in an *environmentally absent* display, it is not acceptable for a cue to direct *perceptual* attention toward a row in a *perceptually absent* display. Tye (2006) claims that the display "*appears* still to be displayed" (p. 511). But this is surely false. With effective delays above, say, 100 ms, subjects clearly hear the cue being sounded *after* they have seen the display vanish. Moreover, at least as we naïvely think of it (though see below), the cue must enter one's conscious experience before one can direct one's attention to the cued row.

Tye might appeal to the phenomenon of visible persistence. To take an example from Coltheart (2009), imagine looking out through a window on a dark and stormy night. Suddenly a flash of lightning illuminates the landscape. The bolt lasts just five-hundredths of a millisecond, yet the experience of the landscape lasts much longer than that. Here, one may well be able to attend to a tree significantly after its 50 μ s illumination has ended. Tye might think that this is just what is occurring with iconic memory.¹¹ However, PR superiority cannot be accounted for in terms of visible persistence. As Coltheart details, there are "two fundamental properties of visible persistence" (1980, p. 183), namely:

- (a) The *inverse duration effect*—the finding that increasing the duration of a stimulus decreases its persistence beyond stimulus offset (visible persistence is "time-locked" to stimulus onset).
- (b) The *inverse intensity effect*—the finding that increasing the intensity of a stimulus decreases its persistence beyond stimulus offset.¹²

Neither of these properties holds for informational persistence as defined by Sperling's PR methods. The duration of iconic memory is not inversely related to stimulus duration nor to intensity. As Coltheart puts it, "informational persistence or iconic memory cannot be identified with visible persistence, since they have fundamentally different properties" (1980, p. 183).

Tye's illusion account thus fails (a) because it does not capture how things seem from the subject's perspective, and (b) because, phenomenology aside, it implicitly treats iconic memory as a form of visible persistence, despite the fact that visible persistence lacks the properties of the construct implicated in Sperling's task.


Imagistic Accounts

Dretske (and many others) reject the illusion account and posit instead a nonperceptual image, or "icon," to serve as the object of attention post-cue. The idea is that after the grid has been offset, it disappears from perceptual consciousness but remains available in the form of a visual memory image. Subjects have access to the grid in visual memory, and they can selectively attend to aspects of their memory image upon cueing. This account need not think of subjects literally attending to images (as if to mental photographs). However, they must do more than claim that subjects can remember *that* the grid contained certain letters (on pain of assigning no nontrivial role to attention). Minimally, the imagistic account requires subjects to visually (i.e., episodically) recall the grid. Arguably, subjects do this by drawing on the same general capacity as they have for visual imagination, the only difference being the particularity of content in case of memory.¹³ It is in this sense that such accounts count as imagistic.


Imagistic accounts seem initially more plausible than perceptual accounts since they allow that the grid does not still *appear* to be displayed. However, they face their own difficulties. Suspicions are first raised by the fact that subjects in Sperling's task are not instructed to make, and do not consistently report making, use of visual imagery. This contrasts with supposedly similar paradigms also appealed to in overflow arguments—for example, Landman, Spekrijse, and Lamme (2003) and Sligte, Scholte, and Lamme (2008), where imagery does seem implicated.¹⁴ Suspicions are strengthened when we appreciate that images take time to form. This is nicely brought out by a series of experiments conducted by Brockmole, Wang, and Irwin (2002).

Brockmole et al.'s first experiment takes the form of a standard temporal integration task in which a stimulus grid is presented briefly (here for 33 ms), followed by a blank, followed by a second stimulus grid in the same location for the same duration (see figure 9.2). The integration task is to determine which grid position does not contain a dot on either trial. The task is trivial if the grids are presented simultaneously, for the grid will appear as a single grid (figure 9.3). The task is also easy when a very small temporal gap is introduced (Di Lollo, 1980; Loftus & Irwin, 1998). At such delays, the displays continue to be "phenomenally integrated" (Coltheart, 2009, p. 429). The time gap cannot be large, however; by 100 ms, Brockmole et al. (2002) found that performance fell from 79% to 21% accuracy.

What Brockmole et al. (2002) discovered, however, was that accuracy does not stay at this low level as longer delays are introduced. Rather,



First Stimulus Grid



Second Stimulus Grid

Figure 9.2 Examples of first and second stimulus grids as used in Brockmole, Wang, and Irwin, 2002.




Figure 9.3 Representation of 'phenomenal integration' of grids in Figure 9.2.

accuracy begins to increase in positive (curvilinear) relation to the inter-stimulus interval, until at a delay of 1500 ms, "Accuracy asymptoted at approximately 68%" (p. 318). Performance then plateaus through delays of 5000 ms.

Brockmole's subjects were encouraged to form a visual image of the first grid and "overlay" the image on the second grid when it was presented. Subjects were clearly capable of doing this with remarkable efficacy. However, what is also clear is that "the generation of the visual image is clearly a very slow and effortful process taking as long as 1500 ms to complete" (Coltheart, 2009, p. 429; cf. Kosslyn, Thompson, & Ganis, 2006). The crucial point for our purposes is that grid images take time to form (i.e., episodic recall takes time). Consequently, if memory imagery accounted for performance in Sperling's task, performance would exhibit an entirely different time-course from that in fact found. We would expect PR superiority to be negligible at, say, 150 ms, but then increase to near ceiling at 1500 ms. This is simply not what occurs.¹⁵ As a result, imagistic accounts are not viable.¹⁶

Grush's Wave-Collapse Illusion

Attention is the undoing of both illusion and imagistic accounts, for they lack a plausible answer to the question: What are the objects of attention? Illusion accounts propose that we attend to a grid that has already slipped from consciousness. Imagistic accounts more plausibly think of our attention as being directed toward a visually recalled grid. However, phenomenological doubts about this idea aside, the proposal appears ruled out when set against what else we know about the time-course of image construction and episodic recall.

In general, accounts of Sperling's paradigm neglect attention at their peril. Let me illustrate by briefly mentioning an account inspired by remarks in Grush (2007). Grush himself agrees with Block that regarding "the phenomena in the Sperling . . . experiments . . . the better analysis is that the items subjects think are phenomenally present are in fact present, but inaccessible in a certain sort of way" (p. 504). Nonetheless, Grush goes on to offer an interpretation of change and inattention blindness that one might think could be applied to Sperling.¹⁷ Grush's idea has three parts: (1) that much of our phenomenology is *generic* as opposed to specific (i.e., detailed and determinate); (2) that we normally think of our experience as highly detailed; and (3) that this is an "illusion," explained by the fact that we do not notice the transition between generic and specific contents when we shift our attention or gaze. Grush thinks of his "wave-collapse illusion" as "a less radical cousin of the famous refrigerator light illusion" (2007, p. 504).

I am very happy to agree that much of our phenomenology is generic.¹⁸ However, it is unclear why we are said, in general, not to notice transitions from generic to specific phenomenology. If we really did have the impression that our experiences were richly detailed throughout, why would we bother shifting our attention (or gaze)? Rather, it is natural to think that we attend when we want to get a clearer view of things. And we are rightly unsurprised when attending yields a more determinate experience of some aspect of the scene: that is what we intended to bring about.

This concern aside, can Grush's suggestion form the basis of an account of Sperling's task? To do so, the account must say what it is that we *attend* to such that our experience of it switches from having generic to having specific phenomenology. Generic-specific shifts are supposed to occur when we shift our gaze, thus the obvious answer is that our perceptual experience of the *grid* shifts as we shift our attention. But, as already much emphasized, we are no longer in visual contact with the grid when the attention-directing cue arrives, so this cannot be what is going on in Sperling's task.

One might try and rehabilitate Grush's idea in terms a failure to appreciate that aspects of our "memory" images are more specific than our generic experiences. No doubt there are some cases where we accurately visually imagine a past scene in detail, despite not having had detailed *conscious* experience of that scene. However, in general, it does not seem attractive to reject the testimony of subjects "that their responses are based on specific phenomenology that was there all along" (Block, 2007, p. 531; Block, 2008, p. 307; Burge, 2007, p. 501). In short, a wave-collapse account of Sperling's task equally struggles with saying what the objects of attention are.

If no standard account of Sperling's PR task can assign a plausible role to attention, the prospects of a postdictive interpretation are bolstered. That is, of course, if we can assign a plausible role to attention on such an interpretation. I now turn to that task.

5. POSTDICTION AND ATTENTION

There are various forms of postdictive account available: “Stalinesque” accounts face their own difficulties in providing a role for attention in Sperling’s task; “extensionalist” models fare better.

Stalinesque Accounts

“Stalinesque” accounts¹⁹ of postdiction posit a substantial (300 ms plus) delay between initial grid stimulation and conscious experience of the grid, allowing for the processing of intervening stimuli. The simplest version of such an account can be sketched as shown in figure 9.4. In the diagram, the bottom line corresponds to the time-course of external stimulation, the top line to the time-course of experiential acts.

On this simple picture, *all* experience is delayed by some fixed amount. This delay affords enough time for the cue tone to occur before we experience the grid (and so for it to be taken into account “in the editing room”). The difficulty this account faces is that our attentional reaction to the cue seems to be a reaction to our conscious experience of that cue. Yet, once the cue tone is in the stream of consciousness, so too are the earlier parts of the stream of consciousness, including our grid experience. Thus, the picture provides no room for a reaction to our *conscious* experience of the tone to affect our experience of the grid.²⁰

One strategy at this stage is to claim that our attentional reaction to the cue is an automatic response driven by unconscious processing of the cue. Thus, we really do “attend” to the grid (as it was a few hundred milliseconds ago), but such attention is not a response to consciously experiencing the tone but, rather, an automatic reaction to nonconscious auditory stimulation.²¹ This reaction leads to an alteration to the way our attentional system is deployed and so to which row of the grid we




Figure 9.4 Representation of our experience in Sperling’s paradigm according to a simple Stalinesque account.

discriminate specific letter identities from. This strategy is coherent. However, to square itself with the subject's perspective, such an account must claim that this reaction *seems* to us in hindsight to be prompted by conscious experience (cf. Dennett, 1991, p. 122). Thus, the account takes us to be subject to a striking illusion regarding attentional control. As such, it is not clearly preferable to orthodox illusion and imagistic accounts above.

A Stalinesque account might hope to avoid this difficulty regarding attentional control by radically pulling apart the structure of the stream of consciousness from the structure represented in that stream, as shown in figure 9.5. Here, the bottom line represents both the order of external stimulation and the order in which the events are *represented* as occurring.²² As before, the top line represents the time-course of experiential acts (i.e., representings). On this picture we do not experience the grid until well after the cue tone has been experienced. This allows plenty of time for us to react consciously to the cue tone and redirect our attention.

However, this approach does not avoid the difficulty. According to it, we now attend *before* we have seen the grid: the cue is, in effect, a standard prior cue. It is just that, on this picture, the grid's appearance is referred back to a time before the cue sounds. Is our act of attention also referred back in time (to the same time as the grid display)? If so, it will seem to us as if we are attending pre-cue. But since our shift of attention is apparently a response to the cue, this makes little sense. Thus, it must be that we seem to attend post-cue. But then it cannot be that our attention is attention to a grid, given that the grid is perceived as occurring well before the cue tone. Once again, then, we face the question: What is it that we attend to? To make progress we need to reconsider the way in which attention might have its effect on postdictive approaches.

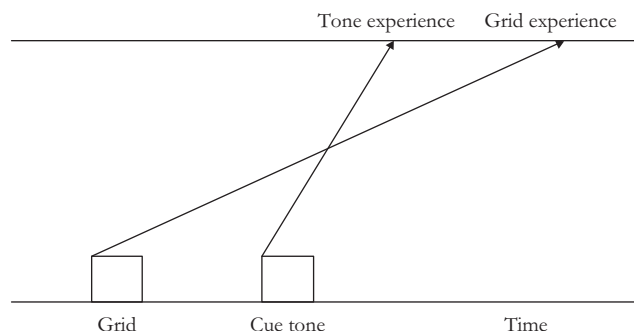


Figure 9.5 Representation of our experience in Sperling's paradigm according to a revised Stalinesque account.

Rethinking the Role of Attention

Psychologists and philosophers have long been interested in the role of attention in modulating conscious experience. Contemporary researchers talk of the role of attention in terms of the facilitation or amplification of internal representations of goal-relevant information and, conversely, in terms of the inhibition of internal representations of goal-irrelevant information. Evidently these effects are not effects at the personal level. Rather, they are the potential results of attending to a location (or object or feature), which is something that occurs at the personal level. Of course, these subpersonal effects may in turn have consequences for our conscious experience.

In the Sperling task, the goal is only determinately specified when the cue arrives after grid offset. Nonetheless, processing of information concerning the grid evidently occurs almost as soon as light from the display reaches the retina. Note two facts about this processing. First, it takes time. Second, and only marginally more controversially, a great deal more information is processed than ever reaches conscious awareness (e.g., Thornton & Fernandez-Duque, 2002). Suppose, then, that our attention is drawn or deployed to some location from which information has already been received via the eyes. This will typically be the case with all the grid positions in Sperling's task, given the way they are presented. It takes time for attention to be shifted (whether voluntarily or otherwise). Thus, if information concerning the relevant location has already been received and is being processed, though without necessarily being "headed" for consciousness, it is plausible to suggest (from an evolutionary design perspective) that an optimal attentional system would seek to capitalize on this information already in the system. After all, in general, attention is drawn or deployed because some location is of potential interest to us. Through attending, we enhance our perceptual acquaintance with the relevant location.²³ It would be wasteful to wait for attention to deploy itself and to receive fresh information when information regarding the relevant location is already being processed. Thus, a natural thought is that, if we direct attention to some currently unattended region, information relating to that region already in the system may be amplified. A consequence of attending may be a boost in the processing not just of *new* input but also of input *already* working its way through visual processing.

If this is right, an attentional effect regarding some stimulus need not arise from attending to that stimulus. We simply need to attend to the place where that stimulus *was*. So perhaps it is a mistake to think that we attend to either display or image. Perhaps, when the cue arrives, we simply *attend to where the relevant row of the grid was*. That is, on hearing a high tone, we attend and "look up" to where the letters were presented. It is very natural for us to do this. What is puzzling is why it should do us any good. After all, there is nothing to be seen there, only a blank screen. The above story provides the answer. Although we do not visually attend

to the letters themselves, attending to where they were (even in their absence) can have an effect on the processing of information relating to that row so long as that processing is ongoing.²⁴ Since processing takes time, this yields a window of opportunity for attention-driven processing consequences.

We can now return to our postdictive account. The proposal in front of us is that in Sperling's task, subjects attend to the place where the indicated row *was* presented, as opposed to the grid itself. Being cued to attend in this way is, it is suggested, sufficient to alter the way that subjects perceive the grid. Subjects do not attend to the grid or any image; rather, a shift in attention has processing consequences that alter the grid percept. We might attempt to understand this in Stalinesque terms. However, any Stalinesque account will retain two problematic features. First, they will require a large delay between grid display and grid experience. Second, they will require a radical divergence between the temporal structure of experiential presentation and the temporal structure of what is presented in experience. As I have argued elsewhere (Phillips, 2009, 2010) it is far from clear that this can be sustained when we reflect on our experience. Instead, I suggest that the proposal before us is better understood in terms of an *extensionalist* model of postdiction that has need for neither such a delay nor such a divergence.

Extensionalist Accounts

On an extensionalist metaphysics, our consciousness essentially "extends a short distance through time" (Dainton, 2008, p. 631; Dainton, 2000; Phillips, 2009). The idea is not merely the uncontroversial thought that experience is a temporally extended phenomenon but, rather, the claim that there are certain durations of experience that are *metaphysically prior* to their subtemporal parts. This is what the extensionalist means when denying that "our consciousness is confined to an instant" (Dainton, 2008, p. 626).²⁵ What does it mean to say that a duration of experience is metaphysically prior to its subparts? It is not to deny that there are facts about instants during our stream of consciousness. It is, however, to insist that such facts are derivative. The most basic facts about our experiential lives are facts about extended periods of the stream of consciousness. What is true at an instant is true only in virtue of that instant being an instant during a certain period of experience.

The extensionalist account of Sperling's task denies that prior stimulus experience is independent of the pitch of the cue tone. But this is not explained in terms of a delay. Rather, what is proposed is that stretches of experience are metaphysically basic. Thus, the nature of our experience at any given moment must be considered in relation to the nature of our experience over an extended period, including experience over, say, the next 300 ms. As a result, the answer to the question, "What is seen at the time when the display is being presented?" is not determined independently

of what is presented over surrounding periods. We may see the grid very shortly after it is presented (i.e., without a Stalinesque delay), but the *way* we see the grid may depend essentially on facts about a stretch of experience, including the subsequent cue tone. According to the extensionalist, the basic events that we perceive are extended in time. The relevant events to consider in Sperling's task are of the form "grid followed by tone." A different pitch of tone means a different event, potentially with a different *visual* (as well as auditory) appearance.

Consider figure 9.6. According to the extensionalist, the circled stretch is a stretch of experience the nature of any part of which is dependent on what else occurs within the stretch. Since experience of the tone falls within the stretch, the nature of our grid experience may be different as a result. Nothing in this picture demands that we attend to the grid, nor to any image of it. Instead, the suggestion above is taken up: the tone leads us to attend to where the relevant row *was*.

At a processing level, we can continue to think of attention as boosting the processing of information relating to a given row of the grid. But at a metaphysical level, this processing underpins a stream of consciousness in which *stretches* of experience are the basic units. Thus, in the first instance, the processing subserves an experience of the event "grid followed by tone". The way we experience the grid, being a subpart of this overarching experience, is metaphysically hostage to the overall experience that is still unfolding at the time the tone sounds. Processing that occurs due to the tone can evidently affect this *overall* experience. Thus, it can affect the way we experience the grid. Yet, as can be seen in the figure, this does not demand a large delay between stimulation and experience, nor does it require us to pull apart the structure of the stream of consciousness from the structure of its objects.




Figure 9.6 Representation of our experience in Sperling's paradigm according to an extensionalist account.

6. CONCLUSION

In the last section, I made the following claims. First, that the objects of visual attention in PR paradigms need be nothing more than the *locations* where a given row *was* presented. Visual attention need not be to the row itself, nor to an image of the row. Second, that attention of this kind can nonetheless have significant implications for grid-related processing. Third, that we can provide a model on which such attention can affect our conscious encounter with the grid. Two were outlined: a Stalinesque account, and an extensionalist account. I indicated that we should prefer an extensionalist model on the grounds that (1) it avoided substantial delays between stimulation and experience; and (2) it avoided pulling apart the temporal structure of the stream of consciousness from the structure of its objects.

Where does this leave us with respect to our starting question concerning the span of apprehension? If postdictive accounts, in particular extensionalist accounts, are not only viable but also preferable to orthodox interpretations of Sperling, then, as discussed above, the argument for overflow collapses: cue independence cannot be assumed. We are left, then, much where we stood at the turn of the last century—at a loss as to how to differentiate experimentally between questions relating to very short term recall, attention, and consciousness. How worried we should be by this state of affairs is another question.

Notes

This paper inherits the debts recorded in my “Perception and Iconic Memory” (Phillips, forthcoming). On this occasion, I am especially grateful to detailed comments from all three editors and, as ever, to Hanna Pickard.

1. Hatfield, 1995, discusses Aristotle’s views. Gill and Dallenbach, 1926, survey early experimental work. See also Fernberger, 1921, and the excellent literature review in Wilken, 2001, to which I am indebted. Much of the literature in this area focuses on immediate numerosity judgments (what Kaufman, Lord, Reese, & Volkmann, 1949, later termed “subitizing”), as opposed to the reporting of specific letters in Cattell and Sperling. The *locus classicus* for numerosity judgments is Jevons’s 1871 investigation inspired by Hamilton, 1859, 252f.

2. There is a further question as to how exactly a capacity limit should be understood; see Wilken, 2001; and Wilken and Ma, 2004.

3. Cf. Oberly, 1924.

4. Gill and Dallenbach’s use of the term “attention” resides within a theoretical framework long discarded. Roughly speaking, they use “attention” to refer to one of two “levels of clearness” within conscious experience. On attention as clearness, see Titchener, 1910. For acute criticism of Titchener connected to the issues at hand, see Woodworth, 1909. In general, discussion of attention and consciousness in this period is unsatisfactory owing to deep disagreement over the nature of attention.

5. See also, for example, Dehaene, Naccache, Sackur, and Sergent, 2006; Dennett, 1991; Papineau, 2002, esp. pp. 182–84; and Weiskrantz, 1997, p. 84.

6. This line is not new. It is found in William Hamilton, 1859, p. 247. The contemporary *locus classicus* is Mack and Rock, 1998. See Mole, 2008, for discussion.

7. The *locus classicus* is Block, 1995; see also Block, 2007; Burge, 2007; Dretske, 2006; Koch and Tsuchiya, 2007; Iwasaki, 1993; Tye, 2006; and Wolfe, 1999.

8. Those unacquainted with the paradigm may wish to try a demonstration such as the one currently available at: <http://www.ulb.ac.be/psycho/fr/docs/museum_en/Experiments/Sperling-exp.html>.

9. This quotation in the text is Block's description of his (1995) claim that Sperling's experiments are a case of P-consciousness without A-consciousness (p. 198). Block, 2007, 2008, offers much the same story. However, now, as he puts it, "the fact of overflow is used to argue for the conclusion that the machinery of phenomenology contains more than the machinery of cognitive accessibility" (2007, p. 487). Closely related claims can be found in Burge, 2007; Dretske, 1996, pp. 151–52, fn.10; and Dretske, 2006; Fodor, 2007; and Tye, 2006. Bayne and Chalmers, 2003, argue that Sperling's work shows a dissociation of A- and P-consciousness for complex (conjunctive) but not simple (individual letter) contents.

10. See Gegenfurtner and Sperling, 1993, for a detailed model of PR findings in terms of selective and nonselective transfer of information. Their picture involves (1) a pre-cue attentional state in which attention is mainly directed to the middle row resulting in nonselective transfer of letters biased toward the middle row; and (2) a post-cue attentional state where "attention shifts to the cued row of the display" (p. 865) and selective transfer of that row occurs.

11. If so, he would not be alone; the conflation of visible persistence and iconic memory was commonplace in the sixties and seventies before Coltheart's work.

12. See Coltheart, 1980, 1983, 2009; and Di Lollo, 1980. Note that these effects are not exhibited if the stimulus in question is so intense as to produce an afterimage. This is not the case in Sperling's experiments.

13. See Martin, 2001, who argues that both visual imagination and episodic memory involve the representation of experiencing. They differ only in the particularity of their contents.

14. For a critique of appeals to these experiments as evidence of overflow, see Phillips, forthcoming.

15. Visual imagery raises large questions beyond the scope of this paper. In particular, it would be interesting to know if those with poor or no (self-reported) visual imagery were equally capable of performing Brockmole's task at 1500 ms delays.

16. Note that neither visible nor informational persistence "has anything to do with positive or negative retinal after-images" (Coltheart, 2009, p. 429); the stimuli used in these experiments are not bright enough to generate after-imagery, and there are other "major differences," for which see Di Lollo, Clark, and Hogben, 1988.

17. Declan Smithies encouraged me to consider this idea. Recent work of Henry Shelvin's defends an account along roughly these lines.

18. Much more needs saying about the nature of generic phenomenology. For an introductory exploration of the related notion of scene "gist," see Oliva, 2005.

19. The term comes from Dennett and Kinsbourne, 1992; they do not endorse such a view.

20. A similar difficulty affects Tye's version of the Stalinesque account on which *momentary* experiences have temporally extended contents (2003, pp. 88–89). Applying Tye's account to Sperling, an experience of "grid followed by tone" comes "all at once" at some small delay after the tone and at a more substantial delay after the grid. In contrast to their simultaneous *representing*, the grid and cue tone are *represented* as occurring in succession, as they do in reality. The difficulty again is that the first moment at which the grid impinges on our conscious lives is with the momentary experience that comes after the tone. At this point, any role for attention is already complete, since the original grid experience is already upon us.

21. This idea is not completely outlandish, given the fact that PR paradigms require substantial training. Chow, 1985, notes that around one hundred trials are needed to obtain a significant effect.

22. Discussing Sperling, Tye writes, "Of course, the time at which the sensory memory is formed [that is, for Tye, when the array *appears*] need not be the same as the time at which the array is displayed or the time at which it disappears. In general, represented time need not be the same as time represented" (2006, p. 511). It is unclear how exactly Tye intends to exploit this freedom.

23. I focus on location (as opposed to objects, features, etc.) here for obvious reasons. However, I am not assuming the specialness or otherwise of location in respect of selective visual processing (see Tsal & Lavie, 1988, 1993; Cave & Pashler, 1995; and for philosophical deployment, Campbell, 2002).

24. This account explains the findings of Thomas and Irwin, 2006, and Irwin and Thomas, 2010, regarding the disruption to partial reporting of middle and upper rows due to blinking. As Coltheart and Coltheart put it, "a blink deflects visual attention downward even before the blink is initiated" (2010, p. 353). These findings are very hard to explain on imagistic accounts.

25. Or so I develop Dainton's view. Dainton draws substantially on Foster, 1979, 1982; see also Soteriou, 2007.

References

- Alpern, M. (1953). Metacontrast. *Journal of the Optical Society of America*, 43, 648–57.
- Averbach, E., & Sperling, G. (1961). Short term storage of information in vision. In C. Cherry (Ed.), *Information theory* (pp. 196–211). Washington, DC: Butterworth & Co.
- Ayer, A. J. (1956). *The problem of knowledge*. London: Macmillan.
- Bayne, T., & Chalmers, D. (2003). *What is the unity of consciousness?* In A. Cleeremans (Ed.), *The unity of consciousness* (pp. 23–58). Oxford: Oxford University Press.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18, 227–47.
- . (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30, 481–99.
- . (2008). Consciousness and cognitive access. *Proceedings of the Aristotelian Society*, 108, 289–317.
- Brockmole, J. R., Wang, F., & Irwin, D. (2002). Temporal integration between visual images and visual percepts. *Journal of Experimental Psychology*, 28, 315–34.
- Byrne, A., Hilbert, D. R., & Siegel, S. (2007). Do we see more than we can access? *Behavioral and Brain Sciences*, 30, 501–2.

- Burge, T. (2007). Psychology supports independence of phenomenal consciousness. *Behavioral and Brain Sciences*, 30, 500–1.
- Campbell, J. (2002). *Reference and consciousness*. Oxford: Oxford University Press.
- Cave, K. R., & Pashler, H. (1995). Visual selection mediated by location: Selecting successive visual objects. *Perception & Psychophysics*, 57(4), 421–32.
- Cattell, J. M. (1886). The inertia of the eye and brain. *Brain*, 8, 295–312.
- Choi, H., & Scholl, B. J. (2006). Perceiving causality after the fact: postdiction in the temporal dynamics of causal perception. *Perception*, 35, 385–99.
- Chow, S. L. (1985). Iconic store and partial report. *Memory and Cognition*, 13, 256–64.
- Coltheart, M. (1980). Iconic memory and visible persistence. *Perception and Psychophysics*, 27, 183–228.
- . (1983). Iconic memory. *Philosophical Transactions of the Royal Society of London, Series B*, 302, 283–94.
- . (2009). Memory, iconic. In T. Bayne, A. Cleeremans, & P. Wilken (Eds.), *The Oxford companion to consciousness* (pp. 427–31). Oxford: Oxford University Press.
- Coltheart, M., & Coltheart, V. (2010). Visual memories. In V. Coltheart (Ed.), *Tutorials in visual cognition* (pp. 349–65). London: Psychology Press.
- Cowan, N. (1984). On short and long auditory stores. *Psychological Bulletin*, 96(2), 341–70.
- Creel, W., Boomslinger, P. C., & Powers, S. R. (1970). Sensations of tone as perceptual forms. *Psychological Review*, 77(6), 534–45.
- Dainton, B. (2000). *Stream of consciousness*. London: Routledge and Kegan Paul.
- . (2008). The experience of time and change. *Philosophy Compass*, 3/4, 619–38.
- Dehaene, S., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, 10, 204–11.
- Dennett, D. (1991). *Consciousness explained*. Boston: Little, Brown.
- Dennett, D., & Kinsbourne, M. (1992). Time and the observer. *Behavioral and Brain Sciences*, 15, 183–247.
- Di Lollo, V. (1980). Temporal integration in visual memory. *Journal of Experimental Psychology: General*, 109, 75–97.
- Di Lollo, V., Clark, C. D., and Hogben, J. H. (1988). Separating visible persistence from retinal afterimages. *Perception & Psychophysics*, 44(4), 363–68.
- Doughty, J. M., & Garner, W. R. (1947). Pitch characteristics of short tones, two kinds of pitch threshold. *Journal of Experimental Psychology*, 37, 351–65.
- Dretske, F. (1969). *Seeing and knowing*. Chicago: University of Chicago Press.
- . (1996). Phenomenal externalism or if meanings ain't in the head, where are qualia? *Philosophical Issues*, 7, 143–58.
- . (2006). Perception without awareness. In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 147–80). Oxford: Oxford University Press.
- Dufour, A., Touzalin, P., Moessinger, M., Brochard, R., & Després, O. (2008). Visual motion disambiguation by a subliminal sound. *Consciousness and Cognition*, 17, 790–97.
- Eagleman, D. M., & Sejnowski, T. J. (2000). Motion integration and postdiction in visual awareness. *Science*, 287, 2236–38.
- . (2003). The line-motion illusion can be reversed by motion signals after the line disappears. *Perception*, 32, 963–68.

- Enns, J. T., & Di Lollo, V. (2000). What's new in visual masking? *Trends in Cognitive Sciences*, 4, 345–52.
- Fernberger, S. W. (1921). A preliminary study of the range of visual apprehension. *American Journal of Psychology*, 32(1), 121–33.
- Fodor, J. A. (2007). The revenge of the given. In B. P. McLaughlin & J. D. Cohen (Eds.), *Contemporary debates in philosophy of mind*. Malden, MA: Blackwell.
- Foster, J. (1979). In self-defence. In G. F. Macdonald (Ed.), *Perception and Identity* (pp. 175–82). London: Macmillan.
- . (1982). *The Case for Idealism*. London: Routledge & Kegan Paul.
- Gegenfurtner, K. R., & Sperling, G. (1993). Information transfer in iconic memory experiments. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 845–66.
- Gill, N. F., & Dallenbach, K. M. (1926). A preliminary study of the range of attention. *American Journal of Psychology*, 37(2), 247–56.
- Grush, R. (2007). A plug for generic phenomenology. *Behavioral and Brain Sciences*, 30, 504–505.
- Hamilton, W. (1859). *Lectures on metaphysics and logic* (vol. 1). Edinburgh: Blackwood.
- Hatfield, G. (1995). Attention in early scientific psychology. In R. Wright (Ed.), *Visual attention* (pp. 3–25). Oxford: Oxford University Press.
- Irwin, D. E., & Thomas, L. E. (2010). Eyeblinks and cognition. In V. Coltheart (Ed.), *Tutorials in visual cognition* (pp. 121–141). London: Psychology Press.
- Iwasaki, S. (1993). Spatial attention and two modes of visual consciousness. *Cognition*, 49, 211–33.
- Jevons, W. S. (1871). The power of numerical discrimination. *Nature*, 3, 281–82.
- Kaufman, E. L., Lord, M. W., Reese, T. W., & Volkman, J. (1949). The discrimination of visual number. *American Journal of Psychology*, 62, 498–525.
- Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences*, 11(1), 16–22.
- Kosslyn, S. M., Thompson, W. L., & Ganis, G. (2006). *The case for mental imagery*. Oxford: Oxford University Press.
- Lamme, V. A. (2003). Why visual attention and awareness are different. *Trends in Cognitive Science*, 7, 12–8.
- Landman, R., Spekreijse, H., & Lamme, V. A. F. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, 43, 149–64.
- Loftus, G. R., & Irwin, D. E. (1998). On the relations among different measures of visible and informational persistence. *Cognitive Psychology*, 35, 135–99.
- Mack, A., & Rock, I. (1998). *Inattention blindness*. Cambridge, MA: MIT Press.
- Mackay, D. M. (1958). Perceptual stability of a stroboscopically lit visual field containing self-luminous objects. *Nature*, 181, 507–508.
- Martin, M. G. F. (1997). The shallows of the mind. *Proceedings of the Aristotelian Society*, 74(1), 75–98.
- . (2001). Out of the past: Episodic recall as retained acquaintance. In C. Hoerl & T. McCormack (Eds.), *Time and memory: Issues in philosophy and psychology* (pp. 257–84). Oxford: Oxford University Press.
- Mole, C. (2008). Attention and consciousness. *Journal of Consciousness Studies*, 15(4), 86–104.
- Naccache, L. (2006). Is she conscious? *Science*, 313, 1395–96.
- Neisser, U. (1967). *Cognitive psychology*. Englewood Cliffs, NJ: Prentice-Hall.

- Oberly, H. S. (1924). The range for visual attention, cognition and apprehension. *American Journal of Psychology*, 35(3), 332–52.
- Oliva, A. (2005). Gist of the scene. In L. Itti, G. Rees, & J. K. Tsotsos (Eds.), *Encyclopedia of neurobiology of attention* (pp. 251–56). San Diego, CA: Elsevier.
- Papineau, D. (2002). *Thinking about consciousness*. Oxford: Oxford University Press.
- . (2007). Reuniting (scene) phenomenology with (scene) access. *Behavioral and Brain Sciences*, 30, 521.
- Phillips, I. B. (2009). *Experience and time*. PhD thesis, University College, London.
- . (2010). Perceiving temporal properties. *European Journal of Philosophy*, 18(2), 176–202.
- . (forthcoming). Perception and iconic memory. *Mind & Language*.
- Pitcher, G. (1971). *A theory of perception*. Princeton, NJ: Princeton University Press.
- Sekuler, R., Sekuler, A. B., & Lau, R. (1997). Sound alters visual motion perception. *Nature*, 385, 308.
- Sligte, I. G., Scholte, H. S., & Lamme, V. A. F. (2008). Are there multiple visual short-term memory stores? *PLoS ONE*, 3, 1–9.
- Soteriou, M. (2007). Content and the stream of consciousness. *Philosophical Perspectives*, 21, 543–68.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, 74(11), 1–29.
- Steinmann, R., Pizlo, Z., & Pizlo, F. (2000). Phi is not beta, and why Wertheimer's discovery launched the Gestalt revolution. *Vision Research*, 40, 2257–64.
- Stevens, S. S., & Davis, H. (1938). *Hearing*. New York: Wiley.
- Thomas, L. E., & Irwin, D. E. (2006). Voluntary eyeblinks disrupt iconic memory. *Perception and Psychophysics*, 68, 475–88.
- Thornton, I. M., & Fernandez-Duque, D. (2002). Converging evidence for the detection of change without awareness. In J. Hyönä, D. P. Munoz, W. Heide, & R. Radach (Eds.), *The Brain's Eyes: Neurobiological and Clinical Aspects of Oculomotor Research* (pp. 99–118). Amsterdam: Elsevier Science.
- Titchener, E. B. (1910). Attention as sensory clearness. *Journal of Philosophy, Psychology and Scientific Methods*, 7(7), 180–82.
- Tsal, Y., & Lavie, N. (1988). Attending to color and shape: the special role of location in selective visual processing. *Perception & Psychophysics*, 44(1), 15–21.
- . (1993). Location dominance in attending to color and shape. *Journal of Experimental Psychology: Human Perception and Performance*, 19(1), 131–39.
- Tye, M. (2003). *Consciousness and persons*. Cambridge, MA.: MIT Press.
- . Tye, M. (2006). Content, richness, and fineness of grain. In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 504–530). Oxford: Oxford University Press.
- Watanabe, K., & Shimojo, S. (2001). When sound affects vision: effects of auditory grouping on visual motion perception. *Psychological Science*, 12, 109–16.
- Weiskrantz, L. (1997). *Consciousness lost and found: A neuropsychological exploration*. Oxford: Oxford University Press.
- Weisstein, N., & Wong, E. (1986). Figure-ground organization and the spatial and temporal responses of the visual system. In E. C. Schwab & H. C. Nusbaum (Eds.), *Pattern recognition by humans and machines, Volume 2, Visual perception* (pp. 31–64). San Diego: Academic Press.

- Wertheimer, M. (1912). Experimentelle Studien über das Sehen von Bewegung. *Zeitschrift für Psychologie*, 61, 161–265.
- Whipple, G. M. (1914). *Manual of mental and physical tests, part 1: Simpler processes*. Baltimore: Warwick & York.
- Wilken, P. C. (2001). *Capacity limits for the detection and identification of change: Implications for models of visual short-term memory*. PhD. thesis, University of Melbourne.
- Wilken, P. C., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, 4(1–3), 1–16
- Wolfe, J. M. (1999). Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting memories* (pp. 71–94). Cambridge, MA: MIT Press.
- Woodworth, R. S. (1909). Review of *a text-book of psychology* by E. B. Titchener. *Journal of Philosophy, Psychology and Scientific Methods*, 6(25), 692–94.