

Internalism, Active Externalism, and Nonconceptual Content: The Ins and Outs of Cognition

Terry Dartnall

Information and Communication Technology, Griffith University, Australia

Received 13 February 2006; received in revised form 31 May 2006; accepted 2 June 2006

Abstract

Active externalism (also known as the extended mind hypothesis) says that we use objects and situations in the world as external memory stores that we consult as needs dictate. This gives us economies of storage: We do not need to remember that Bill has blue eyes and wavy hair if we can acquire this information by looking at Bill. I argue for a corollary to this position, which I call ‘internalism.’ Internalism says we can acquire knowledge on a need-to-know basis by consulting portable, inner analogues of the world. This, however, leads to a dilemma. If the knowledge was stored in memory, it is difficult to see how we could have acquired it by consulting inner analogues, for it would seem that we knew it already. Moreover, if it was stored in memory, we lose the economies of storage that provide much of the rationale for external memory stores and hence for their inner analogues. If, on the other hand, the knowledge was *not* stored in memory, it is difficult to see how we could have acquired it by consulting inner analogues, for it was not there to be acquired. I propose a solution to this problem that turns on the concept of nonconceptual content, and I relate the solution to Stephen Kosslyn’s (1994) architecture of perception and visual mental imagery. Viewed in a broader context, the solution shows that the world leaks into the mind, as well as mind leaking into the world.

Keywords: Active externalism; Extended mind; Internalism; Imagery; Nonconceptual content; Epistemology; Cognitive architecture

1. Introduction

Active externalism, also known as the extended mind hypothesis, says that the mind leaks or extends into the world (Clark, 1997, 2003; Clark & Chalmers, 1998; Dennett, 1996; Donald, 1991; Hutchins, 1995). Clark and Chalmers said that cognitive *processes* extend into the world

when we use pen and paper to work something out, when we use a computer, or even when we use language, which Clark (2003) thought was the first technology. Clark and Chalmers said that cognitive *states* extend into the world when we use objects and states of affairs, or diaries and address books, as external memory stores that we can consult as needs dictate. We do not need to remember that Bill has brown eyes and wavy hair if we can access this information by looking at him.

In Dartnall (2005), I argued for a corollary and complement to this position, which I called ‘internalism.’ Internalism says that the world leaks into the mind. We can perform actions in our heads that we would normally perform in the world, and we typically perform them on inner analogues of external objects. This can lead to empirical discovery: If an operation would have yielded empirical discovery when performed in the world, it will also yield empirical discovery when performed in the mind. We can manually rotate a jigsaw piece to discover where it fits, or we can mentally rotate it, when not in the presence of the puzzle, to discover where it fits. Both cases give noninferential, empirical discovery.

In some ways we would *expect* this to be the case. One problem with using external objects as memory stores is that most external objects are not portable. Having inner analogues of external objects would free us from the here-and-now, the context, and the moment. It would give us the ability to plan and anticipate in the safety of the cave. It would bring together in a single, central location knowledge about the world that is stored in different parts of the brain; and it would enable us to process it using cognitive (most obviously perceptual) resources that are already in place—and the redeployment of available resources is invariably nature’s way.¹

However, now we have a problem. Suppose we scan an inner analogue of Bill and discover that he has brown eyes and wavy hair. If this information was not stored in memory, it is difficult to see how we could have discovered it, for it was not there to be discovered. If it *was* stored in memory, we have two subproblems. The first subproblem is that it is difficult to see how we can discover something that we have stored in memory; for if it is stored in memory, it would seem that we know it already. The second, and probably more important, subproblem is that active externalism says that we use the world as an external memory store to *avoid* storing knowledge in memory. We have high-level knowledge about how to access specific knowledge (such as the color of Bill’s eyes), and we have the ability to zoom in and access this knowledge on a need-to-know basis. If inner analogues are to be useful, they must provide us with portable data stores that we can consult in a similar way. However, storing this knowledge in memory undermines this particular rationale for inner analogues. Put differently, if we use inner analogues as portable memory stores, we would expect to be able to retrieve knowledge from them in the same way that we retrieve it from the world: in a cost-efficient way, on a need-to-know basis, but now with the advantage of portability. However, this leads to a dilemma: If the knowledge is stored in memory, then (a) it seems that we know it already, so that we cannot discover it; and (b) we seem to lose the economies of storage that provide the rationale for external memory stores and that, by analogy, provide a rationale for inner analogues. If, on the other hand, it is *not* stored in memory, it is difficult to see how we can discover it by consulting inner analogues. I mention this problem in Dartnall (2005), but do not pursue it there.

In this article I focus on visual mental images as a subset of inner analogues and outline three ways in which we can acquire knowledge by inspecting inner analogues in this sense,

all of which avoid the dilemma. The first two ways are relatively straightforward. We can further conceptualize, or further process, visual images to give us *relational knowledge*; and we can further conceptualize visual images through a process of *seeing as*. The third case is more complex and is the centerpiece of this article. I believe that we can acquire knowledge by scanning images that we have generated by conceptualizing stored *nonconceptual content*. More on this in a moment.

An example of acquiring relational knowledge by inspecting visual images is discovering that the tip of a horse's tail is lower than its rear knees; that a mug is taller than it is wide; or that a jigsaw piece, when rotated, fits into the puzzle in a certain way. An example of acquiring knowledge through *seeing as* is discovering that Bill is ugly, or that he looks intelligent or tired. Both of these cases (relational knowledge and *seeing as*) avoid the dilemma: The knowledge that we acquire was not stored but was derived by further processing what we already knew.

This gives us a partial symmetry between external objects and inner analogues, both of which we can interrogate on a need-to-know basis. However, the symmetry seems to be incomplete. We can interrogate images to get relational knowledge and *seeing as* knowledge, and in both cases we acquire knowledge that was neither stored nor inferred from knowledge that was stored. It seems, nevertheless, that we need to store 'first-order,' or what I call 'constitutive,' knowledge to generate the images from which to derive the 'second-order' knowledge. It seems, for example, that we need to store information about the shape of Bill's face and the color of his eyes to construct an image of him that we can scan for ugliness or intelligence. To the extent that we have to store this constitutive knowledge we lose the economies of storage that accrue from using the world as an external memory store—economies of storage that should carry over into internal analogues.

I show that we do not have to store this knowledge. One of my main claims in this article is that we can acquire knowledge by scanning images that we have generated by conceptualizing stored *nonconceptual content* (Evans, 1982; McDowell, 1994; Peacocke, 1992). The basic idea behind nonconceptual content is that our mental states can have content that is too detailed and fine grained to be captured in its entirety by the concepts we possess: We can, for example, discriminate more shades of color, more shapes, more textures, and more sounds than we can conceptualize. This does not mean that we cannot conceptualize such content. We can, and the availability of nonconceptual content makes a rich array of conceptualized content and judgment possible. The point is simply that nonconceptual content cannot be captured in its entirety by the concepts we possess.

I argue that we copy nonconceptual content from perception into memory—most obviously into visual memory. On a need-to-know basis we can then generate imagery by conceptualizing this content and scan this imagery for the information we require. Storing nonconceptual content in a format similar to perceptual input enables us to conceptualize the content in multiple ways, on a need-to-know basis, taking advantage of the richness and diversity of the input. Among other things, this respects Marr's (1982) 'principle of least commitment,' which says that we should not discard information that we might need later on, for problems we have not yet encountered.

The inner–outer, internalism–externalism symmetry is now complete. We use objects and states of affairs in the world as memory stores that we consult as needs dictate, and we do

much the same with inner analogues. There are different but probably equivalent economies of storage in the two cases, as we shall see.

The structure of this article is as follows. In sections two and three I revisit active externalism and internalism, respectively. Section four visits the imagery debate and provides *prima facie* reasons for preferring pictorialism (also known as ‘depictivism’) to descriptivism and Perceptual Activity Theory. Section five outlines Stephen Kosslyn’s (1994) pictorialist (depictivist) model of perception and mental imagery in more detail.² Kosslyn (1994) said that perceiving an object and having an image of that object (e.g., a fox’s head) deploy the same mechanisms once the configuration of activity exists in the visual buffer. Our intuition, however, probably tells us that in the case of perception the information is ‘out there’ in the world and that we access it by looking at the fox, whereas with imagery the information is in our heads and we use it to generate an image that we scan for the information we require. I address this intuition in section six by showing that we can acquire relational knowledge and *seeing as* knowledge by further processing inner analogues. When we do this we acquire knowledge that was neither stored in memory nor inferred from what was stored in memory. Now, however, our intuition is probably that we must have stored these analogues as conceptualized images or conceptualized content. I counter this intuition with another one. I then show that Kosslyn’s (1994) model features compressed images that are stored in visual memory and not conceptualized until they are triggered by concepts in associative memory. These compressed images play much the same role in knowledge storage and retrieval as objects and states of affairs in the external world: When conceptualized, they generate patterns of activity in the visual buffer, which can be scanned for information using the same mechanisms that we use in our perception of the world. In section seven I focus on the philosophical issues and show that these compressed images give us stored, nonconceptual content that can support multiple judgments, which we typically make on a need-to-know basis. Section eight provides a discussion and looks at some possible objections.

2. Active externalism

Two broad arguments have been advanced for active externalism: the *parity argument* and the *complementarity argument* (for a comparison, see Sutton, in press). The parity argument (Clark & Chalmers, 1998) says that if something counts as cognitive when it is performed in the head it should also count as cognitive when it is performed in the world. Suppose you have to rotate images of geometrical shapes on a computer screen. You can rotate them using a neural implant or you can rotate them using a ‘rotate’ button in the world. Presumably, we will say that the implant case is cognitive—so why is the button case not cognitive as well? Clark and Chalmers said that epistemic credit is due where epistemic actions are performed, regardless of whether they are performed in the head or in the world.

This, however, only covers cognitive *processes*, and Clark and Chalmers (1998) admitted that the processes might be in the world while all of our ‘truly mental states—experiences, beliefs, desires, emotions, and so on’ (p.12) might be in the head. Therefore, they took the parity argument a stage further and argued that cognitive states can be constituted partly by features of the environment. For example, Otto suffers from Alzheimer’s disease. He hears

that there is an exhibition at the Museum of Modern Art. He consults his notebook, which says that the museum is on 53rd street. He walks to 53rd street and goes to the museum. Clark and Chalmers said that the notebook plays the same role for Otto that biological memory plays for the rest of us. It just happens that ‘this information lies beyond the skin’ (p. 13).

Clark and Chalmers’s (1998) claim is not only that Otto’s belief is out there in the world. They also said that Otto believed the museum was on 53rd street before he looked it up, courtesy of the ‘functional isomorphism’ between the notebook entry and a corresponding entry in biological memory.

The complementarity argument for active externalism (Clark, 2003) has two broad components: the neural opportunism argument and what I call the ‘bioborg’ argument (Dartnall, 2004b). The bioborg argument says that we have neural and biomechanical subsystems that we either launch to do our bidding (when we walk across the room to the fridge, or when we sign a check or tie our shoelaces) or that get on with the job of interacting with the world without reporting back to consciousness at all. The latter is nicely illustrated by the Titchener/Aglioti Circles Illusion (Aglioti, DeSouza, & Goodale, 1995; Gazzaniga, 1998; Goodale & Milner, 2004; Milner & Goodale, 1995). We *consciously* see a plastic disk as larger than it is, but when we reach for it, our thumb and forefinger form the correct sized aperture to pick it up. A plausible explanation is that there are two incoming visual pathways: the *ventral* and the *dorsal*. The ventral system deals with thinking and planning at the conscious level, whereas the dorsal stream deals with the fine-grained details of interacting with the world and does not report back to consciousness. The ventral system is fooled by the illusion, but the dorsal system is not. When we pick up the disk, we launch a ‘pick-it-up’ routine that is guided by information coming in through the dorsal stream—information that we do not have access to at the conscious level.

The neural opportunism argument says that we have opportunistic brains that use enduring features of the world as external memory stores. When we saccade around a room we return to the same place time and again—because *that is where the information is*. We have a high-level map of what is out there, plus the ability to zoom in and retrieve detailed information on a need-to-know basis. This opportunism has operated in consort with our bioborg brains to create what active externalists call ‘cognitive technologies,’ that complement our basic cognitive abilities: We have constructed external aids and devices to complement our already-tool-kit brains. Clark (2003) thought that the first cognitive technology was probably language. Since then we have upgraded our mindware through writing, increasingly sophisticated forms of printing, and are now entering into an intimate relationship with machines.

The upshot is that mind leaks or loops out into the world. Cognitive processes loop into the world when we use technologies such as pen and paper, or pocket calculators, to work something out; and cognitive states loop into the world when we use objects and states of affairs, or diaries or address books, as external memory stores that we can consult as needs dictate.

3. Internalism

Active externalism says that the mind leaks into the world. Internalism says that the world leaks into the mind, so that the leakage is in both directions. In Dartnall (2005), I argued for

this through an example, by looking at the symmetry between internalism and externalism and through considerations of portability and efficiency.

The example is this: We look at a jigsaw puzzle, leave the room, and realize (when we are no longer in the presence of the puzzle) that if we rotate one of the pieces it will fit into the puzzle in a certain way.³ We do not discover this through straightforward empirical discovery because we are not in the presence of the puzzle at the time. We do not *remember* it, because we can only remember what we knew when we were in the room, and when we were in the room we did not know where the piece fitted into the puzzle. We do not *infer* it, for there are no covering laws (true general statements) from which it can be derived—or so I argued in Dartnall (2005). I argued that we discover it by performing an operation in our heads that we would normally perform in the world, so that we make an empirical discovery by performing an operation in our heads. Manually rotating the piece would have given us an empirical discovery in which we discovered something about the world that we did not know before. Mentally rotating the piece while we were looking at the puzzle would also have given us an empirical discovery. What difference does it make if we mentally rotate the piece when we are *not* looking at the puzzle? This will give us an empirical discovery as well.

Now for the symmetry argument: Active externalism says that when we use a rotate button to rotate shapes on a screen we are performing a cognitive act. Internalism says that when we mentally rotate the jigsaw piece we are performing an action in our minds that we would normally perform in the world. That is the first symmetry—performing cognitive actions in the world (active externalism) and performing actions in our heads that we would normally perform in the world (internalism). However, active externalism also says that we use the world as an external data store that we can consult as needs dictate. If the symmetry carries over, we will have portable inner analogues of external data stores that we can consult as needs dictate. This is a weak reason, little more than a hunch, for believing that we have inner analogues.

Another reason that might be proposed for saying that we have inner analogues is this: If we perform operations in our minds that we would normally perform in the world—more specifically, if we perform operations in our minds *in the absence of the original object*—then we must have inner, remembered analogues to perform these operations on. In the case of the jigsaw puzzle, we perform an operation in our minds that we would normally perform in the world. We say that we ‘rotate the piece in our minds.’ What we rotate is an inner analogue of the piece.

This argument is at least problematic. Descriptivists, such as Pylyshyn (1973, 1978, 1981, 2002, 2003), will say that all of our knowledge is either stored as inner symbolism or derived from this symbolism through inference. The symbolism enables us to generate images, but the images themselves play no functional role in cognition and consequently should not be seen as inner analogues of objects and states of affairs in the world.

The final reason I consider in this section for believing that we have inner analogues is that most external objects are not portable. We might use a blackboard or faces as memory stores when they are in front of us, but they are difficult things to carry around. Inner analogues would overcome this problem and enable us to think and plan in an unhurried way, in safety, with the rest of the community, about objects and states of affairs that are not spatio-temporally present.

(Clark & Chalmers, 1998, acknowledge the importance of portability.) Being able to think and plan in such a way would have given us a huge evolutionary advantage. Evolution typically deploys mechanisms and capabilities that are already in place. This suggests that it would have redeployed our perceptual and cognitive capabilities, now applied to inner analogues of the world, and that we use the same strategy with these analogues that we use with their objective counterparts—consulting them as needs dictate. Suppose we have learned to use the world as an external data store and suppose we need to free ourselves from the context and the moment. One option is to develop a propositional, declarative database couched in an inner language of thought. This will require us to develop new cognitive resources, akin to developing pencil and paper if we want to record everything in the room. A more efficient strategy is to deploy cognitive capabilities that are already in place. This immediately leads to the idea of inner analogues that we process in the same way that we process objects and states of affairs in the world. There is a sense in which economies of storage and requirements of portability and safety, when coupled with evolutionary pressure, force inner analogues on us.⁴

4. Analogues, images, and processes

The previous arguments are speculative, problematic, or highly theoretical. They do, however, give us some initial grounds for believing that we have inner analogues that function in a similar fashion to external memory stores. I now focus on visual mental imagery and put a picture together that provides explanatory leverage and shows how inner analogues can give us knowledge that was neither stored in memory nor inferred from what was stored in memory. In this section, I briefly visit the imagery debate (Kosslyn, 1980, 1994; Kosslyn, Thompson, & Ganis, 2006; Pylyshyn, 1973, 1978, 1981, 2002, 2003). I do not wish to become mired in this debate, which Thomas (1994) called ‘a tangled mess of argument’ (p. 291). I outline the two main positions, pictorialism (depictivism) and descriptivism, and the more recent Perceptual Activity Theory (Thomas, 1999). I then give *prima facie* reasons for adopting a pictorialist position. In section 5 I look at Kosslyn’s (1994) pictorialist model in more detail.

Pictorialism (Kosslyn, 1980, 1994; Kosslyn, Thompson, & Ganis, 2006) says that mental images are ‘quasi-pictures’ that are functionally equivalent to inner pictures. These quasi-pictures are generated in a visual buffer by concepts in associative memory activating compressed images stored in visual memory. We scan the quasi-pictures and extract the information we require. If we want to know the shape of a German shepherd’s ears, whether frogs have lips, or whether a collie’s head is higher than the bottom of a horse’s tail, we can construct an image in the visual buffer and scan it for the information we require. Kosslyn said (personal communication, August 7, 2005), ‘Imagery is in most respects just like vision, except that there’s nothing there.’

Descriptivism (Pylyshyn, 1973, 1978, 1981, 2002, 2003) says that our knowledge and cognitive capabilities are encoded in an inner, symbolic language of thought. Descriptivism’s core commitment is that these language-like representations are sufficient to account for all cognitive capabilities: Our knowledge is either stored as inner symbolism or derived from

inner symbolism through inference. This symbolically encoded knowledge, coupled with our inferential abilities, generates images that are the epiphenomenal, causally impotent by-products of our symbolic cognitive processes. Our perception that they play a role in memory retrieval and information processing is simply an illusion. Images exist, at least in the sense that we experience imagery, but they are not *stored* as images and play no functional role in cognition.

A radically different account of imagery is provided by Perceptual Activity Theory, which according to Thomas (1999) dates back to the 1960s and 1970s. Thomas (1999) said that perception is the “continual updating and refining of procedures (or ‘schemata’; see Neisser, 1976) that specify how to direct our attention most effectively in particular situations: how to efficiently examine and explore, and thus interpret, a scene or object of a certain type (Stark & Ellis, 1981).” (p. 218).

We experience imagery when we apply these schemata inappropriately, either to the wrong things in the world or to nothing at all: “We imagine, say, a cat, by going through (some of) the motions of examining something and finding it is a cat, although there is no cat (and perhaps nothing relevant at all) there to be examined. Imagining a cat is seeing nothing-in-particular *as* a cat (Ishiguro, 1967)” (Thomas, 1999, p. 218). If we want to know whether cats have pointed ears we deploy the schema we would have used if the cat was in front of us. If we perform pointed-ear scans we have retrieved the knowledge, implicit in our procedures, that cats have pointed ears.

According to Perceptual Activity Theory there are no images. We have an image-like *experience* of a cat when we employ catty schemata in the absence of a cat, but this only means that we have the same kind of experience that we would have if we were seeing a cat in the world. To say that there is both an image and an experience of an image is a kind of double counting. It is like saying that when we experience pain there is both the pain and the experience of the pain.⁵

In this article I adopt a pictorialist approach toward visual mental imagery. This is for two reasons. The first reason is that pictorialism, generally, and Kosslyn’s (1994) model, in particular, naturally combine with internalism and an account of nonconceptual account memory to give us explanatory leverage and solve the problem we set out to address. Kosslyn (1994) repeatedly said that imagery is generated and interpreted by the same operations and mechanisms that we use in perception. Later we look at Rick Grush’s (2004) emulation theory of representation, which synthesises Kosslyn’s (1994) theory into a broader explanatory framework. Both theories tell the same evolutionary story: We first learned to perceive and interpret objects and situations in the world; then we ran these operations offline to give us internal analogues of the world.⁶

Kosslyn (1994) also said that the visual buffer is a stage in perceptual information processing, so that the contents of the buffer when we scan visual images are much the same as its contents when we perceive the external world. This is naturally compatible with internalism, which says that we can scan an image for information in much the same way that we scan something in the world for information. Adding nonconceptual memory to this framework increases its explanatory leverage.

The second reason for adopting pictorialism is that visual mental images appear to be retinotopically organized in visual cortex, in the sense that the neurons are organized to

preserve the structure of the retina. Kosslyn, Thompson, and Ganis (2006) said, ‘There is good evidence that the brain depicts representations literally, using space on the cortex to represent space in the world’ (p. 15); ‘These areas represent information depictively in the most literal sense; there is no need to talk about an abstract functional space akin to that defined by an array in a computer’ (Kosslyn, 1994, p. 13). Tootell, Silverman, Switkes, and De Valois (1982) injected a monkey with radioactively tagged sugar (2-deoxyglucose) while the monkey stared at a pattern of flashing lights. The more active a neuron was, the more sugar it took up. The amount of radioactivity was then measured across the cortex, and an image of the pattern was clearly visible in area V1 (primary visual cortex), the first cortical area to receive input from the eyes. The results are visually startling and show a pattern of activation in the cortex that clearly resembles the pattern that the monkey was staring at. Fox et al. (1986) used PET scans to show that area V1 is retinotopically mapped in humans. Kosslyn (1994) said that a lot of information runs backwards in the visual system, so that information stored in higher areas might be able to ‘evoke a pattern of activity in at least some of the retinotopically mapped areas—which would produce a mental image’ (p. 15). Kosslyn, Thompson, and Ganis (2006) cited, ‘many studies that have documented that visual imagery activates topographically organized cortex’ (p. 110), including Klein et al. (2004) and Slotnick, Thompson, and Kosslyn (2005). Kosslyn, Thompson, and Ganis (2006) provided a summary of this research in an appendix.⁷

We now look at Kosslyn’s (1994) model in more detail.

5. Kosslyn’s pictorialist model

Kosslyn (1994) said that imagery, rather than ‘piggybacking on perception,’ is an ‘integral part of how perception operates’ (p. 21). Consequently, he devoted much of his work (Kosslyn, 1994) to developing a theory of perception. I provide a greatly simplified account of his model and follow his practice of using italicized phrases to refer to subsystems of the model. (As mentioned above, Kosslyn’s 1994 model is summarized in Kosslyn, Thompson, and Ganis [2006].)

Kosslyn (1994) considered how the visual system identifies a fox whose rear legs and tail are occluded by a barrier. The system begins with a retinotopic image in the *visual buffer*, a set of ‘topographically organized visual areas (in the occipital lobe) that are used to segregate figure from ground’ (p. 70). The *attention-shifting subsystem* shifts the *attention window* to the appropriate place in the buffer (based on prior information about the overall shape that was encoded). In this case, Kosslyn (1994) assumed that the attention window would surround the fox and part of the barrier. The contents of the attention window are sent to the *ventral system* and the *dorsal system* for further processing.

The ventral system sends the image and its extracted properties to the *pattern activation subsystems* where the image is matched against representations stored as compressed images. The compressed image representation for *fox* is activated and provides imagery feedback to the visual buffer, where it is matched against the input image as well as possible. If the overlap is adequate, the *category pattern activation subsystem* sends a pattern code to *associative memory*.

While the ventral system is encoding object properties, the dorsal system is computing the location, size, and orientation of the fox. It sends this information to both the ventral system and the *categorical spatial relations encoding subsystem* and *coordinate spatial relations encoding subsystem*. These subsystems process the information and forward it to associative memory. Associative memory, armed with the information it has received from the ventral and dorsal systems, tentatively identifies the object as a fox (in a different context, such as a fox hunt, it would be less tentative). This hypothesis is now tested. Other subsystems access representations of foxes and pull out the most distinctive feature, which in this case is probably its head. Another subsystem moves the attention window, zooming it in on the fox's head, while the pattern code of a prototypical fox's head is sent to the pattern activation subsystems to augment the image in the visual buffer. The system then engages in a second cycle. All being well, the fox representation is activated in associative memory.

The gist of this complicated story is this: The attention window selects a configuration of activity in the visual buffer and sends it to the ventral and dorsal systems for further processing. The ventral system encodes properties such as shape, color, and texture; the dorsal system encodes properties such as location, size, and orientation. Information from these systems is then sent to associative memory, where it either (uncontentiously) activates a single representation or generates a hypothesis. If it generates a hypothesis, an information lookup system determines the most salient feature of the representation in associative memory (such as the head) and sees if this matches the activation in the visual buffer.

It is now a matter of filling in the dots to see how the model can account for visual imagery. The basic idea is that we access the representation of an object in associative memory and send its 'pattern code' to the pattern activation subsystems. The pattern code activates or 'primes' the modality-specific representation in the appropriate pattern activation subsystem. This representation is primed so much that it sends feedback to earlier areas, engendering a configuration of activity (an image) in the visual buffer.

According to this model, imagery is generated by the same mechanisms that underlie perception. Kosslyn (1994) said:

I have argued throughout this book that imaged objects are interpreted using exactly the same mechanisms that are used to interpret objects during perception. Once a configuration of activity exists in the visual buffer, input is sent to the ventral and dorsal systems and is processed in the usual ways—regardless of whether the activity arose from immediate input from the eyes or from information stored in memory. This notion lies at the heart of the similarities found between imagery and perception. (p. 336; see also pp. 287, 344.)

And, of course, it cuts both ways: The processes that underlie mental imagery play a role in normal perception (Kosslyn, 1994, p. 152). The big picture is that 'Once a pattern of activation is evoked in the visual buffer, it is processed the same way, regardless of whether it was evoked by input from the eyes (perception) or from memory (imagery)' (Kosslyn, 1994, p. 74).

We now return to the problem this article is addressing. Active externalism says that we use things in the world as external memory stores that we consult as needs dictate. However, most of these things are not portable. We cannot carry them with us and consult them on a need-to-know basis. Given that we have inner analogues (here focusing on visual images), we would expect these analogues to give us the same economies of storage as their external

counterparts. However, how can we access the images for information if that information was not already stored in memory? And, if it was stored in memory, what happens to the economies of storage?

According to Kosslyn (1994), perceiving an object and having an image of that object (say, a fox's head) deploy the same mechanisms once the configuration of activity exists in the visual buffer. In the case of perception, we can retrieve information that we did not have before. We might, for example, discover that foxes have pointed ears. When we have an *image* of a fox we can similarly scan it and discover that foxes have pointed ears.⁸ In the case of perception the information is out there in the world, and we access it by looking at the fox; but in the imagery case, our intuition probably tells us that the information is in our heads, and we use it to generate the image of the fox. After all (our intuition says), how can we generate an image of a fox with pointed ears if we did not already know that foxes have pointed ears? In the next two sections I show that this intuition is false.

6. Relational knowledge and seeing as

I now look at two scenarios in which we can acquire new knowledge by inspecting visual mental images. This is knowledge that was neither stored in memory nor inferred from knowledge that was stored in memory. I cite these scenarios as examples rather than an exhaustive classification. There might be other types of knowledge that we can acquire in this way.

6.1. Relational knowledge

Kosslyn (1994) said that 'perhaps the most basic property of visual mental imagery' (p. 335) is that images make the local geometry of objects available to us: "For example, one can determine from memory whether the tip of a racehorse's tail is above its rear knees, whether the uppercase version of the letter *a* contains an enclosed region, or whether a mug is taller than it is wide. The image depicts the spatial relations among portions of an object or scene, allowing one to interpret them in a novel way" (Kosslyn, 1994, p. 335). The jigsaw puzzle case is another example. We mentally rotate one of the pieces when we are not looking at the puzzle and discover that it fits in a certain way.

Kosslyn (1994) said that at first blush these results seem to be inconsistent with Chambers and Reisberg's (1985) well-known results. Chambers and Reisberg showed participants ambiguous pictures, such as the duck-rabbit picture, and asked them to rotate the image in their minds and provide the other interpretation of the picture. The participants were unable to do this, but they were able to draw the picture and discover the other interpretation by rotating it. Kosslyn (1994) questioned Chambers and Reisberg's methodology and reported Finke et al.'s (1989) strikingly different results. Finke et al. asked participants to close their eyes and juxtapose or superimpose mental images, to describe the result, and then to draw the pattern in their image. A well-known example is that participants were able to rotate a 'D' and add it to an upright 'J' to get an image of an umbrella. Pinker and Finke (1980) found that participants could imagine a three-dimensional scene as viewed from a novel perspective. Kosslyn (1994) cited

similar results by Anderson and Helstrup (1993); Brandimonte, Hitch, and Bishop (1992); Finke and Slayton (1988); and others. Rollins (1999) provided an alternative interpretation of Chambers and Reisberg's results that combines the role of representations with attentional and other strategies.

6.2. *Seeing as*

Suppose somebody asks you if Bill is ugly. You generate an image of him, scan it, and discover he is ugly. You have not thought about Bill's appearance before, so this is something new that you discover by inspecting an image. This knowledge is not obviously relational. You might decide that Bill is ugly because his eyes are too close together (which is a kind of relational knowledge), but you might equally decide that he is ugly because he has beady eyes and a big nose.

This is a case of what Wittgenstein (1963) called *seeing as*. You have seen Bill but have not seen him as ugly before. Perhaps you were too busy when you saw him, or perhaps it is not in your nature to judge people on their appearances. However, now that somebody asks you, you apply the concept of ugliness and discover that it fits.

These relational and *seeing as* scenarios are straightforward cases of acquiring knowledge by scanning or manipulating visual images, just as we can scan and manipulate things in the world. We acquire this knowledge by scanning the images on a need-to-know basis, just as, according to active externalism, we acquire knowledge about things in the world on a need-to-know basis. Kosslyn (1994) said that we acquire the knowledge by performing the same operations and using the same mechanisms that we use when we interpret objects during perception. We scan and zoom in on images, just as we saccade and foveate onto things in the world. We scan images either by shifting the attention window (as we might saccade around a room or scan a face) or by performing the inner analogue of turning our head or body so that we see something different but contiguous to what we saw before (Kosslyn, 1994, p. 367). We zoom in on an image 'until the sought part or characteristic is clearly visible' (Kosslyn, 1994, p. 369). Kosslyn (1994) assumed that the visual buffer has only limited resolution but he says that a part or characteristic will 'pop into view' when we foveate onto it as more detailed information is supplied by the encoding. In doing this, we use the same mechanisms that we use in perception.

If this is correct, the external–internal symmetry carries over, at least in some cases. We use objects and situations in the world as external memory stores and we use the same mechanisms to acquire knowledge from inner analogues. This knowledge is not explicitly stored. It is not clear to me whether it is implicitly stored or not stored at all. I prefer to say that it is not stored at all, but some might think that this is too paradoxical. This issue might be a red herring, however. When we infer something from what we already know, was the derived knowledge implicitly stored or not stored at all? The question probably strikes us as odd. What matters is that we should be able to explain how we acquired the knowledge. I suspect that we have the same situation with the kind of noninferential knowledge we are dealing with here. I explain how we acquire this knowledge later on.

The knowledge that we acquire by consulting inner analogues is computationally cheap. All that we need in the Bill example is a representation of Bill's face and the ability to apply

the concept of ugliness. It is true that we have the computational cost of applying the concept, but we can apply it on a need-to-know basis when we know exactly what information we require. In addition to this, we can scan the same representation for many things—ugliness, bad temper, boredom, fatigue—and we can do so for the same computational cost and using the same mechanisms as scanning Bill's face in the world, but now with the advantage of portability.

So far so good. However, this is only the tip of the iceberg.

6.3. *Concepts and content*

In the cases we have considered we acquire new knowledge by further processing what we already know. Having generated an image of a face or a coffee mug we can further process the images by applying new concepts to them—concepts such as ugliness and width compared to height. We apply the concepts to the images in much the same way that we apply them to faces and mugs in the world.

The idea of further processing the image probably leads to an intuition—that we store a minimally conceptualized image and apply what we might call 'second-order concepts' to it. These are concepts such as ugliness, fatigue, and boredom that (according to the intuition) we can only apply once we have a more or less complete image to apply them to. Such concepts stand in contrast to what we might call 'constitutive' concepts, such as shape and color (or complexion), that are used to generate the image in the first place.⁹ The intuition is that a minimally conceptualized image has to be stored. We might be able to retrieve the knowledge that Bill has blue eyes by scanning an image of Bill; but, unlike the knowledge that he is ugly, this knowledge has to be stored. So according to the intuition, the object–image symmetry breaks down here: We must have stored constitutive knowledge about Bill's face to generate an image of it that we can apply second-order concepts to. It seems that evolution has given us (a) the ability to use objects in the world as external memory stores, and (b) inner analogues that we can use with the same economies of storage; but (according to the intuition) nature did not finish the job. Constitutive knowledge has to be stored; therefore, we lose at least some of the economies of storage.

Let me try to counter this intuition with another one. To acquire information about the world we need high-level knowledge and the ability to zoom in and grab the details when we need them. Similarly with images and second-order concepts, all that we need is the 'lay of the land'¹⁰ and the ability to zoom in and grab the details when necessary. Two principles seem to be at work in both cases. The first is a principle of parsimony, which says, 'do as little processing as possible and leave it to the last minute.'¹¹ The second is Marr's (1982) 'principle of least commitment,' which says that we should not commit ourselves early on to a subset of possible solutions by discarding information that we cannot recover later. Kosslyn (1994) explicitly appealed to this principle, as I show later.

The two principles seem to work in consort. As much information as possible should be available to us (we should not lose information by prejudging a situation and committing ourselves to a subset of possible solutions), but we should only access it when we know what information we require.

Now we come to the point: If these principles are operative, and given that we do have portable memory stores, we would expect nature to have evolved a way of storing inner analogues that we do not conceptualize *at all* until we need to do so. We can acquire knowledge by applying second-order concepts to images and we can do so on a need-to-know basis. To complete the job (and the internalism–externalism symmetry), nature would have given us the ability to withhold *all* conceptualization until the last minute; enabling us to acquire knowledge on a need-to-know basis from mental imagery—just as we acquire it from external memory stores.

So far we have followed a hunch and seen how things would be if some rather vague principles were in place. However, the hunch (and the principles) is, at least implicitly, implemented in Kosslyn's (1994) model. In the model, conceptualization takes place in associative memory.¹² The compressed images stored in the pattern activation subsystems are preconceptual—they are not conceptualized until they are triggered by concepts in associative memory. In this sense they play the same role as objects and scenarios in the external world: They engender a pattern of activation in the visual buffer. The pattern in the visual buffer is processed in the same way, regardless of whether it is engendered by patterns stored in the pattern activation subsystems or through interaction with the world. Kosslyn (1994) said:

Once a configuration of activity exists in the visual buffer, input is sent to the ventral and dorsal systems and is processed in the usual ways—*regardless of whether the activity arose from immediate input from the eyes or from information stored in memory* [italics added]. This notion lies at the heart of the similarities found between imagery and perception. (p. 336)

Perception and image generation are not disjoint processes. When we perceive something, imagery typically augments the input to the visual buffer, which is activated by both external and internal stimuli; patterns stored in the pattern activation subsystems collaborate with input from the world to generate images in the visual buffer. When we generate an image, the buffer is only activated by internal stimuli: The stored patterns do all the work. As Kosslyn (personal communication, August 7, 2005) said, 'imagery is in most respects just like vision, except that there's nothing there.'

I suggest that these stored patterns supply us with *nonconceptual content*, just as our perception of things in the world supplies us with nonconceptual content.

7. Nonconceptual content

Philosophers sometimes talk about 'nonconceptual content' (Evans, 1982; McDowell, 1994; Peacocke, 1992). The basic idea is that our mental states can have content that is not adequately captured by the concepts we possess. The concept of nonconceptual content has application in several areas, such as the status of representational states in infants and nonhuman animals and the status of representational states implicit in our tacit knowledge of the rules of syntax. Its most obvious application, however, is to perception. Wright (2003) said, 'The most straightforward argument for embracing nonconceptual content is based on the fine-grained nature of experience. The rich detail of features presented in our experience outstrips the conceptual capabilities of even the most sophisticated creatures' (p. 39). This is

not to say that we cannot conceptualize hitherto unconceptualized experience. We can, and the process produces conceptual content:

The process of conceptualization or judgment takes the participant from one kind of state (with a content of a certain kind, namely nonconceptual content) to his being in another kind of cognitive state (with a content of a different kind, namely, conceptual content. (Evans, 1982, p. 227)

The point is simply that perception is too fine grained to be conceptualized in its entirety. We can discriminate more shades of color, shapes, textures, and sounds than we can conceptualize.

Nonconceptual perceptual experience makes a diversity of perceptually based concepts available to us—concepts such as *that shape, that color, diamond shaped, red, Cambridge blue, and yellow ochre*, to name but a few (Peacocke, 1999). It does this in two ways. First, the fine-grained nature of experience, in all its detail and diversity, supports a multitude of judgments. We say that a picture is worth a thousand words. Another way of saying this is to say that it supports a thousand judgments. Second, the same content can be conceptualized in different ways. We might hear a sound and think of it as coming from the south or from the football stadium; but Fido, who has a limited understanding of geography (and no understanding of football), might think of it as just ‘coming from over there.’¹³ Human infants would probably conceptualize it in the same way, and we can think of it that way ourselves, of course.

Now suppose that we store some of our nonconceptual content in the pattern activation subsystems (also known as visual memory; assuming, for simplicity’s sake, that the modality is vision). We copy it into visual memory early on, without conceptualizing it. This stored nonconceptual content now supports a diversity of judgments, just as our nonconceptual experience of things in the world supports a diversity of judgments. We can make these judgments in context, on the spot, at the right time, and on a need-to-know basis.

Martin (1992), building on Dretske’s (1969) notion of ‘nonepistemic seeing,’ told us about Mary. Mary is a keen board-games player who uses a 12-sided die but does not have the concept of a *dodecahedron*. (She does not like counting above 6 and treats all die with more than 6 faces as the same.) She later acquires the concept of a dodecahedron and realizes, in retrospect, that the die she has been using is a dodecahedron. This is a simple illustration of how we can acquire new knowledge by applying new concepts to stored, remembered content. In the same way, Fido might come to realize that the sound came from the south, if he acquires an understanding of the points of the compass.

I invite you to try the following exercise. Remember a rich experience, such as the wild flowers you saw when you were bushwalking in the hills, the taste of a fine wine, listening to the cello, or the sensation of someone running their fingers through your hair. Is it not the case that the remembered experience is underdetermined by the concepts you possess? Is it not the case that (as the saying goes) your memory is ‘rich beyond description’? We typically say, referring to a past experience, that we ‘cannot really (or cannot fully) describe it.’ This suggests that our memory of the experience contains fine-grained, nonconceptual content.

If this is the case, at least some of our memories are not conceptualized until we recall them. We might not conceptualize Bill’s face *as* a face or his eyes *as* blue until we conceptualize the stored nonconceptual content, and by doing so, generate memories. In one sense, this collapses the intuitive distinction between constitutive and second-order concepts. We do not

think of Bill as ugly until we apply the concept of ugliness to a generated image of him; and it may also be the case that we do not think of him as blue-eyed or brown-haired until we apply the appropriate concepts to the stored, nonconceptual content. In another sense, however, it *explains* the distinction. Constitutive concepts apply to nonconceptual content to give us an image of Bill's face, whereas second-order concepts apply to the image of the face once it has been generated.

The situation, however, is not that we *might* store nonconceptual content in memory but that we *must* do so according to Kosslyn's (1994) model. There is a general reason and a specific, architectural reason for this. The general reason has to do with Kosslyn's (1994) claim that the same mechanisms are at work in both imagery and perception. If the mechanism that is used in perception is redeployed to generate imagery, we would have to give it much the same input that it receives in the case of perception. That, after all, is what it does: It processes that kind of input. If your car runs on petrol you cannot give it something different to process, such as water or kerosene. If the mechanism processes perceptual input, the only thing we can give it to generate imagery is something very much like (a copy or close copy of) perceptual input—and this is what stored nonconceptual content is.

Now for the architectural reason. Kosslyn (1994) did not talk about nonconceptual content, but he did talk about representations stored as *compressed images* in the pattern activation subsystems. These representations are not concepts because conceptualization occurs later in associative memory. The pattern activation subsystems (category pattern activation and exemplar pattern activation) are downstream from the visual buffer, but upstream from associative memory. The representations in question are copied from input early on. Kosslyn (1994) called them 'visual memories' and 'visual representations of patterns' and said that they are stored in the anterior portion of the inferior temporal lobes (at least in the case of monkeys: Fujita, Tanaka, Ito, & Cheng, 1992). This region is not used to store other types of representations, such as those used to understand speech or guide movement. Kosslyn (1994) said that the 'long-term representation of an image' (p. 119) is probably stored as a feature vector across columns of neurons, with each neuron participating in many representations (Desimone & Ungerleider, 1989; Rolls, 1987; Young & Yamane, 1992).

Kosslyn (1994) provided an argument strikingly similar to the one I outlined earlier: Suppose that to distinguish between similar shapes we have to compare the properties of a stored image to the input. We cannot 'know in advance what properties will be useful' (p. 118). Consequently:

It is important to store information that will allow one to recover a wide range of properties; these properties need not be explicit in the representation, but must be able to be reconstructed from it. Ideally, this information will allow one to reconstruct the image itself: Not only are there an infinite number of propositions one can assert about an object (Goodman, 1983), but also one must note the spatial relations among these properties. (p. 118)

I agree. One of the advantages of storing analogues is that they support a multitude of judgments, which we can make on a need-to-know basis. Here Kosslyn (1994) said that there are an infinite number of judgments we can make about an object. What he meant, in this context, was that there are an infinite number of judgments we can make about an *image*.

In terms of the picture I am putting together, stored nonconceptual content makes an infinite number of judgments possible. Kosslyn (1994) also said that storing a representation that specifies enough information to reconstruct an image respects Marr's (1982) principle of least commitment. As we have seen, this principle says that we should not commit ourselves early on to a subset of possible solutions by discarding information that we cannot recover later.

Let us take our theory for a test drive. We want to know whether foxes have pointed ears. We have seen foxes or pictures of foxes, and patterns of activation (representations with nonconceptual content) have been copied across and stored in our pattern activation subsystems (i.e., visual memory). The concept *fox* activates these patterns to generate an image of a fox in the visual buffer. We vicariously experience, courtesy of the stored content, a reddish-brown creature with pointed muzzle and ears and a bushy tail slinking along low to the ground and looking at us with startled eyes (because the last time we saw one it was trapped in the headlights). We scan this image and retrieve the information we require. This simple story does not commit us to saying that we stored the information that foxes have pointed ears. What we stored was nonconceptual content, stored as patterns of activation. We bring concept and content together (in a very Kantian way, as seen in the next section) to generate an image of a fox in the visual buffer. We scan the image and discover that foxes have pointed ears. We could have scanned it for other information, such as the color of the fox's eyes or what kind of a tail it has.

8. Discussion

I began this article with a strange, almost paradoxical problem. Active externalism says that we use objects and situations in the world as external memory stores that we can consult as needs dictate. However, most objects and situations are not portable. We cannot carry them with us and consult them on a need-to-know basis. Given that we do have inner analogues we would expect them to give us economies of storage, as their external counterparts do. However, now the problem appears. If the information we retrieve was stored in memory, we have lost the economies of storage. Also, it is difficult to see how we can discover the information, for if it was stored in memory there is presumably a sense in which we knew it already. If, on the other hand, the information was not stored in memory, it is difficult to see how we could have discovered it, for it was not there to be discovered.

We can now answer these questions. With relational knowledge and *seeing as* we acquire new knowledge by further processing what we already know. Having, for example, generated an image of a face we can bring new concepts to it—concepts such as ugliness and intelligence. We apply these concepts to the image in much the same way that we apply them to faces in the world. We interrogate objects and images in similar ways—on a need-to-know basis, using the same mechanisms and performing the same saccade-and-foveate operations. Both cases give us economies of storage. If Bill is in front of us we can look at him to see if he is ugly, so that we do not need to store this information. If he is not in front of us, we can consult an inner analogue, performing the same operations and applying the same concepts that we would apply if we were looking at him in the world—but now with the advantage of portability. We have not stored this knowledge. We acquire it by applying concepts to the generated image.

This gives us a partial symmetry between external objects and inner analogues, and to some extent vindicates the intuition that inner analogues function in the same way as external memory stores. However, the symmetry seems to be incomplete, because it *seems* that we need to store knowledge about things in the world to construct images that we can apply second-order concepts to. It seems, for example, that we need to store information about the shape of Bill's face and the color of his eyes if we are to construct an image that we can scan for ugliness and intelligence.

I began by suggesting that this might not be the case. There are *a priori* reasons for believing that we store nonconceptual content in memory, because neither our perception nor our memory of what we perceive can be adequately captured by the concepts we possess. However, the case is stronger than this. According to Kosslyn's (1994) model, representations in visual memory must be nonconceptual both because of their location in the cognitive architecture and because of the role they play in cognition. In terms of the cognitive architecture, visual memory (i.e., the pattern activation subsystems) is situated between the visual buffer and associative memory, so that representations are stored in visual memory before they are conceptualized. There is good reason for this. Conceptualizing them early would prejudice solutions to problems we have not yet encountered. Storing compressed images in a format similar to the original input enables us to conceptualize them in multiple ways, on a need-to-know basis, taking advantage of the richness and diversity of the input.¹⁴

If this is correct, the inner–outer symmetry is complete. Stored representations supply us with content, just as external objects do. We have economies of storage in both cases, because we do not acquire the knowledge until we conceptualize the content. When we scan Bill's face in the world we are not tempted to say that the information we acquire was stored in our heads. It is the same with the image. We scan it and discover that Bill is ugly. This information was no more stored in our heads than it was when we scanned Bill's face in the world. We acquired it by applying the concept to the stored content and scanning the resultant image.

In terms of economies of storage, it is true that we have to store the nonconceptual content. However, we need to see the big picture here. According to Kosslyn (1994), imagery augments and fills in the gaps in perception so that when we look at Bill's face in the world, stored imagery is already at work. Consequently, storing nonconceptual content in memory for the sake of subsequent reflection comes at little or no additional cost to storing it to facilitate perception. Kosslyn (1994) also said that we use the same mechanisms for imagery that we use for perception, so we get these for free as well. The evolutionary story implicit in Kosslyn's (1994) account is that we developed repositories of compressed images to aid perception and learned to deploy them offline, using in-place perceptual mechanisms, so that we could remember, plan, and anticipate in the safety of the cave. If this story is correct, imagery comes at little or no additional cost, in terms of either processing or storage.

I now look at some issues and some possible objections.

8.1. Knowledge and content

At the end of the previous section I said that we bring concept and content together 'in a very Kantian way' to generate an image of a fox in the visual buffer. Kant (1964) famously said that concepts without percepts are empty; percepts without concepts are blind. Knowledge for Kant

was conceptualized content—content that is originally derived from sensory input. Treating knowledge as conceptualized content is key to my solution to the problem. We can acquire knowledge by scanning inner images that we have generated by conceptualizing stored nonconceptual content. We did not store this knowledge. We stored the content and acquired the knowledge when we scanned the image that was generated by conceptualizing the content. In this context, the stored nonconceptual content plays a similar role to sensory input from the world.

Kosslyn (1994) said a compressed image is probably stored as a feature vector across columns of neurons, with each neuron participating in many representations (p. 119). I assume that this gives us what connectionists call ‘superpositional storage,’ where the same artificial neurons participate in multiple representations. Cats, dogs, and possums, for example, have overlapping microfeatures that can be instantiated by the same set of neurons. This gives us economies of storage. We get additional economies of storage because a single compressed image can be conceptualized in different ways, on a need-to-know basis, depending on our requirements at the time. Arthur Markman (editorial comments, April 21, 2006) captured some of this picture when he said, ‘the processes and storage of information responsible for classifying images are separate from the processes and storage of information responsible for aspects of the analysis of images.’

An anonymous referee was sceptical, however, and suggested that when we generate an image of a face, ‘it would seem to be the case that (at least implicit) information about the face shape is stored. The shape seems like a basic feature, which, though not labelled, is a first-order characteristic.’

Let me try to loosen the grip of this intuition in another way. Suppose we look at the ambiguous old woman/young woman picture for the first time and store it in visual memory before we have seen the old woman interpretation of it. There is no sense, as far as I can see, in which we now know the shape of the old woman’s face. We have not seen her face; we have not seen this interpretation of the picture. We can, however, acquire this knowledge by retrieving the picture and conceptualizing it as the face of an old woman. It does not follow from this, of course, that we have stored the image as nonconceptual content. I am only saying that we do not know the shape of the old woman’s face until we have conceptualized the stored image in a certain way. We can extend this to nonconceptualized content as follows. We look at Bill and store a compressed image of his face as a feature vector across columns of neurons. The vector is tagged so that different aspects of it can be activated by associative memory. Later, when we are talking or thinking about Bill, we activate the compressed image and discover that he is ugly—and we discover ‘first-order’ things about him as well, such as the shape and color of his face. We acquire this knowledge by conceptualizing the stored nonconceptual content.

Markman raised another possible objection. Some people, he said, might cite the change-blindness literature as showing that ‘much less information about the perceptual world may enter memory than we perceive in the moment as we look around the world’ (editorial comments, April 21, 2006). He added, ‘It seems to me that it is still likely that we store much more perceptual information than is analyzed conceptually, and so your main point still holds’ (editorial comments, April 21, 2006). I agree with him on both counts. Activist externalists (e.g., Clark, 2003) like to cite the change-blindness literature as evidence that we extract and process less information about the world than we might expect. However, ironically, they

also cite the Titchener/Aglioti circles illusion (see section 2 of this article) as evidence that we have subsystems that get on with the job of interacting with the world without reporting back to consciousness. In this case, more information is processed through the dorsal stream than is presented to consciousness through the ventral stream; therefore, in this case, *more* information about the perceptual world is processed than we consciously perceive in the moment.

8.2. *Knowing when and what to conceptualize*

Knowing when and what to conceptualize seems to be a fine art. For example, assume someone with autism or Asperger's disorder goes into a supermarket and is overwhelmed because he or she tries to conceptualize, categorize, and remember everything all at once and in great detail. (For a study of someone who is overwhelmed in this way, see Haddon, 2003.) Such people are unable to effectively deploy an attention window that they can selectively move around the supermarket. They cannot maintain a high-level map and 'conceptually foveate' on a need-to-know basis—and everything crowds in on them. Most of us maintain a partial, high-level map of where things are. If we need help, we look at the signs and labels at the end of the aisles. This might be seen as illustrating the importance of maintaining a 'separation of powers' between perception and cognition that ensures that we do not conceptualize too much or too soon. Presumably, this applies to stored content as much as it does to external perception, because we have stored a vast number of compressed images from different angles, under different light conditions, and so forth. Without a separation of powers we would be overwhelmed by our memories, just as an autistic person is overwhelmed by the blooming, buzzing confusion in the supermarket. Haugeland (1998) said, 'It would be silly, for most purposes, to try to *keep track* of what shelf everything in the refrigerator is currently on; if and when you want something, just *look*' (p. 219).

8.3. *Nonconceptual content versus less determinate concepts*

Some people have argued that the fine-grained nature of perceptual experience can be accommodated at the conceptual level. John McDowell (1994), for example, suggested that the conceptual content of perceptual experience is given by demonstrative concepts such as *that shade* and *that color*. According to McDowell, the information loss in the transition from perception to perceptual belief is not a sign that there are two types of content in play, but is due to the transition from a more determinate to a less determinate type of conceptual content. Consequently, we do not need the concept of nonconceptual content.

Marr's (1982) principle of least commitment, together with considerations of symmetry and parsimony and Kosslyn's (1994) detailed architecture, all gently militate against McDowell's (1994) position. There is room for reciprocity and feedback in Kosslyn's (1994) model so that, in principle, perceptual input could be conceptualized early on—but the fact remains that visual memory is not far downstream from the visual buffer, so that it is most naturally seen as being not conceptualized at all (even by concepts such as *that color* and *that shade*). As for Marr's principle of least commitment, demonstrative concepts

commit us early on (albeit in an attenuated way) to a subset of possible solutions more than nonconceptual content does.

8.4. *Dispositional knowledge*

Someone might take a dispositional stance toward knowledge and say that we know something if we can derive it through generating and scanning. It is not that we know it *once* we have generated and scanned it. We know it already *because* we can derive it by generating and scanning. We implicitly know that foxes have pointed ears because we can generate an image and scan it for the explicit knowledge we require. Put differently, we know that foxes have pointed ears because we can represent this to ourselves without having to look in the world. In this sense, knowledge is dispositional.

I am happy to concede that we know these things in this dispositional sense. Clark (2003) argued that we know something if we have quick and easy access to it, so that we know what the time is if we can easily look at our watch. (If someone asks us if we know the time, we say that we do and *then* we look at our watch.) In this dispositional sense we know that foxes have pointed ears and that Bill has blue eyes before we consult our images of them. I am happy to concede this, because it underscores the symmetry between internalism and externalism: In this dispositional sense, we know what the time is before we look at our watch; and in the same dispositional sense, we know that foxes have pointed ears and that Bill has blue eyes before we consult our images of them.

8.5. *Passive structures versus active expectations*

Kosslyn (1994) repeatedly said that, despite the complexity of his model, it almost certainly fails to do justice to the complexity of the brain. One way in which it fails to do justice to the complexity of the brain is to treat the contents of visual memory as passive structures unrelated to motor imagery, motor control, and our interaction with the world. Rick Grush (2004) said, 'Surely, to treat perception, imagery, and motor control as functionally distinct modules, as though any of them could do its job without the others, is to significantly distort the genuine neurophysiological phenomena' (p. 393).¹⁵ Grush (2004) thought that he could synthesize 'at least the main aspects' of Kosslyn's (1994) 'very detailed' theory into his own framework to provide 'a yet broader account of the brain's information-processing structure' (p. 392). Grush's (2004) *emulation theory of representation* says that, for reasons of motor efficiency, we have developed online emulators that enable us to anticipate and emulate the world. During sensorimotor engagement, these models are driven by efference copies of motor commands, operating in parallel with the body and environment, to provide expectations of sensory feedback and to enhance and process sensory information. Perception involves 'a content-rich emulator-provided expectation that is corrected by sensation' (Grush, 2004, p. 394). In evolutionary time we learned to run these emulators offline 'to produce imagery, estimate outcomes of different actions, and evaluate and develop motor plans' (Grush, 2004, p. 377).

Grush's (2004) theory spoke of motor efficiency and sensorimotor engagement, but in many ways it was strikingly similar to Kosslyn's (1994) account. Perception is inherently

anticipatory and driven by content-rich, emulator-provided expectations that are corrected by sensation. The emulation framework explains ‘exactly how it is that the same process that can operate off-line as imagery, can operate on-line as the provider of expectations that fill in and interpret bare sensory information’ (Grush, 2004, p. 392). The emulation framework’s explanation of imagery as the driving of emulators by efferent motor areas ‘predicts what Kosslyn’s theory posits, that in many cases the imagery that is used to aid in perception is the product, in part, of the activity of motor areas’ (Grush, 2004, p. 392). When we manually rotate the jigsaw piece, we run the emulator online to anticipate, fill in, and enhance sensory feedback. When we mentally rotate it (having left the room), we run the emulator offline, using efference copies of motor commands. The same motor emulation is at work, and it gives us content that is similar to the content it gave us in the presence of the puzzle, but now without any input or feedback from the world: We rotate the jigsaw piece in our minds—but now there is nothing there. In so doing, we perform an operation in our heads that we would normally perform in the world and make an empirical discovery in both cases. I say this in my commentary on Grush’s (2004) article (Dartnall, 2004a). In his reply, Grush (2004) said, ‘I agree entirely with this suggestion, and in fact in Grush (2003) I discuss this briefly’ (p. 430).

Grush’s (2004) framework points to the same evolutionary story as Kosslyn’s (1994). We began with perception being anticipatory and driven by content-rich, emulator-provided expectations. Survival and natural selection pressured us to plan and anticipate in the safety of the cave—or at least a long way from the sabre-toothed tiger—and nature deployed the only tools that were available: motor emulators, which we learned to run offline to produce imagery, estimate outcomes of actions, and evaluate and develop motor plans.

9. Conclusion

We began this article by looking at a partial symmetry between active externalism and internalism: Active externalism says that we can perform cognitive processes in the world, and internalism says that we can perform operations in our minds that we would normally perform in the world. Active externalism, however, also says that we use objects and states of affairs as external memory stores that we can consult as needs dictate. If the symmetry carries over, we will have portable inner analogues of the world that we can consult as needs dictate. In this article we have looked at a major problem for this suggestion and have seen not only that the problem can be solved but that the solution suggests that nonconceptual content is stored and conceptualized in the brain. This completes the symmetry: We use the world as a memory store that we consult on a need-to-know basis, and we do much the same with inner analogues. The big picture is that the mind leaks into the world and the world leaks into the mind—and they do so in comparable ways.

Notes

1. Stephen Kosslyn (referee’s report, April 21, 2006) pointed out that working with mental representations has other advantages. It is easier, for example, to imagine moving a sofa around the room rather than actually moving it.

2. Kosslyn's 1994 model is summarized in Kosslyn, Thompson, and Ganis (2006), along with the standard arguments against depictive representations, the standard replies to such arguments, and a wealth of empirical findings. I refer mainly to the Kosslyn 1994 in this article because the 1994 account is more detailed.
3. Shepard and associates (Cooper & Shepard, 1973; Shepard & Metzler, 1971) showed in a series of classic experiments conducted in the early 1970s that we can perform such operations. Finke, Pinker, and Farah (1989) showed that we can rotate objects in our minds and assign a different meaning to them on the basis of the rotation. For example, we can rotate the letter 'D,' add it to an upright 'J,' and realize that we have an image of an umbrella.
4. Guy van Orden (referee's report, April 21, 2006) pointed out that some memory stores are portable, such as wristwatches, wallets, notebooks, and computers. I think he would agree that these are the exception rather than the rule. We cannot carry faces, bodies, furniture, and the medium-sized objects that make up our world around with us, tucked under our arms. Yet all of these things, according to active externalism, can serve as external memory stores. Also, the portable objects that van Orden mentioned have appeared very recently on the evolutionary scene and were not available to our early ancestors.
5. Hebb (1968) said we experience 'apparent' or 'imagined' objects rather than images:

... one is not describing the image but the apparent object One does not perceive one's perceptions, nor describe them; one describes the object that is perceived, from which one may draw inferences about the nature of the perceptual process. In the case of imagery, one knows that the apparent object does not exist, and so it is natural to think that it must be the image that one perceives and describes, but this is unwarranted.' (pp. 467–468)

Kosslyn (referee's report, April 21, 2006) pointed out that people can reliably rate the vividness of the image itself (not the vividness of the object) and estimate the visual angle within which objects can appear (treating the image itself as an entity to be evaluated). In both cases, he said that it is the image and not the object that is being addressed.

6. Kosslyn, Behrmann, and Jeannerod (1995) said:

In some ways imagery may have the same relation to other processes that glasses have to the nose: Once certain properties of a nose were present (presumably to warm air, direct scents to different nostrils, and so on), they could be used to hold up glasses. Similarly, once specific properties of basic sensory and response mechanisms were in place, they could be recruited into imagery—which then in turn may have become a function in its own right, thereafter being the subject of natural selection. (p. 6)

Kosslyn (personal correspondence, October 6, 2005) said that he wrote this passage and wonders why he did not refer to Stephen Jay Gould's (1991) concept of 'exaptation.' Exaptation is the co-opting of previously evolved functions to do new things (Gould, 1991; Gould & Vrba, 1982).

7. Kosslyn and Thompson (2003) addressed the problem that many neuroimaging studies have not shown activation in early visual cortex (Areas 17 or 18). They provided an overview and analysis that 'makes sense of what at first glance might appear a chaotic set of empirical findings' (p. 740). Mazard, Tzourio-Mazoyer, Crivello, Mazoyer, and Mellet

(2004) found that the topographically organized visual cortex was activated when participants performed depictive imagery tasks, but was deactivated when they performed spatial imagery tasks. They also found that the activity of the early visual cortex during some imagery tasks varied among participants and said that this variability should be taken into account in explaining the divergent results found in previous studies.

8. Kosslyn (personal communication, October 6, 2005) pointed out that images are not entirely picture-like, principally because they lack resolution.
9. Kosslyn (1994), following Marr (1982), wrote about ‘nonaccidental properties’ (e.g., pp. 108–114), which loosely correspond to my ‘constitutive properties.’ These are identified and extracted early on by the *preprocessing subsystem*. Kosslyn (personal correspondence, October 6, 2005) said he believes the notion originated with David Lowe.
10. The expression is Kosslyn’s, although it might well have been Andy Clark’s. Kosslyn (1994) said, ‘I assume that ‘global precedence’ is useful in both image generation and perceptual encoding, for the same reason: In both situations, having the ‘lay of the land’ helps one to organize the details’ (p. 292).
11. The second part of the principle might be unnecessary, because leaving things to the last minute might be the most efficient course of action. It is only then that we know what information we require: ‘Where *exactly* is the sabre-toothed tiger?’ ‘Where *exactly* are my fellow hunters?’
12. Kosslyn (1994) said: Associative memory contains not only associations between perceptual representations, but also more abstract ‘conceptual’ information (names, categories, parts of speech, and so forth). Information in associative memory can be accessed by input from all sensory systems; once one has accessed the appropriate information, the object has been identified. (p. 73)
13. Peacocke (1999) said, ‘Some of the nonconceptual content of our experience can be identical with the representational content of the experience of creatures that either possess no concepts, or possess only a set of concepts far more rudimentary than our own’ (p. 5). In terms of Kosslyn’s (1994) model we might speculate that such creatures have limited associative memories.
14. Strictly speaking, as Guy van Orden pointed out in his referee’s report, I am not committed to many of the details of Kosslyn’s (1994) model: My theory of nonconceptual memory might be correct even if the details of Kosslyn’s (1994) model turn out to be false. While this is true, I do think that the details matter. It is not a coincidence, for example, that Kosslyn’s (1994) model requires nonconceptual memory. I think this shows that his model, and my account of nonconceptual memory, are on the right track.
15. Kosslyn (personal correspondence, October 6, 2005) said he agreed with Grush on this point and said he had addressed it in some of his own publications (Kosslyn, Thompson, Wraga, & Alpert, 2001; Wraga, Thompson, Alpert, & Kosslyn, 2003).

Acknowledgments

I am doubly indebted to Stephen Kosslyn, first for his comments on an earlier version of this article and for a helpful discussion, and then for suggesting more improvements

when he refereed the article. Kosslyn has independently developed his own version of active externalism, which focuses on what he calls ‘social prosthetic systems.’ These are ‘human relationships that extend one’s emotional or cognitive capabilities’ (Kosslyn, 2006) so that other people ‘serve as prosthetic devices, filling in for lacks in an individual’s cognitive or emotional abilities.’ The self, Kosslyn (2006) said, ‘becomes distributed over other people who function as long-term social prosthetic devices.’ I had the feeling when I was reading the article that this was active externalism seen through the eyes of the later Wittgenstein.

I also thank Guy van Orden, Arthur Markman, and an anonymous referee for their constructive criticism and support.

References

- Aglioti, S., DeSouza, J. F., & Goodale, M. A. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology*, 5, 679–685.
- Anderson, R. E., & Helstrup, T. (1993). Visual discovery in mind and on paper. *Memory and Cognition*, 21, 283–293.
- Brandimonte, M. A., Hitch, G. J., & Bishop, D. V. M. (1992). Influence of short-term memory codes on visual image processing: Evidence from image transformation tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 157–165.
- Clark, A. (1997). *Being there: putting brain, body and world together again*. Cambridge, MA: MIT Bradford.
- Clark, A. (2003). *Natural-born cyborgs*. Oxford, England: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Cooper, L. A., & Shepard, R. N. (1973). Chronometric studies of the rotation of mental images. In W. G. Chase (Ed.), *Visual information processing* (pp. 75–176). New York: Academic.
- Dartnall, T. H. (2004a). Epistemology, emulators and extended minds. *Behavioral and Brain Sciences*, 27, 401–402.
- Dartnall, T. H. (2004b). We have always been . . . cyborgs. *Metascience*, 13, 139–148.
- Dartnall, T. H. (2005). Does the world leak into the mind? Active externalism, ‘internalism’ and epistemology. *Cognitive Science*, 29, 135–143.
- Dennett, D. (1996). *Kinds of minds*. New York: Basic Books.
- Desimone, R., & Ungerleider, L. G. (1989). Neural mechanisms of visual processing in monkeys. In F. Boller & J. Grafman (Eds.), *Handbook of neuropsychology* (pp. 267–299). Amsterdam: Elsevier.
- Donald, M. (1991). *Origins of the modern mind*. Cambridge, MA: Harvard University Press.
- Dretske, F. (1969). *Seeing and knowing*. Chicago: University of Chicago Press.
- Evans, G. (1982). *The varieties of reference*. Oxford: Oxford University Press.
- Finke, R. A., Pinker, S., & Farah, M. (1989). Reinterpreting visual patterns in mental imagery. *Cognitive Science*, 13, 51–78.
- Finke, R. A., & Slayton, K. (1988). Explorations of creative visual synthesis in mental imagery. *Memory and Cognition*, 16, 252–257.
- Fox, P. T., Mintun, M. A., Raichle, M. E., Miezin, F. M., Allman, J. M., & Van Essen, D. C. (1986). Mapping human visual cortex with positron emission tomography. *Nature*, 323, 806–809.
- Fujita, I., Tanaka, K., Ito, M., & Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex. *Nature*, 360, 343–346.
- Gazzaniga, M. (1998). *The mind’s past*. Berkeley: University of California Press.
- Goodale, M. A., & Milner, A. D. (2004). *Sight unseen: An exploration of conscious and unconscious vision*. Oxford, England: Oxford University Press.
- Goodman, N. (1983). *Fact, fiction, and forecast* (4th ed.). Cambridge, MA: Harvard University Press.
- Gould, S. J. (1991). Exaptation: A crucial tool for evolutionary psychology. *Journal of Social Issues*, 47, 43–65.
- Gould, S. J., & Vrba, E. S. (1986). Exaptation: A missing term in the science of form. *Paleobiology*, 8(1), 4–15.

- Grush, R. (2003). In defense of some 'Cartesian' assumptions concerning the brain and its operation. *Biology and Philosophy*, 18, 53–93.
- Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27, 377–442.
- Haddon, M. (2003). *The curious case of the dog in the night-time*. London: Doubleday.
- Haugeland, J. (1998). Mind embodied and embedded. In J. Haugeland (Ed.), *Having thought: Essays in the metaphysics of mind* (pp. 207–237). Cambridge, MA: Harvard University Press.
- Hebb, D. O. (1968). Concerning imagery. *Psychological Review*, 75, 466–477.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Ishiguro, H. (1967). Imagination. *Proceedings of the Aristotelian Society*, 14(Suppl.), 37–56.
- Kant, I. (1964). *Critique of Pure Reason* (N. Kemp Smith, Trans.). London, England: Macmillan. (Original work published in 1781).
- Klein, I., Dubois, J. M., Kherif, F., Flandin, G., Poline, J. B., Denis, M., et al. (2004). Retinotopic organization of visual mental images as revealed by functional magnetic resonance imaging. *Cognitive Brain Research*, 22(1), 26–31.
- Kosslyn, S. M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.
- Kosslyn, S. M. (1994). *Image and brain: The resolution of the imagery debate*. Cambridge, MA: MIT Press.
- Kosslyn, S. M. (2006). On the evolution of human motivation: The role of social prosthetic systems. In S. M. Platek, J. P. Keenan, & T. K. Shackelford (Eds.), *Evolutionary cognitive neuroscience* (pp. 541–554). Cambridge, MA: MIT Press.
- Kosslyn, S. M., Behremann, M., & Jeannerod, M. (1995). *The neuropsychology of mental imagery*. Oxford, England: Elsevier/Pergamon.
- Kosslyn, S. M., & Thompson, W. L. (2003). When is early visual cortex activated during visual mental imagery? *Psychological Bulletin*, 129, 723–746.
- Kosslyn, S. M., & Thompson, W. L., & Ganis, G. (2006). *The case for mental imagery*. Oxford: Oxford University Press.
- Kosslyn, S. M., Thompson, W. L., Wraga, M. J., & Alpert, N. M. (2001). Imagining rotation by endogenous and exogenous forces: Distinct neural mechanisms for different strategies. *Neuroreport*, 12, 2519–2525.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: Freeman.
- Martin, M. G. F. (1992). Perception, concepts, and memory. *The Philosophical Review*, 10, 745–763.
- Mazard, A., Tzourio-Maoyer, N., Crivello, F., Mazoyer, B., & Mellet, E. (2004). A PET meta-analysis of object and spatial mental imagery. *European Journal of Cognitive Psychology*, 16, 673–695.
- McDowell, J. (1994). *Mind and world*. Cambridge, MA: Harvard University Press.
- Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action*. Oxford, England: Oxford University Press.
- Neisser, U. (1976). *Cognition and reality*. San Francisco: Freeman.
- Peacocke, C. (1992). *A study of concepts*. Cambridge MA: MIT Press.
- Peacocke, C. (1999, May). *Does perception have nonconceptual content?* Talk given at the Certosa di Pontignano, Sienna, XXCountryXX.
- Pinker, S., & Finke, R. A. (1980). Emergent two-dimensional patterns in images rotated in depth. *Journal of Experimental Psychology: Human Perception and Performance*, 6, 244–264.
- Pylyshyn, Z. W. (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin*, 80, 1–24.
- Pylyshyn, Z. W. (1978). Imagery and artificial intelligence. *Minnesota Studies in the Philosophy of Science*, 9, 19–55.
- Pylyshyn, Z. W. (1981). The imagery debate: Analogue media versus tacit knowledge. *Psychological Review*, 88, 16–45.
- Pylyshyn, Z. W. (2002). Mental imagery: In search of a theory. *Behavioral and Brain Sciences*, 25, 157–238.
- Pylyshyn, Z. W. (2003). Return of the mental image: Are there really pictures in the brain? *Trends in Cognitive Sciences*, 7, 113–118.

- Rollins, M. (1999). Pictorial representation: When cognitive science meets aesthetics. *Philosophical Psychology*, 12, 387–413.
- Rolls, E. T. (1987). Information representation, processing, and storage in the brain: Analysis at the single neuron level. In J.-P. Changeux & M. Konishi (Eds.), *The neural and molecular bases of learning* (pp. 503–539). New York: Wiley.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701–703.
- Slotnick, S. D., Thompson, W. L., & Kosslyn, S. M. (2005). Visual mental imagery induces retinotopically organized activation of early visual areas. *Cerebral Cortex*, 15, 1570–1583.
- Stark, L., & Ellis, S. R. (1981). Scanpaths revisited: Cognitive models direct active looking. In D. F. Fisher, R. A. Monty, & J. W. Senders (Eds.), *Eye movements: Cognition and visual perception* (pp. 193–226). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Sutton, J. (in press). Exograms and interdisciplinarity: history, the extended mind, and the civilising process. In R. Menary (Ed.), *The extended mind*. Aldershot, England: Ashgate.
- Thomas, N. (1994). Review of Michael Tye: The imagery debate. *The Journal of Mind and Behavior*, 15, 291–294.
- Thomas, N. (1999). Are theories of imagery theories of imagination? *Cognitive Science*, 23, 207–245.
- Tootell, R. B. H., Silverman, M. S., Switkes, E., & De Valois, R. L. (1982). Deoxyglucose analysis of retinotopic organization in primate striate cortex. *Science*, 218, 902–904.
- Wittgenstein, L. (1963). *Philosophical investigations* (G. E. M. Anscombe, Trans.). Oxford, England: Blackwell. (Original work published 1953).
- Wraga, M. J., Thompson, W. L., Alpert, N. M., & Kosslyn, S. M. (2003). Implicit transfer of motor strategies in mental rotation. *Brain and Cognition*, 52, 135–143.
- Wright, W. (2003). McDowell, demonstrative concepts, and nonconceptual representational content. *Disputatio*, 14, 39–55.
- Young, M. P., & Yamane, S. (1992). Sparse population coding of faces in the inferotemporal cortex. *Science*, 256, 1327–1331.