



Attention as a patchwork concept

Henry Taylor¹ 

Received: 27 January 2023 / Accepted: 29 June 2023
© The Author(s) 2023

Abstract

This paper examines attention as a scientific concept, and argues that it has a patchwork structure. On this view, the concept of attention takes on different meanings, depending on the scientific context. I argue that these different meanings vary systematically along four dimensions, as a result of the epistemic goals of the scientific programme in question and the constraints imposed by the scientific context. Based on this, I argue that attention is a *general reasoning strategy* concept: it provides general, non-specific guidance that aids scientific work. As well as shedding light on attention as a scientific concept, this theory has two wider consequences. First, it demonstrates that the patchwork approach is a fruitful way to think about psychological concepts. Second, it provides novel resources to resist eliminativism about attention.

Keywords Attention · Psychology · Patchwork concepts · Eliminativism

1 Attention as a scientific concept

Attention is one of the central concepts in modern cognitive science, and has been deployed in accounting for psychological phenomena as diverse as consciousness (Sergent et al., 2013), the perceived properties of external objects (Carrasco et al., 2004), visual binding (Treisman, 2006), and much else (see Carrasco, 2011 and Chun et al., 2011 for surveys). William Uttal lists 47 different roles that attention has been put to (2011, p.231).

Philosophical work on attention has mostly focussed on two inter-connected projects (cf. Dicey Jennings, 2020, p.199). The first project attempts to answer the question ‘what *is* attention?’. The proposed answers are diverse. They include the neural amplification of input signals (Fazekas & Nanay, 2021); subject-level mental selection (Dicey Jennings, 2012, 2020) and cognitive unison (Mole, 2010). The second project aims to understand the relation between attention and faculties such

✉ Henry Taylor
j.h.taylor.1@bham.ac.uk

¹ University of Birmingham, ERI Building, Edgbaston, B15 2TT Birmingham, UK

as demonstrative reference (Dickie, 2015), consciousness (Prettyman, 2018; Taylor & Sayim, 2018; Watzl, 2017), perceptual justification (Silins & Siegel, 2013), social sensitivity (Magrì, 2020), and so on.

This paper will draw on the philosophy of science to examine attention *as a scientific concept*. Rather than asking what attention *is*, this paper will ask how the *concept* attention is used in scientific research. Attention as a scientific concept is often discussed alongside the metaphysical question of what attention *is*. Indeed, philosophers who present a unifying metaphysics of attention often do so partly with the aim of reinforcing its unity as a scientific concept (see, e.g. Wu, 2014, pp.5–6; Dacey Jennings, 2020, p.3 and Mole, 2010, vii).¹ This paper is different in analysing attention as a scientific concept independently of debates about its metaphysics. This paper's central aim is not to give a theory of what attention *is*, but to analyse *the way that the concept is used* in scientific research. I will focus on how the concept is structured, and how this structure contributes to the success of psychological research that uses it. I will focus on two questions. First, given that the concept is deployed in so many theoretical contexts in psychology, how do these shifts in context affect the concept, and why is this important for understanding how the concept aids scientific work? Second, what light can the philosophy of science shed on whether attention is still a useful scientific concept? To answer these questions, I will draw on the patchwork approach to scientific concepts (Cartwright, 1994; Wilson, 2006, 2017).

The results of this paper are important for philosophers of science. Such philosophers have long been interested in how scientific concepts contribute to scientific success, how they shift in meaning to accommodate novel scientific contexts, and the conditions under which they should be eliminated (e.g. Irvine, 2013). The present paper will draw upon the patchwork approach to provide an analysis of these issues with relation to attention. This provides novel support for the patchwork approach generally (by showing its application to attention). Historically, the patchwork approach has been most at home in materials science, and (as we will see) has been further extended to biology and neuroscience. This paper will be the first to apply the approach to *psychology*. Showing that the patchwork approach can help us understand psychological concepts further demonstrates its power, which in turn provides additional reason to accept the patchwork approach itself.

The paper is also important for thinkers primarily interested in attention. Any thinker interested in attention should also be interested in the scientific *concept* of attention. In pursuit of this, this paper will identify several subconcepts of attention, and explore the key features in virtue of which they aid scientific work. As we will see, the patchwork approach is essential to uncovering and explicating these subconcepts. The approach is also crucial for understanding how these subconcepts are impacted by features of a scientific context like epistemic goals and constraints. The paper will show that the concept of attention is more complex than we have so far realised. It will also show that this complexity stems from systematic application

¹ Wu (2014) discusses attention's metaphysics, its relation to other faculties, and the scientific concept of attention. Thanks to an anonymous referee for drawing my attention to this work.

of certain core principles to a range of scientific contexts. Navigating this mixture of complexity with systematicity is one of the main theoretical roles that the patchwork approach is uniquely well positioned for, as we will see. Finally (as I argue below), the application of the patchwork approach generates a novel response to accusations that attention is an outdated and scientifically useless concept.

Hereon, I use capitals for concepts. ATTENTION is a concept, whilst attention is the faculty. I take concepts to be mental representations, which can be used to pick out groups of entities. No particular theory of concepts is assumed. A concept's *extension* is just the set of things that the concept correctly applies to. So the extension of DOG is the set of dogs. I refer to the individual meanings that ATTENTION takes on in particular contexts as *subconcepts*, and to ATTENTION itself as the *overarching* concept.

In Section 2 I further motivate the project by outlining the threat of eliminativism. I then explain the patchwork approach to scientific concepts (Section 3) and develop a theory according to which ATTENTION is a patchwork concept (Section 4). According to this view, ATTENTION has multiple different meanings, which vary dependent on scientific context. I argue that these different meanings vary systematically in four ways: the scale they operate at, the measurement/manipulation technique of the paradigm that they are embedded in, the properties that something must have to fall into the extension of the subconcept, and the extension of the subconcept. I argue that the variations between these subconcepts are systematic, and should be explained by two factors: the different *epistemic goals* and *constraints* that operate in distinct scientific contexts. In Section 5, I argue that the overarching ATTENTION concept is a general reasoning strategy concept: it provides general guidance to scientists on how to manipulate the presence and absence of attention. I show that this result can be used to resist eliminativism about ATTENTION. Section 6 concludes.

I confine my discussion to the scientific (as opposed to folk) concept ATTENTION. In bracketing off the folk concept, I do not claim that the folk concept is unconnected to the scientific concept, or that the scientific concept is privileged. It could be that the folk concept derives its meaning from the scientific concept, in a way that is plausibly the case with neologistic concepts like ARTHRITIS (Burge, 1979). Alternatively, quite the reverse might be true: that the scientific concept derives its meaning from the application of a *folk* concept to a scientific context. In support of this latter view, note that *attention* is typically introduced, even in scientific papers, by appeal to the reader's folk intuitions (e.g. Carrasco, 2011, p.1484). Yet a third view is that the folk use and scientific use are divergent, carving out different but overlapping extensions, in a way arguably characteristic of concepts like WATER and FISH. (Malt, 1994, Dupré 1993).² However, the position of this paper is that there is much to learn about the scientific concept of ATTENTION, even independently of complex questions over how it relates to the folk concept, so I set aside the folk concept hereafter.

² This view is naturally suggested by some of the experimental work on the folk concept ATTENTION (Mole, 2008; De Brigard, 2010).

2 The threat of eliminativism

Despite the centrality of ATTENTION in cognitive science, many psychologists have expressed scepticism about the concept. Hommel et al. (2019) claim that it has too many meanings to be useful for science, and should be eliminated. Di Lollo (2018) agrees that the concept should be dropped, arguing that its definitions rely too much on metaphor. Anderson laments the ‘inconsistent usage and confusion’ besetting ATTENTION, and suggests replacing it with a ‘new technical vocabulary’ (2011, p.1). Uttal calls ATTENTION ‘a residual concept left over from earlier mentalist psychologies’ (2011, p.235). This scepticism is not new. Writing in 1982, Donald Broadbent suggests that avoiding ATTENTION would be ‘a step towards clarity’ (1982, p.253). Call these views ‘eliminativist’ about ATTENTION.

These critiques raise the question of whether ATTENTION is a scientifically useful concept, or one that should be done away with. There is a distinction between *entity* eliminativism (that a certain faculty does not exist) and *conceptual* eliminativism (that a certain concept or term should be eliminated from our scientific discourse). These kinds of eliminativism must be kept separate, since one can consistently hold one of them without the other (cf. Irvine & Sprevak, 2020). The eliminativists I mentioned are arguing in favour of a form of *conceptual* eliminativism, without accompanying commitment to *entity* eliminativism. This can be difficult to see as the distinction is usually not marked, and metaphysical claims about attention often accompany claims about the concept ATTENTION. For example, Britt Anderson claims that there is no such thing as attention (2011). This can make the view appear to be entity eliminativism. However, the substance of the argument is that attention is an ‘effect’ rather than a ‘cause’, so Anderson clearly believes that attention exists.³ In later work, this issue is clarified and the position is described as ‘terminological eliminativism’ (2021, p.1), and Anderson suggests that ATTENTION ‘has no clear *conceptual* association and no explanatory power’ (2021, p.3, my emphasis).⁴ Similarly, Hommel et al. are explicitly concerned with the concept ATTENTION, or term ‘attention’, saying: ‘we take the position that the term “attention” should be abandoned’ (2019, p.2288).

As I am concerned with the *concept* ATTENTION, I will be concerned with *conceptual* eliminativism in this paper. I assume a minimal form of entity realism about attention, which simply claims that attention exists. This minimal form of realism is consistent with the view that attention is actually many different faculties (Taylor, 2020). This form of realism doesn’t imply, for example, that attention is a natural kind (Prinz, 2012, p.90). As noted above, this minimal form of entity realism is consistent with the *conceptual* eliminativism that this paper rebuts, so assuming it in this context is not question begging.

It is tempting to think of conceptual eliminativism as worrying for pluralists about ATTENTION. Some pluralists claim that ATTENTION is a polysemous concept, with many disparate different meanings (Taylor, 2020). Other pluralists claim that

³ They link this idea to Chris Mole’s adverbial account of attention (Anderson 2011, p.2).

⁴ Thanks to an anonymous referee for pointing this out to me.

the concept ATTENTION is univocal (conceptual monism), but that it refers to many distinct kinds of cognitive processes (Pashler & Johnston, 1998, p.156). Some pluralists have suggested that the different meanings associated with the concept, or the different faculties that the concept refers to, share a family resemblance (Duncan, 2006; Taylor, 2015). But this raises the threat of eliminativism: if ATTENTION is just a family resemblance concept, wouldn't it be far more precise to eliminate ATTENTION from cognitive science, and replace it with a collection of more precise concepts?

However, eliminativism is not just a threat to pluralists. Monism also comes in a variety of forms, but all monists accept that the different expressions of attention share a common set of properties, which together spell out what it is for something to be an instance of attention (e.g. Prinz, 2012, pp.90–95). To see that monism is not immune to the threat of eliminativism, compare ATTENTION with the concept SUPERLUNARY OBJECT.⁵ This concept refers to all objects beyond the orbit of the moon. This concept has a univocal meaning, and all of its referents share properties that are necessary and sufficient for something to be a superlunary object: they are all objects, and they are all beyond the orbit of the moon. Still, the concept is not useful for modern astronomy, and we do not frame our astronomical theories in terms of it. Astronomists are right to eliminate it. Similarly, embracing monism *on its own* does not protect ATTENTION from the threat of eliminativism, any more than it does for SUPERLUNARY OBJECT. Rather, in order to answer the threat of eliminativism about ATTENTION, both monists and pluralists must show that ATTENTION fulfils some important scientific use, in virtue of which it is worth preserving.

3 Patchwork concepts

In this section, I summarise the general patchwork approach to scientific concepts, before applying it to ATTENTION in Section 4. Section 5 returns to eliminativism.

3.1 Conceptual spread

Sometimes, a scientific concept will be applied to a novel context, resulting in small alterations to make the concept suitable for application to the novel context. The result will be two different meanings, one of which is more suitable for the original context, and the other suited to the novel context. One example is HARDNESS in materials science (Wilson, 2006, ch.6; Haueis, 2022). Suppose scientists wish to measure the hardness of some metal. They would use an indentation test: indent the metal with a known amount of force, and measure the indentation. The shallower the indentation, the harder the metal. However, consider a shift in context, where we attempt to measure the hardness of rubber. The indentation test will not be useful, because the elastic properties of rubber will mean that if we indent it, the material will bounce back into place, incorrectly giving us the result that rubber is harder

⁵ Example from Griffiths (1997)

than steel. As a result, when applied to rubber, hardness is measured by a durometer. It's not just the method of measurement that changes, but the actual property that is measured shifts as well. In the steel case, the measured property is yield strength, whilst in elastomers like rubber, it's Young's modulus of elasticity (Haueis, 2022, cf. Wilson, 2006, ch.6). The result is two different meanings of the concept *HARDNESS*: one of which applies to metals, and the other to elastomers. These different meanings are what I call *subconcepts* of *HARDNESS*.

This kind of conceptual spread can occur again. When applied to ceramics, hardness is measured using a dry sand wear test, and the property tested for is wear rate (Haueis, 2022). As this process proliferates, we get a large number of different subconcepts of the overarching concept. Each subconcept is like one patch on a patchwork quilt: there are many different patches, which are more or less similar to one another, but no one is 'the correct one'. The overarching concept itself is like the quilt: it is constructed out of the individual subconcepts, as well as their relations. The patchwork approach is currently undergoing a surge in popularity, and has been applied to *HOMOLOGY* and *SPECIES* in biology (Novick, 2018; Novick & Doolittle, 2021); *CORTICAL COLUMN* in neuroscience (Haueis, 2021), and *GOLD* in chemistry (Bursten, 2018).

Why think that the patchwork approach is a promising way to approach *ATTENTION*? Conceptual spread typically causes confusion. It gives rise to large numbers of differing meanings, which can easily give the impression of a disordered and random mess of concepts. The patchwork approach is designed to help with this. As well as being a fruitful approach to scientific concepts generally, it also presents a comprehensive framework for understanding *why* this spread occurs, and for uncovering an underlying systematicity in the apparent mess. This ability to navigate apparent messiness is making the patchwork approach increasingly attractive in a philosophical environment where the complexity of central scientific concepts is becoming increasingly appreciated (e.g. Taylor & Vickers, 2017). As we saw in Section 2, it is precisely the apparent messiness of *ATTENTION* that motivates eliminativism. This indicates that the patchwork approach will be a fruitful framework through which to look at *ATTENTION*.

3.2 Features of subconcepts

There are four features of subconcepts or 'patches' that will be important for this paper.⁶ The first is the *scale* of the scientific investigation in which the subconcept applies. For example, *HOMOLOGY* in biology refers to any characteristics that are descended from a common ancestor. Some subconcepts of *HOMOLOGY* apply to body parts of organisms (e.g. a human arm and a bat's wing). The overarching *HOMOLOGY* concept also has subconcepts that apply at the genetic level, as when two genes are

⁶ The following analysis is inspired by Haueis (2022), Novick (2018) and Wilson (2006). I have modified Haueis' taxonomy in two ways for application to *ATTENTION*. First, Haueis talks about the property that is *measured* by a patchwork concept, rather than the properties that qualify something into inclusion into the extension of the concept. Haueis also frames his discussion in terms of *domains* rather than extensions.

descended from a common ancestor (Novick, 2018). Here there is a scale difference between two different *HOMOLOGY* subconcepts: body parts and genes.

The second feature of a subconcept is the *measurement/manipulation technique*. This is the approach to measurement or manipulation that is used in the scientific context in which the subconcept is deployed. For example, the indentation test is the measurement/manipulation technique for the subconcept of *HARDNESS* that applies to metals. For the subconcept of *HARDNESS* as applied to rubber, the measurement/manipulation technique would involve a durometer. With subconcepts of *ATTENTION*, the measurement/manipulation technique will be the experimental paradigms that the subconcept features in.

The third feature of a subconcept is the *property* or properties in virtue of which something falls into the extension of the subconcept. These properties can be relational. For example, the subconcept of *HOMOLOGY* that applies to genes spells out the relational properties required for two genes to be homologous (being descended from a common ancestor).

The final feature of a subconcept is its *extension*, i.e. the things that it applies to. As I show below, the extension of the various *ATTENTION* subconcepts will involve both the subject and the item(s) that are attended.

Different subconcepts can overlap in some of their features. For example, the different subconcepts of *HARDNESS* all concern macroscopic objects, and in that sense they all share the same scale. In this paper, a scientific concept is *patchwork* if and only if it has at least two subconcepts, (which are specific meanings that it takes on in particular scientific contexts) which vary in at least one of the four features just outlined.


4 ATTENTION as a patchwork concept

ATTENTION is operationalised in different ways in different scientific contexts. By examining three such contexts, I show that *ATTENTION* has various subconcepts that vary systematically in the way that qualifies *ATTENTION* as a patchwork concept. In each of the below subsections, I introduce a scientific context and outline a subconcept of *ATTENTION* that operates in that context, then show that it varies in terms of the four properties outlined above. Finally, I argue that its features can be explained in terms of the epistemic goals of the scientists, combined with certain constraints.

4.1 *ATTENTION*₁: The Posner paradigm

The first context is work in perceptual psychology that uses the Posner paradigm to study spatial and object-based attention in adult humans (Posner, 1980). Figure 1 represents the exogenous (A) and endogenous (B) conditions of the paradigm. Subjects are asked to foveate (focus their eyes on) a central point. This is the plus sign in Fig. 1, in between the two boxes. Then, a cue will instruct them where to direct their attention. In the exogenous condition (A), the cue is the black square that appears around one of the other squares. Then, a stimulus will appear at the

Fig. 1 The Posner paradigm. **A** is the exogenous condition, **B** is the endogenous one. A third element of the original figure has been removed, as it is not relevant. Reprinted from Vecara, S. and Rizzo, M. 2003. Spatial attention: normal processes and their breakdown. *Neurologic Clinics of North America*. 21: 575–607. Reproduced with permission. © Elsevier, 2003. All rights reserved



location where the cue is indicated. This is the ‘target’ (the * in Fig. 1). Subjects are instructed to respond, typically through button presses, when they see the target. In the endogenous condition (B), the cue appears at the fixation point (where subjects

are focussing their eyes). In Fig. 1, the cue in the exogenous condition is the ‘←’ that points to the relevant part of the display where the stimulus will subsequently appear.

Figure 1 depicts the valid conditions (the cues correctly indicate where the target will appear). There are also conditions where the cue incorrectly indicates where the target will appear (invalid condition). In some versions, there are also cases where no target appears (target absent condition). In this paradigm, if subjects are *faster* at reporting the presence of the target in the valid condition as opposed to the invalid one, then this is taken to be evidence that the subject paid attention to the stimulus (in the valid case). In other words, being faster to identify the target when the cue correctly indicated the presence of the target as opposed to when it incorrectly indicated its location is the indicator of attention. This is known as a reaction time advantage. That is, the marker of attention is a reaction time advantage when responding to valid target stimuli.

Recall the four features of patchwork subconcepts: scale, measurement/manipulation technique, properties that qualify something for inclusion in the subconcept’s extension, and the subconcept’s extension. The scale here is the macroscopic level. The measurement/manipulation technique is to cue subjects to particular areas of the visual field, and then measure a subject’s reaction times when responding to targets that appear there. The properties that qualify something for inclusion into the extension of the subconcept can be spelled out as follows: if a subject S has a reaction time advantage with relation to stimulus X when the cue validly indicates the location of X (as opposed to invalidly) then S counts as attending to X. Similarly, the *extension* of the subconcept would include anyone that meets this criterion (which thereby count as attending), as well as the relevant stimuli (that count as attended). Call this subconcept ‘ATTENTION₁’.

Now we should explain how these features make ATTENTION₁ scientifically useful in this particular context. This can be understood by looking at the *epistemic goals* of the scientific context. Let me explain. The Posner paradigm provides a clear way to experimentally demonstrate that a subject S attended some stimulus X: they need to have a reaction time advantage with relation to X in the valid condition, as opposed to the invalid condition. This makes it relatively straightforward to establish the presence of attention to some stimulus. This straightforwardness makes ATTENTION₁ very useful in research programmes that need a simple and relatively uncontroversial way of demonstrating that some stimulus was attended.

Here’s an example. One series of experiments investigated whether subjects could pay attention to stimuli that they sincerely denied awareness of (Kentridge et al., 2008; Norman et al., 2013). In order to fulfil this epistemic goal, scientists need two things:

- 1 An experimental method for ensuring that the subject sincerely denies awareness of the stimulus.
- 2 Clear experimental criteria for the subject to count as attending to the stimulus.

Both (1) and (2) are required in order to test whether subjects could attend to a stimulus of which they sincerely denied awareness.

A variety of techniques have been taken as good candidates for (1). In one instance (Kentridge et al., 2008) meta-contrast masking was used, which is a masking technique where one stimulus is swiftly followed by a second stimulus. The inner contours of the second stimulus coincide with the outer contours of the first stimulus, meaning that (when the timing of the stimuli is correct) signals concerning the first stimulus are confused with signals concerning the second stimulus, with the result that the first stimulus escapes awareness (Breitmeyer and Ogmen, 2004).⁷

(2) is the role that is fulfilled by $ATTENTION_1$. The subconcept is able to fulfil this role, by providing an appropriate measurement/manipulation technique to test for attention (cueing subjects to respond to target stimuli in valid and invalid conditions) and also a clear set of properties that need to be fulfilled for a subject to count as attending to a stimulus (a reaction time advantage in valid conditions, as opposed to invalid ones). So, as long as those properties can be shown in the absence of awareness of that stimulus, then attention without awareness would be demonstrated, which makes $ATTENTION_1$ perfect for role (2). The techniques appropriate for abolishing awareness (role (1)) were incorporated into a Posner paradigm (which fulfils role (2)), allowing attention without awareness to be tested. Subjects demonstrated a reaction time advantage with respect to the target stimuli, even when the stimulus escaped awareness (Kentridge et al., 2008). This was taken to demonstrate attention without awareness.

The important point for this paper is as follows: what is needed for this scientific context is a subconcept of $ATTENTION$ that can fulfil role (2). $ATTENTION_1$ can fulfil this role *because* of the very particular measurement/manipulation technique that it uses, and the very specific properties that qualify something for inclusion into its extension. It is these specific features of the subconcept that make it useful for this scientific context.

4.1.1 An alternative interpretation?

Jesse Prinz (2012) suggests that the subjects in the above experiments are not attending to the stimuli, even though they demonstrate a decreased reaction time (2012, p.115).⁸ Prinz defines $ATTENTION$ as accessibility to working memory (2012, p.92) and then points out that reportability is linked with working memory. From this, he concludes that, since subjects cannot report the stimuli (they sincerely deny awareness of it), this indicates information about the stimuli was not accessible to

⁷ Norman et al. (2013) used a different technique to fulfil (1). Gabor patches (rippled textures tilted to a particular orientation) were used to define the contours of an object. For example, one set of Gabor patches would be horizontal in the shape of a square, against a background of vertically oriented patches. They then caused the orientations to flip back and forth so quickly that *all* of the patches appeared as a blur to viewers (whilst always ensuring that the patches that defined the shape were perpendicular to the patches in the background). The result was a shape that subjects were unaware of.

⁸ Thanks to a referee for suggesting that I discuss this. Prinz's view appeared before the Norman et al. (2013) studies, so strictly speaking, it is directed at earlier work.

working memory and hence (according to Prinz's definition of ATTENTION) they are unattended. He suggests instead that it may be due to a pre-attentional system he calls 'orienting'. (2012, p.114). He argues against other experimental paradigms on similar grounds (2012, p.116 and p.96).⁹

I do not have space to fully consider this debate but it is worth a brief discussion, because the patchwork approach presents a natural diagnosis of the disagreement. The patchwork approach would claim that the disagreement arises (at least in part) from the use of different patches: one patch cashes ATTENTION out as accessibility to working memory (2012, p.95). Call this ATTENTION_{WM}. The ATTENTION_{WM} patch has a different set of properties that qualify something for inclusion into its extension from ATTENTION₁. In other words, according to ATTENTION_{WM}, for a subject S to count as attending to a stimulus X, S must exhibit properties like being able to verbally report information about X, being able to use information about X for high level reasoning, controlled and deliberative action, and so on (Prinz, 2012, p.92). This is in stark contrast to the properties that qualify something for inclusion into the extension of ATTENTION₁, which requires only that there be decreased reaction times in valid conditions, as explained above. Similarly, the measurement/manipulation techniques that ATTENTION_{WM} relies on are very different from the Posner paradigm: they are paradigms that probe working memory function (Prinz, 2012, pp.90–97).

Having noted these differences between ATTENTION_{WM} and ATTENTION₁, we have the resources to explain the disagreement. According to ATTENTION_{WM}, in order to fall under the extension of ATTENTION, we would need relatively demanding conditions to be in place, such as reportability, availability for controlled action, etc. Subjects in the above experiments do not instantiate these properties, so this sub-concept doesn't count them as attending. Conversely, ATTENTION₁ does not have this stipulation (it requires only that subjects show decreased reaction times in the valid as opposed to invalid conditions) so counts the subjects as attending (cf. Taylor, 2013). Here the disagreement stems (at least in part) from the different patches in play, not merely the empirical details. This is another advantage of the patchwork approach: it has the resources to account for differing interpretations of experimental results, as not wholly stemming from empirical matters of fact, but also the use of different subconcepts of ATTENTION.

An obvious question is: *why* would ATTENTION_{WM} be a good subconcept for thinking about attention? The view I am developing in this paper suggests that the answer requires examining the project's epistemic goals. The project in Prinz's case is to help explain phenomenal consciousness (2012). More specifically, the project requires a subconcept of ATTENTION that can be used to draw the distinction between stimuli that are perceived consciously and unconsciously.¹⁰ We need not get embroiled in the details of this view in order to see that ATTENTION₁ would be

⁹ Prinz also suggests that the results may be explained by eye saccades (2012, p.113). I do not have space to fully examine this claim, see Mole (2014) and Taylor (2013) for thorough discussion.

¹⁰ Prinz's view is that attention to intermediate-level representations are necessary and sufficient for consciousness. Intermediate-level representations are the layer of representational complexity above edges, blobs and local light conditions, but below conceptual representations (2012, pp.50–51).

unsuitable for this purpose, as it counts some stimuli as attended to that are not consciously perceived (this is what the studies outlined above demonstrate). Rather, a much more suitable subconcept of ATTENTION for this purpose pitches it at a much higher-level, at the level of working memory function. Thus, as the patchwork approach sees it, this the structure of ATTENTION_{WM} is explained in terms of epistemic goals (marking the distinction between consciously and unconsciously perceived stimuli).

4.2 ATTENTION₂: Infants

I have given a case where ATTENTION₁ is appropriate given the epistemic goals of the scientific context. In this section, I explore a subconcept of ATTENTION that aids scientific work by balancing the epistemic goals with the *constraints* of the scientific context.

ATTENTION₁ applies to adult humans. Its measurement/manipulation technique reflects this fact. It requires subjects to understand relatively complex task instructions, such as ‘maintain focus on the central spot, and then direct your attention where the arrow points’. This is appropriate given that the subconcept applies to adult humans. ATTENTION₁ becomes inappropriate when we change the scientific context to one involving human infants. Such subjects cannot understand complex task instructions, so the ATTENTION₁ subconcept becomes inappropriate. For this reason, a new subconcept is needed.

One option for measuring attentional fixation in infants uses looking behaviours. For example, one study investigated the extent to which faces attracted infants’ attention, as opposed to non-face images. They presented six-month old infants with an array of objects arranged in a circle around the display. One of these was a human face, and the others were random distractor objects such as alarm clocks and birds (Gliga et al., 2009). The experiment tracked where infants fixated their gaze, and how long they fixated it for. Here, eye-gaze fixation is taken to be a marker of the infant’s attention. Using gaze fixation to measure attention in infants has enjoyed widespread use for decades (e.g. Baillargeon, 2004; Spelke & Kinzler, 2007; Spelke et al., 1992).

Let’s call this subconcept ATTENTION₂. The scale involved is the same as that we encountered with ATTENTION₁, the macroscopic level of human subjects. However, the measurement/manipulation technique is different. It involves placing stimuli in front of the infant and tracking their *eye gaze*. This is very different from ATTENTION₁, which involves keeping subjects’ gaze fixated on the centre of the display (*away* from where the stimulus will appear), and measuring *reaction times*, rather than eye fixation. The property in virtue of which items fall into the extension of the subconcept has also shifted. For ATTENTION₂, if an infant fixates their gaze on item X, rather than other items in the display, then the infant is attending to X. Finally, the extension of the subconcept is very different from ATTENTION₁. Not only does one of them only apply to adults, but the extensions of the subconcepts vary in more fine-grained ways as well. For example, take a case where a subject maintains their gaze on one point in a display throughout, but they are cued to attend to a stimulus that appears in their

peripheral field, and have a reaction time advantage when responding to that stimulus. They would count as attending to that stimulus according to $ATTENTION_1$ (because of the reaction time advantage) but not $ATTENTION_2$ (because their eye-gaze was not directed there). So, the two subconcepts have different extensions.

Why does $ATTENTION_2$ differ from $ATTENTION_1$ in these particular ways? And how do these differences make $ATTENTION_2$ scientifically useful? We need to look at both the epistemic goals and the constraints of the scientific context in order to answer this question. The epistemic goal for which $ATTENTION_2$ is deployed involves infants. But this places *constraints* on the research programme, because infants cannot understand the kind of complex task instructions involved in $ATTENTION_1$. $ATTENTION_1$ presents a useful way to demonstrate the presence of attention to some stimulus in adult humans, but its measurement/manipulation technique becomes inappropriate in the context of testing on infants, because infants cannot understand task instructions. So, a new subconcept is required for this context, which is $ATTENTION_2$. The new subconcept will require a new measurement/manipulation technique, which involves eye-gaze fixation tracking. In this way, the features of an $ATTENTION$ subconcept can be useful by balancing the constraints of a scientific context with the need to fulfil its epistemic goals.

$ATTENTION_2$ is not the only subconcept of $ATTENTION$ that has a structure designed to overcome constraints placed by infants. Other subconcepts of $ATTENTION$ use heart rate changes to study attention in infants (Richards, 2010). This subconcept also finds a way to avoid the constraint that young infants cannot understand task instructions. The property that qualifies something for inclusion into the extension of the subconcept is again different (increase in heartrate), as is the extension, though the scale (macroscopic) and the set of individuals it is deployed on (human) are the same. Certain nonhuman animals present similar constraints, for which new subconcepts are again needed. When measuring attention in bees, scientists don't use reaction times, looking behaviours, or heartrate changes, but rather track which stimuli the bees fly toward (Morawetz & Spaethe, 2012). What we see with infants and non-human animals is consistent shifting in our subconcepts of $ATTENTION$, in order to overcome the challenges that those subjects present.

4.3 $ATTENTION_3$: Neuroscientific markers of attention

When we move to a neuroscientific context, we again find different subconcepts. Consider a project that has the epistemic goal of finding a neuroscientific marker of visual attention that is present in the absence behavioural measures, such as button presses or looking behaviours (Datta & DeYoe, 2009). Given this constraint, both of the previous subconcepts are clearly inappropriate, since they make use of these behavioural techniques.

Previous work on the neural underpinnings of visual attention has suggested that paying attention to a particular area of the visual field corresponds reliably to increased activity in the corresponding location in the visual cortex (Somers et al., 1999, Brefczynski et al., 2009). Mapping these relations is known as 'retinotopic

mapping'. Datta and DeYoe (2009) used this to develop alternative ways of measuring for the presence of attention, in the absence of behavioural cues.

In the first experiment, subjects were placed in an fMRI machine, and activity in areas V1 and V2 of visual cortex was monitored while they were presented with an array of visual stimuli. The array was a series of concentric circles, divided into wedge-shaped segments, similar to a dartboard. Subjects were given a fixation point in the centre of the array (the 'bullseye' of the dartboard). The fact that subjects did not move their gaze from the bullseye was confirmed with an eye-tracker. In one condition of this experiment, subjects were instructed to attend to a particular location, and report (using button presses) the colour of the segment, and the orientation of a pattern of lines that appeared at that segment.¹¹ Subjects were instructed to do this, with each of the segments on the 'dartboard', and their fMRI activity was catalogued.

This first experiment clearly makes use of behavioural measures such as button presses, but it was used as a stepping stone to develop an alternative measurement/manipulation technique that does not rely on such behavioural measures. In a second experiment, subjects were asked to 'secretly' attend to one of the segments, and provide no behavioural cues of which segment it was until after the fMRI analysis could be carried out. Using fMRI analysis, the experimenters were able to determine which segment was attended to with 100% accuracy (Datta & DeYoe, 2009).

The shift from behavioural/psychophysical measures of attention to neuroscientific ones will always involve a shift in scale. In this case, the scale has moved from the macroscopic subject level to the level of activation patterns in certain portions of visual cortex. The measurement/manipulation technique involves using fMRI to monitor blood oxygenation levels in areas V1 and V2 of visual cortex. As for the property that qualifies someone as attending to stimulus X, clearly an important property is increased activity in the areas of visual cortex that are correlated with the visual location of X. This is not itself sufficient for attention to X, since it is possible for there to be increased firing in V1 and V2 for other reasons. To qualify as attending to X, a subject must have a specific kind of increased activity in V1 and V2, which has been correlated with visual attention to a particular area in the first experiment, within the context of an experimental setup where they have been instructed to attend to a certain object. The extension of the concept (those who count as attending) are human adults, and the stimuli in question are the segments that correlate in location to the areas of the cortex that have increased activity. Call this subconcept ATTENTION₃.

It is potentially misleading to say that this subconcept dispenses with behavioural data. The initial condition, which allowed the fMRI findings to be *calibrated*, used subjects' button presses (in order to fix the patterns of activation that are correlated with visual attention to a particular area). After the experiment, subjects' button presses were used to *verify* the accuracy of the neuroscientific measure of attention (Datta & DeYoe, 2009). However, the properties that qualify something for inclusion into ATTENTION₃ do not themselves require behaviour by the subject (indeed, the subconcept is designed to dispense with them).

¹¹ In some conditions, subjects were asked to maintain attention on the central 'bullseye', instead of one of the segments.

Here is a case where the changes in the subconcept are driven by the epistemic goals of the scientific project. The epistemic goal of the project is to develop a neuroscientific measure of attention. This also acts as a self-imposed constraint, as it forces the experimenters to develop a subconcept that uses neuroscientific markers as its primary measure of attention. This epistemic goal then determines the subconcept's measurement/manipulation technique, which is purely neuroscientific, and involves no behavioural measures. The epistemic goal also determines the scale (areas of V1 and V2 in visual cortex), and the properties that qualify something for inclusion into the extension of the subconcept are likewise non-behavioural. These properties, in turn, will determine the subconcept's extension.

4.4 The same concept?

The patchwork view I have been developing holds that there are a range of different scientific contexts, and within these contexts there are different subconcepts of ATTENTION, which differ along the four dimensions explained above. These differences (I claim) are crucial for explaining how the different subconcepts are scientifically useful in the various contexts.

An opponent could resist the claim that there is a difference in the concept ATTENTION in each of the three scientific contexts. In this subsection, I outline two ways that this suggestion could be developed, which I call strong and weak views. I show that the strong view is implausible, and the weak one is compatible with the patchwork view.

The strong version of this objection accepts that the cases discussed above involve different scientific contexts, but claims that they do not involve *any* changes in the concept ATTENTION *at all*. On this view, the concept of ATTENTION is *precisely* the same in all three cases. By analogy, suppose I utter the sentence 'the ball is red' in the street and then utter the same sentence in my house. The current suggestion is that a shift in scientific context no more leads to a change in the concept of ATTENTION than moving from the street to my office changes the concept RED.

This strong view implies a conditional: if we were to take the concept of ATTENTION deployed in one experimental context, and deploy it in another, then *all* differences between the two cases would disappear. This suggestion is implausible because as I have argued, the subconcepts in each case place different criteria on the properties that qualify a subject as attending, and these different criteria change the extension of the subconcepts as well. We can see this very clearly by holding the context fixed, and seeing that the different subconcepts still diverge in their extensions. Take a context like the Posner paradigm, where we present subjects with a central cross and then flash up a stimulus outside of the subject's gaze. Suppose a subject looks directly at the stimulus, but does not click any buttons to indicate that they saw it. In this case, they could count as attending according to ATTENTION₂ (on the grounds that they fixated the stimulus), but they would not count as attending by ATTENTION₁ (on the grounds that there is no reaction time advantage). The subject would count as attending by one subconcept but not the other, so there is a difference in the subconcepts' extensions, as a result of the fact that they place

different conditions on inclusion into their extensions. This shows that the strong view explained above is implausible.¹²

However, there is also a weak interpretation of the view that the concepts are the same in all three cases, which is to accept that the three cases do involve different subconcepts, which are importantly different as a result of the epistemic goals and constraints of the different scientific contexts. However (the suggestion goes), at a coarse-grained level, there are similarities between the different subconcepts, which indicate that they are all different ways of sharpening the same overarching ATTENTION concept. One way of developing this weak suggestion would be to point out that all three of the subconcepts involve the subject mentally selecting some information in some way, and so it could be suggested that the three subconcepts are all different ways of sharpening the basic idea of ATTENTION as mental selection by the subject (Dicey Jennings, 2012, 2015). The suggestion need not be linked to this particular view of attention. On a suitably broad sense of ‘action’, it might also be claimed that all of the subconcepts are different ways of sharpening the notion of ATTENTION as selection for action (Wu, 2014).

This weak suggestion does accept that the subconcepts are *different* from one another, but also claims that, at a suitably coarse level of grain, the different subconcepts all share certain features in common, and each subconcept would be a way of sharpening these coarse-grained similarities, to make it suitable for different experimental contexts.

This weak suggestion is much more plausible, but it is entirely consistent with ATTENTION being a patchwork concept. Return to the various subconcepts of HARDNESS, as explained in Sect. 3.1. All of these subconcepts share similarities at a coarse-grain, which is that they describe a way in which a material is resistant to external pressure, and they could be thought of as alternative sharpenings of this shared similarity. The point of the patchwork approach is to understand how these concepts change their structure as they’re applied to new scientific contexts, and then to understand how these features of subconcepts make them scientifically useful in those contexts. The patchwork approach does not deny that they may share similarities.

4.5 Summary

I started the paper with two questions about ATTENTION. The first was, given that the concept is deployed in so many theoretical contexts in psychology, how do these shifts in context affect the concept, and why is this important for understanding how the concept aids scientific work? The theory developed in this paper can answer this.

¹² ATTENTION₁ and ATTENTION₃ can also differ in their extensions, at least in principle. For example, in a case where we observe the relevant firings in V1 and V2, but no reaction time advantage, the subject would count as attending according to ATTENTION₃, but not ATTENTION₁. The two subconcepts could then differ in their extensions. Similarly, in a case where a subject directly fixates a stimulus, but there is no increase in activity in the V1 and V2 (or that there is increased activity in another area of V1 or V2), then they would count as attending by ATTENTION₂ and but not ATTENTION₃.

The fact that ATTENTION is such a versatile concept in psychology is to be explained by its patchwork structure. Patchwork concepts can shift their meaning in at least four ways: their scale, the measurement/manipulation technique that they use, the property that qualifies something for inclusion into the subconcept's extension, and the subconcept's actual extension. By shifting their meaning in this way, new subconcepts fit with various different epistemic goals, and constraints. By utilising variation in these four properties, scientists are able to fashion new subconcepts of ATTENTION, and by doing this, the new subconcepts are useful for many different contexts. Each subconcept is a finely-honed tool, suitable for very specific contexts.

We are now in a position to see how the patchwork view differs from the view that ATTENTION is a family resemblance concept (Duncan, 2006; Taylor, 2015). The patchwork approach sees more systematicity than mere family resemblance. It claims that the different ATTENTION subconcepts vary in terms of four specific features, and it explains the scientific usefulness of the different subconcepts in terms of these four features. Finally, it provides an explanation of *why* there is this kind of systematic variation across these features: because of the various epistemic goals of practicing scientists, and specific constraints placed on them by particular scientific contexts. The patchwork framework exposes more complexity in subconcepts of ATTENTION than was previously apparent, but also shows that the complexity is systematic. Each subconcept is different from the others, but each one results from taking the epistemic goals and constraints of the scientific context, and using them to determine the subconcept's four specific properties. Where family resemblance sees only loose clusters of similarity, the patchwork view sees systematic and scientifically important variation.

The patchwork approach demonstrates its power by helping identify novel subconcepts of attention, and explain their structure and scientific usefulness. A more general result is as follows. The patchwork approach has not previously been applied to concepts in psychology. By demonstrating its power at helping us understand psychological concepts, this paper provides additional reason in support of the approach itself. The second question I started this paper with was: what does all of this shed on the *usefulness* of ATTENTION as a scientific concept? It is to this that I now turn.

5 Eliminativism answered

Recall the threat of eliminativism: why not simply do away with the overarching concept ATTENTION? We cannot respond to the threat of eliminativism by pointing out how useful ATTENTION's *subconcepts* are, because then the eliminativist will simply suggest that we do away with ATTENTION, and make do with ATTENTION₁, ATTENTION₂, ATTENTION₃. Rather, we have to find a useful theoretical role for the overarching ATTENTION concept, which justifies its preservation.

The patchwork approach developed above provides the resources to do this. In this section, I argue that ATTENTION is a *general reasoning strategy* concept. This view has two elements: first, the concept provides tacit guidance to scientists in how to go about manipulating the presence and absence of attention; second, ATTENTION

is an essential part of the development of ATTENTION'S *subconcepts*. I argue that this constitutes good reason to preserve it in cognitive science.

5.1 ATTENTION is a general reasoning strategy concept: tacit guidance to scientists

Return to HARDNESS. This concept encompasses many different subconcepts. However, there is some unity to the ways in which different HARDNESS subconcepts are used. The scientific contexts all involve *physically intervening* on a material, and *testing the manner in which it resists that intervention*. The way in which this kind of intervention takes place is different in the case of the different HARDNESS subconcepts, but they all reflect this general approach in different ways. Haueis calls this a 'general reasoning strategy' (2022). This is guidance, specified at a coarse-grained level, that helps scientists test for or manipulate the property in question. These are not explicit instructions written down anywhere, but tacit assumptions that scientists associate with the concept of HARDNESS. This general reasoning strategy will need to be refined and specified in different ways for each individual subconcept, but it is scientifically useful because it provides scientists with general guidance in how to intervene on the world in order to productively measure hardness.

Not every patchwork concept is associated with a general reasoning strategy (cf. Haueis, 2022). However, there are good reasons to think that ATTENTION is. In cognitive science, attention is typically introduced by reference to selection (e.g. Carrasco, 2011, p.1484; Knudsen, 2007, p.57). Of course, selection is not *sufficient* for something to be an instance of attention. A machine designed to sort out gum balls into different sizes is selecting from amongst the gum balls, but it is not attending (Example from Wu, 2011, p.97). So selection cannot provide a satisfactory *analysis* of attention, but it can form the basis of a general reasoning strategy for the overarching concept ATTENTION. The general reasoning strategy would be as follows: in order to test for attention, provide the subject with a range of different information to select, and then find a way to measure which of that information they have selected.

The idea is that scientists approach a new scientific context with this very general, coarse-grained feature of ATTENTION, and that guides them in the way they design experimental paradigms to test for and manipulate the presence and absence of attention. With ATTENTION₁, the subject is provided with a range of alternative locations to focus on, directed to focus on one, and a reaction time advantage is taken to be sufficient for attention. In the case of ATTENTION₂, the infant is provided with a range of alternative objects to fixate their gaze upon, and the one that is selected is taken as a marker of attention. For ATTENTION₃, the subject is given an array of segments, invited to select one, and then fMRI is used to measure which one that was. All of these subconcepts encompass the same general reasoning strategy, involving selection. In this sense, attention is a *general reasoning strategy concept* (it is consistent with this that the concept fulfils some other roles as well).

This constitutes the first point against eliminativism: that ATTENTION as a general reasoning strategy provides tacit guidance to scientists in how to go about testing

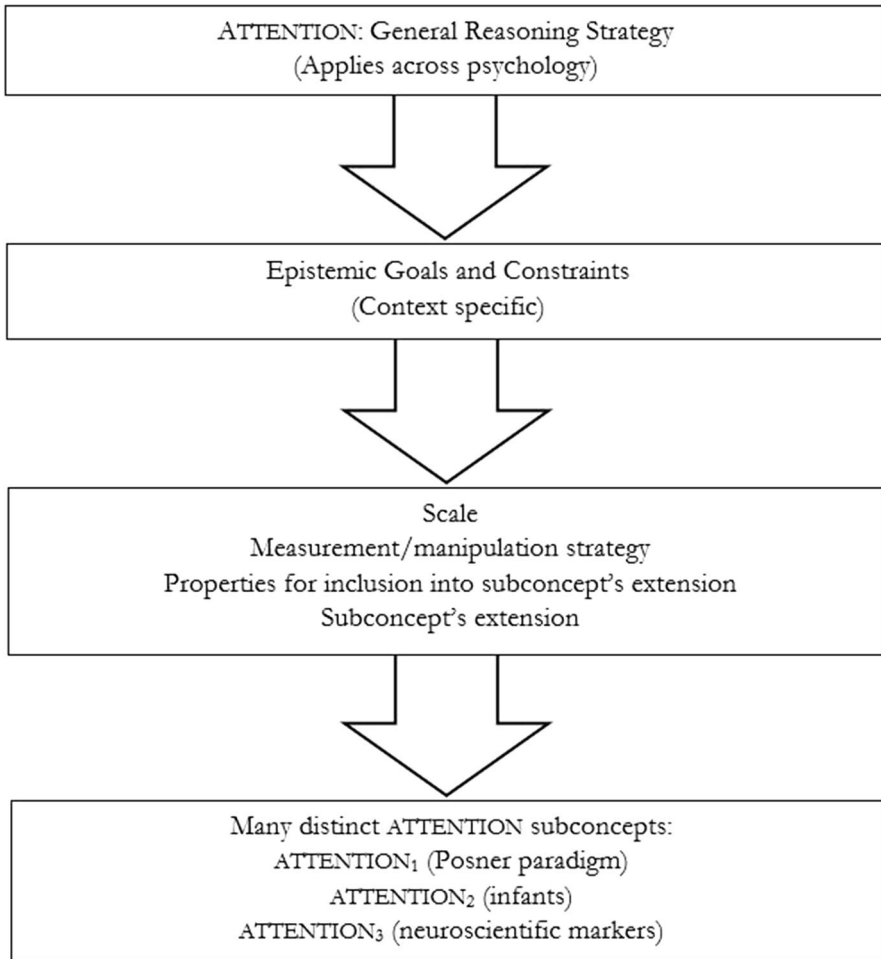


Fig. 2 How ATTENTION as a general reasoning strategy contributes to the development of ATTENTION’s subconcepts

for and manipulating the presence and absence of attention. To eliminate ATTENTION would be to lose this role.

5.2 ATTENTION is a general reasoning strategy concept: the development of subconcepts

There is also a second role for ATTENTION as a general reasoning strategy concept which is that it is essential for the development of the ATTENTION’s *subconcepts*. We can see this by drawing together several threads from this paper to get a picture of how ATTENTION’s subconcepts are developed (Fig. 2).

Figure 2 depicts a four-step process. First, start with ATTENTION, which provides general guidance to practicing scientists concerning how to test for attention: provide subjects with a range of alternatives, and then get them to select one of them above others. Second, draw on a particular scientific context to establish what its epistemic goals and constraints are. Third, specify the subconcept's properties (such as scale and measurement/manipulation technique) by deciding which properties are most appropriate given the general reasoning strategy and goals and constraints of the context. Fourth, by combining all of these properties, we generate new subconcepts for use in different experimental settings. On this view, we should understand the subconcepts of ATTENTION in terms of the fact that the same general reasoning strategy has been filtered through a variety of epistemic goals and constraints, to determine four specific properties of each subconcept, in virtue of which they are useful.¹³

The role of ATTENTION in the development of its subconcepts is especially important because it provides real bite in responding to the challenge of eliminativism. Eliminativism looks at the different subconcepts of ATTENTION, and (correctly noting their usefulness) asks why we can't do away with the overarching ATTENTION concept and just make do with the subconcepts. The picture of ATTENTION as a general reasoning strategy uses the process outlined in Fig. 2 to show that the overarching ATTENTION concept is an important part of the development of these subconcepts in the first place. For this reason, we cannot rely on the usefulness of the subconcepts to justify eliminating ATTENTION, because eliminating ATTENTION would undercut the development of ATTENTION subconcepts themselves, by knocking away the first step of the process outlined in Fig. 2.

Where the eliminativists see an unruly proliferation of concepts, the patchwork approach sees deep systematicity. According to this picture, the many different ways that ATTENTION has been defined and understood are the result of applying the same general reasoning strategy to a range of alternative scientific contexts, and the importance of this role is not at odds with the large amount of variation in the subconcepts.

It's important to distinguish the notion of ATTENTION as a general reasoning strategy from a philosophical *analysis* of attention. As already noted, some philosophical analyses of attention place selection at the core of their analysis (e.g. Dacey Jennings, 2012, 2020; Wu, 2014). The current suggestion does not imply anything as strong as this. It doesn't use selection as the basis for a reductive analysis of ATTENTION. Rather, it only accepts that selection is something that tacitly guides scientists in sharpening the concept so that it can be experimentally manipulated in a range of scientific contexts. The presence of a general reasoning strategy provides a kind of unity to the different subconcepts, but this does not commit us to a kind of monism about attention. Monists claim that the different instances of attention all share a set of properties that qualify something as an

¹³ This four-step process constitutes one of the most important novel contributions of the patchwork framework. This novel contribution is clearest when we contrast the patchwork approach with another view on the scientific use of ATTENTION. Wayne Wu has suggested that a sufficient condition for something to be an instance of perceptual attention is that the subject perceptually selects some item for the guidance of some experimental task (2014, p.11 and p.39). This picture of attention has some things in common with the analysis given above, especially the emphasis on selection. However, the patchwork approach also provides a picture of how this general reasoning strategy is involved in the development of the ATTENTION subconcepts that (as we have seen) are so useful for psychology.

instance of attention. That is, that they share a common structure, to be specified at the neural, psychological, or personal level. For ATTENTION to be a general reasoning strategy concept only implies that the notion of selection gives scientists general tacit instructions about how to go about testing for the faculty. It does not make assumptions about the nature of the faculty that is thereby tested.

This paper has pursued a novel application of the patchwork approach to psychology, and in doing so, has exposed and explained various ATTENTION subconcepts. The explanation has identified four core properties of each subconcept in virtue of which they are scientifically useful, and provided an explanation of why the subconcepts have these properties. Here, I have explored how ATTENTION provides tacit guidance to scientists in testing for attention, and I have used the approach to explain how the overarching concept interacts with the subconcepts (Fig. 2).

5.3 Non-selective attention?

An opponent might offer the following objection. I characterised the general reasoning strategy partially in terms of selection. Many accounts of attention place selection centre stage (e.g. Wu, 2014, pp.95–99, Dacey Jennings, 2020, p.120). However, some have suggested that attention need not involve selection, or at least not obviously so.¹⁴ Above I mentioned Jesse Prinz (2012), who suggests that diffuse attention does not involve selection (2012, pp.90–91 and p.95). Sebastian Watzl claims that attention involves the organisation of mental states into a priority structure, which either does not involve selection, or at least a unique kind of selection (2017). Adrienne Prettyman notes that diffuse attention can be spread across a group of items as a whole, which does not seem to involve *selecting* any one object (2021). Datta and DeYoe talk about an ‘attentional landscape’ with peaks and valleys (2009). More directly, Kranti Saran argues that selection is neither essential nor necessary for attention (2018). Most starkly, Fazekas and Nanay (2021) claim that ‘attention is amplification, not selection’. It might be suggested that really, attention can manifest in a vast array of different ways, and at least some cases are non-selective.

‘Non-selective attention’ might be raised as a counterexample to my view of ATTENTION as a general reasoning strategy concept, since I claimed that one aspect of that involved giving subjects a range of stimuli to select from, and testing which one(s) they selected.

There are two points that can be made in response to this worry. The first is that many of the thinkers in question do not deny that attention involves selection. They just emphasise other explanatorily important features of attention as well. For example, Fazekas and Nanay’s view is that the common neural underpinning of attention is to be explained by neuronal amplification, not that attention doesn’t involve selection (2021). Prettyman argues that even diffuse attention involves selection of objects for thought and action (indeed, she *defines* attention in terms of selection) (2021, pp.374–376). Similarly, Datta and DeYoe often use ‘attention’ and ‘selection’ interchangeably, sometimes saying ‘attentional selection’ (e.g. 2009, p.1037).

¹⁴ Thanks to an anonymous referee for pushing me on this.

A more general response to this objection relies on the distinction drawn in Section 5.2 between ATTENTION as a general reasoning strategy and a philosophical *analysis* of attention. A general reasoning strategy is intended only to give tacit instructions on how to *test for* the presence of attention. It does not aim to give an analysis of *what attention is*, either in terms of necessary or sufficient conditions.

In other words, we need to distinguish two claims:

- i) Selection is a necessary property of attention.
- ii) Selection is a necessary part of the experimental paradigms that help us test for attention.

The first claim is about the metaphysics of attention, and properly belongs to the project of giving a philosophical analysis of attention. I am entirely agnostic about (i) here: attention may necessarily involve selection, or it may not. When we are interested in ATTENTION as a general reasoning strategy, we are only trying to uncover the ways that attention is experimentally operationalised in psychology, and the ways in which this strategy explains the diversity of subconcepts uncovered in this paper. So my view is only committed to (ii). It is consistent to claim that selection is not a necessary property of attention, whilst accepting that the psychological paradigms that manipulate attention do involve selection. Indeed, when we examine these paradigms (as we saw in Sect. 4), we see that they do: they involve the subject taking some information from the environment, and using them to execute a task. To say that selection is necessary for scientifically manipulating attention is very different from saying that it is a necessary property of attention.

This highlights another novel contribution of the patchwork approach: the distinction between selection as part of ATTENTION as a general reasoning strategy concept and the notion of selection as a necessary property of attention allows us to reject eliminativism, without taking a stand on metaphysical issues such as (i).

6 Conclusion

I have argued for three things:

- 1 ATTENTION is a patchwork concept. It encompasses distinct subconcepts, which differ in terms of their scale, measurement/manipulation techniques, the properties in virtue of which something falls into their extension, and their extension.
- 2 The difference in these subconcepts is to be explained by a combination of the epistemic goals and the constraints of specific scientific contexts.
- 3 ATTENTION is a general reasoning strategy concept. At least one of its core roles is to provide general guidance to scientists on how to manipulate and measure attention, and another of its roles is to participate in the development of its subconcepts. Recognising this provides us with the resources to resist eliminativism about ATTENTION.

Much has been written about attention. The approach developed in this paper is unique in that it exposes a high level of diversity in the subconcepts of ATTENTION, whilst also emphasising the systematic connections between these subconcepts. This paper also serves as indirect support for the patchwork approach itself. This paper shows the usefulness of the patchwork approach in helping us understand a central concept in psychology, which is not a field that the approach has been applied to before. Finally, this paper leads us to a general conclusion: when interpreting scientific work on attention, it is important not only to concentrate on the empirical results, but also on the subtle and often overlooked ways in which scientific context have shaped our own thinking about attention, and about ATTENTION.

Acknowledgements Thanks to Ali Boyle, Philipp Haueis, Rose Novick and Katie Robertson for helpful conversation. Thanks to two anonymous referees and the editor for detailed comments that improved the paper immensely.

Author contribution The author performed the research for this study, and wrote the paper. The paper contains no work by other authors, except where this is appropriately cited in the text.

Data Availability No data was generated for this study.

Declarations

Ethical approval The paper is a purely literature-based philosophical piece, and does not require ethical approval.

Informed consent No participants were involved in the writing of this paper.

Competing interests No competing interests to declare, financial or otherwise.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anderson, B. (2011). There is no such thing as attention. *Frontiers in Psychology*, 2, 1–8.
- Anderson, B. (2021). Stop paying attention to “attention.” *Wires Cognitive Science*, 14(1), e1574.
- Baillargeon, R. (2004). Infants' reasoning about hidden objects: Evidence for event-general and event-specific expectations. *Developmental Science*, 7(4), 391–414.
- Brefczynski, J., Datta, R., Lewis, J., & DeYoe, E. (2009). The topography of visuospatial attention as revealed by a novel visual field mapping technique. *Journal of Cognitive Neuroscience*, 27(7), 1447–1460.
- Breitmeyer, B., & Ogmen, H. (2004). Recent models and findings in visual backwards masking: Comparison, review, and update. *Perception and Psychophysics*, 62(8), 1572–1595.

- Broadbent, D. E. (1982). Task combination and selective intake of information. *Acta Psychologica*, 50, 253–290.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4, 72–122.
- Bursten, J. (2018). Smaller than a breadbox: Scale and natural kinds. *The British Journal for the Philosophy of Science*, 69, 1–23.
- Carrasco, M. (2011). Visual attention: The last 25 years. *Vision Research*, 51, 1484–1525.
- Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, 7, 308–313.
- Cartwright, N. (1994). Fundamentalism vs. the patchwork of laws. *Proceedings of the Aristotelian Society*, 103, 279–293.
- Chun, M., Golomb, J., & Turk-Browne, N. (2011). A taxonomy of internal and external attention. *Annual Review of Psychology*, 62, 73–101.
- Datta, R., & DeYoe, E. (2009). I know where you are secretly attending! The topography of human visual attention revealed with fMRI. *Vision Research*, 49, 1037–1044.
- De Brigard, F. (2010). Consciousness, attention and commonsense. *Journal of Consciousness Studies*, 17(9–10), 189–201.
- Di Lollo, V. (2018). Attention is a sterile concept; iterative re-entry is a fertile substitute. *Consciousness and Cognition*, 64, 45–49.
- Dacey Jennings, C. (2012). The subject of attention. *Synthese*, 189, 535–554.
- Dacey Jennings, C. (2015). Consciousness without attention. *Journal of the American Philosophical Association*, 1(2), 276–295.
- Dacey Jennings, C. (2020). *The attending mind*. CUP.
- Dickie, I. (2015). *Fixing reference*. OUP.
- Duncan, J. (2006). Brain mechanisms of attention. *The Quarterly Journal of Experimental Psychology*, 59, 2–27.
- Dupré, J. (1993). *The Disorder of Things*. Harvard Press.
- Fazekas, P., & Nanay, B. (2021). Attention is amplification, not selection. *The British Journal for the Philosophy of Science*, 72, 299–324.
- Gliga, T., Elsabbagh, M., Andravizou, A., & Johnson, M. (2009). Faces attract infants' attention in complex displays. *Infancy*, 14(5), 550–562.
- Griffiths, P. (1997). *What emotions really are*. Chicago Press.
- Hauéis, P. (2021). The death of the cortical column? Patchwork structure and conceptual retirement in neuroscientific practice. *Studies in History and Philosophy of Science*, 85, 101–113.
- Hauéis, P. (2022). A generalised patchwork approach to scientific concepts. *The British Journal for the Philosophy of Science*. <https://doi.org/10.1086/716179>
- Hommel, B., Chapman, C., Cisek, P., Neyedli, H., Song, J., & Welsh, T. (2019). No one knows what attention is. *Attention, Perception and Psychophysics*, 81, 2288–2303.
- Irvine, E. (2013). *Consciousness as a Scientific Concept*. Springer.
- Irvine, E., & Sprevak, M. (2020). Eliminativism about consciousness. In U. Kriegel (Ed.), *The Oxford Handbook of the philosophy of consciousness* (pp. 348–370). OUP.
- Kentridge, R., Nijboer, T. C. W., & Heywood, C. A. (2008). Attended but unseen: Visual attention is not sufficient for visual awareness. *Neuropsychologia*, 46(3), 831–869.
- Knudsen, E. (2007). Fundamental components of attention. *Annual Review of Neuroscience*, 30, 57–80.
- Magri, E. (2020). Social sensitivity and the ethics of attention. *European Journal of Philosophy*, 30, 725–739.
- Malt, B. (1994). Water is not H₂O. *Cognitive Psychology*, 27, 41–70.
- Mole, C. (2008). Attention and consciousness. *Journal of Consciousness Studies*, 15(4), 86–104.
- Mole, C. (2010). *Attention is Cognitive Unison: An essay in philosophical psychology*. OUP.
- Mole, C. (2014). Attention to unseen objects. *Journal of Consciousness Studies*, 21(11–12), 41–56.
- Morawetz, L., & Spaethe, J. (2012). Visual attention in a complex search task differs between honeybees and bumblebees. *The Journal of Experimental Biology*, 215, 2515–2523.
- Norman, L. J., Heywood, C. A., & Kentridge, R. W. (2013). Object-based attention without awareness. *Psychological Science*, 24(6), 836–843.
- Novick, A. (2018). The fine structure of 'homology.' *Biology and Philosophy*, 33, 6.
- Novick, A., & Doolittle, W. F. (2021). 'Species' without species. *Studies in History and Philosophy of Science*, 87, 72–80.
- Pashler, H., & Johnston, J. (1998). Attentional limitations in dual-task performance. In H. Pashler (Ed.), *Attention*. Psychology Press.
- Posner, M. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3–25.

- Prettyman, A. (2018). Seeing the forest and the trees: A response to the identity crowding debate. *Thought*, 7, 20–30.
- Prettyman, A. (2021). What is diffuse attention? *Mind and Language*, 38, 374–393.
- Prinz, J. (2012). *The Conscious Brain: How attention engenders experience*. OUP.
- Richards, J. (2010). The development of attention to simple and complex visual stimuli in infants: Behavioural and psychophysiological measures. *Developmental Review*, 30, 203–219.
- Saran, K. (2018). Meditative attention to bodily sensations: Conscious attention without selection? *Journal of Consciousness Studies*, 25(5–6), 156–178.
- Sergent, C., Wyart, V., Babo-Rebelo, M., Cohen, L., Naccache, L., & Tallon-Baudry, C. (2013). Cueing attention after the stimulus is gone can retrospectively trigger conscious perception. *Current Biology*, 23(2), 150–155.
- Silins N, Siegel S (2013) Attention and perceptual justification. In Blockheads! Essays on Ned Block's Philosophy of Mind and Consciousness. MIT Press
- Somers, D., Dale, A., Seiffert, A., & Tootell, R. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proceedings of the National Academy of Science*, 96(4), 1663–1668.
- Spelke, E., & Kinzler, K. (2007). Core knowledge. *Dev Sci*, 10(1), 89–96.
- Spelke, E., Breinlinger, K., Macomber, J., & Jacobsen, K. (1992). Origins of knowledge. *Psychological Review*, 99, 605–632.
- Taylor, H. (2013). Is attention necessary and sufficient for phenomenal consciousness? *Journal of Consciousness Studies*, 20(11–12), 173–194.
- Taylor, H. (2015). Against unifying accounts of attention. *Erkenntnis*, 80, 39–56.
- Taylor, H. (2020). Attention, psychology and pluralism. *The British Journal for the Philosophy of Science*, 69(4), 935–956.
- Taylor, H., & Sayim, B. (2018). Crowding, attention and consciousness: In defence of the inference hypothesis. *Mind and Language*, 33, 17–33.
- Taylor, H., & Vickers, P. (2017). Conceptual fragmentation and the rise of eliminativism. *The European Journal for the Philosophy of Science*, 7, 17–40.
- Treisman, A. (2006). How the deployment of attention determines what we see. *Visual Cognition*, 14, 411–433.
- Uttal, W. (2011). *Mind and brain*. MIT Press.
- Vecara, S., & Rizzo, M. (2003). Spatial attention: Normal processes and their breakdown. *Neurologic Clinics of North America*, 21, 575–607.
- Watzl, S. (2017). *Structuring mind*. OUP.
- Wilson, M. (2006). *Wandering significance*. OUP.
- Wilson, M. (2017). *Physics avoidance and other essays in conceptual strategy*. OUP.
- Wu, W. (2011). Attention as selection for action. In C. Mole, D. Smithies, & W. Wu (Eds.), *Attention: Philosophical and psychological essays*. Oxford University Press.
- Wu, W. (2014). *Attention*. Routledge.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.