There is no such thing as attention

Britt Anderson*

Department of Psychology and Centre for Theoretical Neuroscience, University of Waterloo, Waterloo, ON, Canada

Given that the core issues of attention research have been recognized for millenia, we do not know as much about attention as we should. I argue that the reasons for this failure are (1) we create spurious dichotomies, (2) we reify attention, treating it as a cause, when it is an effect, and (3) we equate a collection of facts with a theory. In order to correct these errors, we need a new technical vocabulary that allows for attentional effects to be continuously distributed, rather than merely present or absent, and that provides a basis for quantitative behavioral predictions that map onto neural substrates. The terminology of the Bayesian decision process has already proved useful for structuring conceptual discussions in other psychological domains, such as perception and decision making under uncertainty, and it had demonstrated early success in the domain of attention. By rejecting a reified, causal conception of attention, in favor of theories that produce attentional effects as consequences, psychologists will be able to conduct more definitive experiments. Such conceptual advances will then enhance the productivity of neuroscientists by allowing them to concentrate their data collection efforts on the richest soil.

Keywords: attention, Bayes, reification

INTRODUCTION

Attention may have to go, like many a faculty once deemed essential, like many a verbal phantom, like many an idol of the tribe. It may be an excrescence on Psychology. No need of it to drag ideas before consciousness or fix them, when we see how perfectly they drag and fix each other there.

<div align="right">James' (1890/1950).</div>

Attention has been a cornerstone of psychological interest since antiquity, and was one of the three pillars on which modern experimental psychology was erected (Titchener, 1908). And yet, despite all this time and the many investigations, our knowledge of attentional phenomena is little changed from that of the ancient Greeks. Quoting Aristotle ("…is it possible or not that one should be able to perceive two objects simultaneously in the same individual time?") and Lucretius ("…things are not seen sharply 'save those for which the mind has prepared itself.'"), Hatfield (1998) provides the evidence in his review on attention and classical thought.

These ideas have remained preoccupations of experimental psychology ever since. Echoing Aristotle, Angell and Pierce (1892) wrote: "The essential question is whether we can interpret *as simultaneous* two or more disparate simultaneous sensations…" and recently the same idea was central to Huang and Pashler (2007) who stated "This question about [the] possibility of simultaneous selection of two feature values is very fundamental." Many of the current papers on multiple object tracking can also be understood as investigations into this same, 2000 year old, question (e.g., Drew et al., 2009).

The Lucretian idea of expectancy was present in Helmholtz (1881): "…we must form as clear a notion as possible of what we expect to see. Then it will actually appear." And recently, "See What You Want to See: Motivational Influences on Visual Perception" (Balcetis and Dunning, 2006). The quote by Lucretius also emphasizes the idea that objects which are the focus of attention might in some way be perceived better; yet the claim that attention alters perceptual quality is still actively debated (e.g., Carrasco, 2009).

Although we currently cite Cherry (1953) for the cocktail party effect, the problem and its paradox was presented a few hundred years earlier by Stewart (1792/1866): "When two persons are speaking to us at once, we can attend to either of them at pleasure…This power, however, of the mind to attend to either speaker at pleasure, supposes that it is, at one and the same time, conscious of the sensations which both produce."

Thus, the core phenomena that motivate attention research have been clearly appreciated for a very long time. However the progress we have made has been slight. Contrast it with the progress made on another preoccupation of Lucretius', the atomic nature of matter. Our poorer progress might be because attention is the harder problem. Or it might be because of embedded misconceptions about the nature of the term attention itself. If our metaphors for attention guide our thinking (Fernandez-Duque and Johnson, 1999), then we should look to these same metaphors to account for our lack of progress.

I argue that our slow progress in understanding attention can be attributed to three sources, two general and one specific. A general fault of psychology is its predilection to binarize empirical phenomena; we try and shoehorn everything into being either this or that. As a result we construct and pursue false dichotomies. A more specific error is that we misuse the word attention. Attention has been plurally defined and this leads to inconsistent usage and confusion. More importantly, attention has been reified; it is used as a concrete concept that can act in a causal fashion, e.g., "Attention helps optimize the use of our system's limited resources…" (Carrasco, 2009). This logical fallacy leads to misplaced empirical

efforts. In fact, attention never causes anything, because there is no such *thing* as attention. There are, however, many empirical findings that can be accurately labeled attentional. In a phrase, attention is more adjectival than nominal (for an adverbial account see Mole, 2011). The third impediment to faster progress in understanding attentional effects is the equating of data with theory. While neuroscientists reify too ("Attention increases sensitivity of V4 neurons"; Reynolds et al., 2000), the more important problem for neuroscience is the tendency to emphasize data collecting over theory. As Poincaré said, "*Le savant doit ordonner; on fait la science avec des faits comme une maison avec des pierres; mais une accumulation de faits n'est pas plus une science qu'un tas de pierres n'est une maison.*"[1] Studies reporting the patterns of functional brain activations or neuronal firing while subjects perform tasks that yield attentional effects will ultimately tell us much about the *how* of attention, but not the *what is* attention; that kind of understanding will require the combination of neuroscientific data with psychological theory. The challenge for a terminology of attention is to put attention in the context of something to be explained rather than axiomatically assumed. This conceptual transition from cause to effect must also provide a space for the sensible generation of experiments and the deriving of mechanistic accounts from neural data.

The basic plan for the paper is as follows. First, I examine how we invent false dichotomies and then how we have misused the term attention. I then expand on the claim that neural data does not enhance our understanding of attention *per se*. I conclude with remarks on what an alternative approach to attention might look like and highlight research results, old and new, that benefit from this perspective.

FALSE DICHOTOMIES ARE DETRIMENTAL TO ATTENTIONAL RESEARCH

The allure and the limits of a binary approach to psychological science were highlighted by Newell (1973) 40 years ago. In spite of the caveats, the history of attentional research is a lengthening list of putative either-or components (e.g., top-down/bottom-up; see **Table 1**). This approach to attention research is harmful because instead of addressing the core phenomenal components we pursue a pseudo-question: is our dichotomy true?

While the statement that attentional cues are *either* endogenous *or* exogenous is phrased as a choice between two distinct alternatives, this appearance is an illusion. When we bisect outcomes in this way, our terms do not provide a complete and disjoint partition of the relevant space. The question as to whether an object is pink or purple is clearly poorly phrased. Because we know many objects that are neither pink nor purple, and some that are both, we can quickly detect our error. Pink and purple are not the names of crisp, disjoint sets in color space. Membership in the class of pink things is fuzzy. There are objects that have a non-zero membership value in both the set of pink things and the set of purple things. The same flaws are present, but not as obvious, when approaching attentional topics because our terms are variably defined and

[1]The Scientist must set in order. Science is built up with facts, as a house is with stones. But a collection of facts is no more a science than a heap of stones is a house. (Wikiquote, http://en.wikiquote.org/wiki/Henri_Poincaré).

Table 1 | The bifurcation of attention.

Top-down	Bottom-up
Exogenous	Endogenous
Divided	Focused
Feature	Object
Local	Global
Automatic	Controlled
Voluntary	Involuntary
Non-conscious	Conscious
Peripheral	Central
Intraperceptual	Extraperceptual
Passive	Active
Pre-attentive	Attentive

the phenomena are much less familiar. To demonstrate the case, I will briefly discuss two examples of false dichotomies in attention research: pre-attentive or attentive processing and endogenous or exogenous cuing.

PRE-ATTENTIVE OR ATTENTIVE PROCESSING

As the name suggests, pre-attentive processing is "...temporally prior to attention." (Logan, 1992). The division presumes a serial process. First pre-attentive processing occurs and then *subsequently* attention happens. By hypothesis it would be contradictory to speak of attentional effects on pre-attentive processing. A quick march through the work on pre-attentive processing shows that the assumption of a dichotomy has led to interpretative stances that were more pronounced than the data warranted.

Some of the early evidence that led to this division is reviewed in Treisman (1985). In the beginning, the conceptual division into pre-attentive and attentive processing was clear, but the data supporting such a crisp separation were not. One application of this binary division was feature integration theory (FIT; Treisman and Gelade, 1980). The early version of FIT proposed that identifying targets defined by conjoined features required attention for binding, and this in turn required spatial localization. Since identifying targets defined by a single feature did not require binding, and therefore did not require spatial localization, it could be asserted that only single feature targets could be correctly identified when mislocalized. Since localization was an obligatory preliminary stage in the attentionally dependent binding process, targets defined by conjoined features could not be identified when mislocalized.

While the early evidence tilted in this direction, it did not provide the clear separation the dichotomy required. For example in Experiment VIII of Treisman and Gelade (1980) participants correctly identified conjunction targets and feature targets about 80% of the time. If the accuracy focused on trials where the mislocation was one position away from the true target position, conjunction targets were accurately reported 72% of the time and feature targets 82%. For mislocations of two or greater positions it was 50 and 68% respectively. Clearly, there was a relationship between spatial location and accuracy of detection that was stronger for targets defined by conjoined features, but it was not dichotomous.

Another implication of the original version of FIT was that "...[for] targets defined by a single feature, the [spatial] cue should have very little effect." Figure 4 in Treisman (1985) shows an effect of cue validity on the d' for feature search that is similar in shape, though less extreme, to that for conjunction search; a quantitative more than a qualitative effect. The finding of quantitative rather than qualitative differences led Prinzmetal et al. (1986) to conclude that "Contrary to Treisman and Gelade (1980), features are not registered without attention." Subsequently, researchers were forced into contradiction when attentional cues were shown to speed target identification in pre-attentive searches (e.g., Theeuwes et al., 1999). Ultimately, the result was that we generated a lot of data. A million trials later we conclude that in fact there are not two distinct kinds of searches (Wolfe, 1998).

ENDOGENOUS OR EXOGENOUS CUE

The contrast between endogenous and exogenous cuing gives another example of the inefficiency of pursuing axiomatic, dichotomous divisions in attention research. In a common version of a cuing task a participant reports the side of a target. Luminance increments briefly precede target presentation and are deemed to be exogenous cues: cues that automatically attract attention. Centrally presented arrows that point toward or away from the target's location are an example of an endogenous cue: cues where the symbolic content mediates their effect.

In fact, it is not the case that endogenous cues must have a learned, symbolic value. Centrally presented gaze direction is an effective cue for locating a peripheral target (Ristic et al., 2002; Brignani et al., 2009). Endogenous cues do not have to be predictive (Dodd and Wilson, 2008). On the other hand, the automatic response to exogenous cues is not necessarily automatic. In a classic study, Yantis and Jonides (1990) showed that when subjects are engaged in a demanding task and have allocated their processing resources to a particular location (or are using a valid symbolic cue) the effect of a sudden luminance increment (the ultimate exogenous cue) may be muted or absent (Yantis and Jonides, 1990). Thus, our binary construction of cues as belonging to either one of two disjoint types: endogenous or exogenous, is false. After a number of studies, we now know that the delineation of cues into exogenous or endogenous is imprecise, but despite this we don't know much more about attention *per se.*

These two examples illustrate the general problem. We have a natural impulse to subdivide our observations into binary states. When this designation is done prematurely, we spend time investigating the classification instead of the underlying phenomenon. Further, it hinders us from considering alternatives where a continuous, graded, or fuzzy description would be more accurate.

ATTENTION HAS BEEN REIFIED

While our limited progress to date can be partially explained by our inclination to conduct research as a game of 20 questions, an approach ill-suited to psychological phenomena, a more important problem is our language. We use the term attention to mean different things and, more importantly, we have reified attention. The result is research incorrectly focused on explaining attention as a causal agent rather than the more correct conception where attention is seen as a convenient semantic label for a category of experimental result. We need to recognize attention is an effect and not a cause.

The fact that attention is variably defined, and that these multiple definitions obscure the implications of experimental work is old news. For example, Allport (1993) writes "More fundamentally what is meant by *selection*?" and a couple of paragraphs later, "Similar, if not more confusing ambiguities surround the usage of the term *attention.*"

One of our responses to this plurality of definitions has been to rely on "Jamesian Confidence." James' (1890/1950) famous phrase: "Every one knows what attention is." is implicitly (often explicitly) invoked whenever researchers report work on attention without supplying a concrete definition. The presumption is that the researcher's sense of attention will be clear from context and that James' exhortation can be taken to imply more than it says, that is that everyone knows what attention is, *and they all think it is the same thing.* In fact, communal practice reveals that we do not. Attention is subdivided by modality (visual versus auditory), level of analysis (feature versus object), and spatial extent (focal versus global). Attention is invoked as the label for a general preparedness to respond. Attention as vigilance has both a negative aspect (in that one may fail to detect a target) and a positive aspect (where one fails to inhibit a response when presented a non-target item). Attention is treated as a vector where there can be deficiencies in magnitude (implied by the phrase attention deficit) or direction (implied by a term like disengage deficit). Attention can also be given a temporal dimension when people speak of an attention span. None of these senses seem quite what James had in mind as the *obvious* one. Of course, some current workers still restrict the word attention to James' sense. Huang and Pashler (2007) refer to attention in terms of conscious awareness and not the more common selection. The problem with having so many definitions extends beyond the risk of confusion. The fact that there are so many variable definitions empowers researchers to create newer, eclectic ones. Colby and Goldberg (1999) write that "...one could say that their [lateral intraparietal neurons] activity selects targets from the environment for possible but easily cancelable saccades. The latter statement is a good definition of visual attention in the primate." I am not aware of anyone else using the selection of easily cancelable saccades as a good definition for attention, but given that we have so many definitions for attention already, why not one more?

On the one hand, attention's many definitions may impede research progress because they foil clear communication, but it may also be the case that the existence of so many apparently equally good definitions is a symptom of a more general problem: a fundamental misconstrual of what is attention. This idea has also been around since the advent of modern experimental psychology, but it has not received the consideration it warrants. Cattell and Farrand (1896) wrote that "...if we undertake to study attention or suggestibility we find it difficult to measure definitely a definite thing..." and this idea is echoed in more recent times by Johnston and Dark (1986) "It is difficult to conceptualize a process that is not well defined, and it is difficult to falsify empirically a vague conceptualization, especially one that relies on a homunculus."

The difficulty with most of our modern work on attention is that it inverts the relationship between cause and effect. Most

modern work takes a causal position: attention causes faster reaction times or enhanced perceptual processing. I assert attention should be treated as an effect. When we do X in the laboratory and see that responses become faster we might, for convenience, label such an observation (when it is not explained by differences at the level of primary sensory or motor systems) as "attentional." But the important question is what was it about the experimental situation that produced this attentional effect?

The effect interpretation of attention has never been the popular orientation of psychologists when they design and interpret their experiments, but it has been long recognized. In fact, in the James' quote that begins this paper, James is summarizing the effect position as a prelude to attacking it. Others have treated the idea more positively (Johnston and Dark, 1986). Hebb (1949), in the same work so often quoted for its foresight on synaptic learning mechanisms, writes "When an experimental result makes it necessary to refer to. . .'attention,' the reference means, precisely, that the activity that controls the form, speed, strength, or duration of response is not the immediately preceding excitation of receptor cells alone. The fact that a response is not so controlled may be hard to explain, theoretically; but it is not mystical. . ."

The false attribution of causal agency to an abstract concept is known as reification (sometimes also called hypostatization). Reification is used effectively in many literary constructions where the technique is seen to provide an economical and evocative phrasing that few would read literally. For example, "Love conquers all." In scientific parlance reification is both less obvious, and more pernicious. And it is clearly at play in the way we most frequently see the word attention used. It is not difficult to find quotes like Spitzer et al. (1988) "It is concluded that increasing the amount of attention directed toward a stimulus can enhance the responsiveness and selectivity of the neurons that process it." or Colby and Goldberg (1999) "In other words, paying attention to information in the receptive field drives the neuron no matter what direction of saccade the monkey is planning." We may deride theories relying on homunculi, but we employ them when ever we use the term "attention" for an unspecified causal agent. The fact that many of uses of the term are vacuous can be demonstrated by simply deleting the term and seeing whether the explanatory content is significantly reduced. For example, when Treisman (1985) writes: "Some discriminations appear to be made automatically, ~~without attention~~ and spatially in parallel across the visual field. Other visual operations ~~require focused attention and~~ can be performed only serially." the references to attention can be struck out without losing any understanding of the empirical results, and their inclusion doesn't deepen our theoretical understanding; clearly something explains the empirical differences, but the word "attention" as used here is just a place holder for that "something." We need to move beyond rhetorical accounts to more specific causal theoretical accounts. We need to know what is it in a subject's experience that changes reaction time slopes or what it is about the construction of our stimuli that enhances the responsivity, and selectivity of some neurons. We should not be satisfied with the insertion of the term attention as a theoretical wildcard. We should be particularly careful of this practice when evaluating the theoretical import of neuroscience data. Because data on firing rates and BOLD activations are "hard" data, we feel more secure with

them, and we may fail to realize that *by themselves* neuroscience data tell us nothing about what attention is.

NEUROSCIENCE, THEORY, AND ATTENTION

> But psychology may find it dangerous to turn to neurology for help. Once you tell the world that another science will explain what your key terms really mean, you must forgive the world if it decides that the other science is doing the important work.
>
> Skinner (1987).

The problem of dichotomies and reification, though mostly illustrated by examples from psychology, are similarly present in neuroscience studies. For example, when multiple visual elements are simultaneously present on a computer display, the investigator may speak of a cue directing attention to one of these and not the others ("To address this question, we trained monkeys to covertly deploy their visual attention from a central fixation point to one of three objects displayed in the periphery. . ." (Zhang et al., 2011). This language demonstrates the use of dichotomy; attention is either here or there, and not, for instance, some continuous distribution over space or objects. Also, it reifies attention, speaking of it as a thing that can be deployed. Instead of developing hypotheses in terms of the experimental variables, the effects of cues on neural firing rates, the authors interpose an explanatorily empty term: attention.

This is not simply someone else's problem. Anderson and Sheinberg (2008) manipulated the timing of targets while recording from anterior inferior temporal (aIT) neurons. aIT neurons fired more when a visual target was presented at the more likely time. An easy account of these data is to report (as I did) the "effects of temporal attention on spike rates." But, to be precise, Anderson and Sheinberg (2008) does not report a manipulation of attention. It reports a manipulation of the lag between cues and visual targets, and the validity of the cues for predicting delay. What is gained by inserting the intermediary attention? Would not it be more direct to assert that there are effects of prior probability on firing rates since the manipulation of prior probability was a concrete fact of the experimental design? This would lead to a prediction that the change in firing rate should be proportionate to the change in the predictive validity of the temporal cues. Such clear predictions do not emerge if we are satisfied with asserting the presence of attention, as a sort of mental phlogiston, that makes neat work of our results. Perhaps because their data are "hard," neuroscientists seem less occupied with these semantic issues, but data cannot disambiguate a reified term. Without a clear, predictive account of what behavioral situations produce attentional effects, and how they occur, our data collection will merely be piling up stones. There is no EKG long enough to tell us what love is, nor can any number of spike trains tell us what attention is.

Neuroscience data is the only type of data that will ever be able to characterize the mechanisms by which brains yield attentional phenomena. But without strong theoretical motivations for our experiments and data analyses, we risk making spurious causal claims. We do not want to do the neuroscience equivalent of concluding that pneumonia drives an increase in bacterial division because pneumonia-in cases always have higher bacterial counts than pneumonia-out cases. Without a theory of pneumonia by

which to interpret the significance of changes in bacterial counts we risk this sort of error. Without a theory of attention by which to interpret neural activity we risk the same sort of error. Computational models of attention that are framed without recourse to an attentional black box can serve this purpose.

The term attention developed colloquially to describe an introspectively accessible human experience. It was drafted as a technical term to characterize a pattern of behavioral observations. For neither colloquial nor technical use can we determine the nature of attention from its neural correlates. Characterizing the terminology of human behavior and experience is the domain of psychology. Cognitive neuroscience contributes once a consistent experimentally meaningfully terminology has been agreed. Neuroscience will tell us much of what we want to know about attention, but only once it has been yoked to a robust theory of what it is about particular experimental conditions that yields attentional effects. The limitations of neuroscience are not technical, but epistemic. Reviewing the vast and interesting neuroscience literature on attention will not advance us toward our goal of developing a terminology for attention that is experimentally and theoretically fecund.

BIASED COMPETITION: SPOTLIGHT ON A NEW METAPHOR

The biased competition model of attention is sometimes offered as a new theory of attention that reflects our growth in knowledge. However, it is really a new metaphor and not a new theory. Biased competition captures two cardinal phenomena of visual attention: limited capacity and selectivity. From these are posited the general schema where elements (visual objects or pools of neurons) compete for control. The competition is biased in favor of the objects that are most behaviorally relevant (Desimone and Duncan, 1995). From a neural perspective, biased competition is a theory of implementation. How do neurons adjudicate their joint response? Through a system of competition and modification that yields a winner take all result. From a psychological perspective, biased competition is a summary label for attentional phenomena. We have known for at least a century that a research subject's report can be adjusted by instruction (Helmholtz, 1881). Using biased competition to summarize this idea is a succinct terminology for an attentional effect, we have a change in awareness that cannot be ascribed to a change in receptor stimulation (Hebb, 1949). But having characterized the basic experimental data thus, we have not advanced theory. This does not seem to be a conclusion that is at odds with recent summaries of the biased competition concept (Duncan, 2006). After beginning with a quote from Wittgenstein, Duncan (2006) states that attention is probably undefinable. Duncan writes that "[while]. . .'attentional' phenomena. . .have family resemblances, it seems unlikely that they share any one defining component. . ." In general, one does not put scare quotes around words for which one is pronouncing a theory. Like the spotlight metaphor, biased competition nicely summarizes and emphasizes key aspects of the experimental phenomena. Duncan writes "At least in some form, these ideas [biased competition] are implicit in any reasonable account of attentional limits."

It is not a criticism of biased competition to point out that it is not a theory of attention. Just as the spotlight metaphor has its role as a convenient summary of data and as a way to structure our

intuitions, so does the metaphor of biased competition. The central question if we want to construct a causal theory of attention on the basis of biased competition is to determine what constitutes bias? What does it mean to be relevant? How do we avoid the circularity of asserting that behavioral relevance produces bias and that observing bias indicates the relative relevance of simultaneously present stimuli? In considering whether we have learned anything new about attention via the terminology of biased competition, we must ask ourselves whether, if we knew the answer to Duncan's question: ". . .how does the brain establish what is relevant and what is not?", we would have any residual need for the terms bias or attention? The ideas of biased competition would remain but once we can describe behavior, perception, and neural activity from ". . .the antecedents of the attending process in the individual himself and in the material that offers itself from the outside world." (Pillsbury, 1922) we will have eliminated the need for a theory of attention at all.

WHAT AN ALTERNATIVE TERMINOLOGY OF ATTENTION SHOULD SOUND LIKE

To copy Duncan in copying Wittgenstein: "*Wovon man nicht sprechen kann, darüber mußman schweigen.*"[2] We need the right terms if we are to say something meaningful, and, to this point, I have not offered such a terminology. My argument has been primarily critical. I have claimed that our use of the word attention as referring to a particular individual thing that directly causes changes in perception and neuron firing is flawed. I would like to conclude on a more constructive note by offering an alternative conceptual foundation that does not dichotomize or reify, and yields a terminology in which attentional effects are seen to emerge as a consequence of specific and experimentally manipulable factors. To be useful experimentally, our terminology should predict how quantitative variation will translate into behavioral and neural effects. A terminology that meets all these requirements and which has been increasingly used in psychology is the Bayesian terminology. Jones and Love (2010) assert that "Bayesian analysis can serve as a useful starting point when investigating a new domain. . ." This can be extended to the claim that Bayesian analysis is useful for attention if we change "a new" to "anew a." Early applications of Bayesian ideas to the study of attentional phenomena reveal the utility of this language.

BAYESIAN DECISION ACCOUNTS OF ATTENTIONAL PHENOMENA

A Bayesian decision process (BDP) is a computational structure that provides, as output, an estimate of the cost (or gain) to be expected from undertaking specified actions in an uncertain setting. As input the BDP needs some specification of the costs associated with particular actions in particular states, an estimate of the prior probability of those states, and some data from which it can compute the likelihood of that data assuming a particular state. These components map cleanly onto the experimental circumstances that are commonly employed in attentional research and they fit the requirement that they implicitly capture the common sense features required of any "reasonable account of attentional

---

[2]Whereof one cannot speak, thereof must one be silent.

limits" (Duncan, 2006). The mapping to the ideas of biased competition is direct. The likelihood function of the BDP establishes the basis for competition: how likely is what I am seeing, assuming any one of the possible causes? This competition is biased by expectations and the accumulation of evidence over time. From experience the subject has an estimate of how likely are the possible causes. Those things which happen more often receive a bias in proportion to their relative frequencies. The notion of relevance is further handled by the cost/gain function. Things are important because they yield high rewards or run substantial risks. The result of this computation is a decision. Our perception is the result of this covert deliberation. The fuzzy concepts of relevance and bias implicit in instructions and history are made numerical and unambiguous in the BDP framework.

Bayesian decision process components also map well onto experimental variables. This makes our common experimental approaches amenable to objective assessment within the BDP terminology. As an example, consider the generic experiment where a participant has to report some aspect of a visual stimulus (such as presence versus absence or brightness). We can manipulate how often different classes of stimuli appear, how certain are environmental features (i.e., cues) to predict targets, and how much reward, either in juice, money, or auditory feedback, we provide for correct responses (and how we punish errors, too). These are the sorts of experiments we use to demonstrate attentional effects. These manipulations change prior probability, likelihood, and cost/benefits.

Experimental results confirm the relevance of these components. For example, visual search tasks are commonly used to assess attention. According to the BDP formulation, locations of high probability should have a higher search priority (Koopman, 1956). In a classic result, Shaw and Shaw (1977) confirmed this finding. In their experiment subjects searched for a letter after learning the probabilities for letter locations. The conditional probability for detecting stimuli given their location was, for three of four human subjects tested, consistent with the model of optimal performance that allocated search resources in proportion to probability.

What is contextual cuing (Chun, 2000) but a demonstration that visual search prioritizes location based on prior probability? The same basic result has also been given an explicitly Bayesian formulation by Eckstein et al. (2006). Their subjects looked at pictures with elements in expected or unexpected locations or, critically, absent. The pattern of first saccades in target absent images was directed toward probable locations and was well described by a differential, Bayesian, weighting model.

The likelihood function in a Bayesian formulation can be captured by the idea of noise and reliability. When uncertainty increases visual detection declines. For example, Lasley and Cohn (1981) briefly illuminated an LED and manipulated the number of non-overlapping temporal intervals in which the stimulus could appear. Stimulus discriminability, reported as d', decreased monotonically with increasing stimulus uncertainty. A likelihood function provides a link to data on visual image discriminability and target ambiguity (Duncan and Humphreys, 1989).

Bayesian probability models are increasingly common in psychology (Jones and Love, 2010), and are beginning to be used

for predicting attentional effects. For example, salience, a common term for the tendency of some regions of space to attract gaze, has often been stated in the past to attract attention. Najemnik and Geisler (2005) and Zhang et al. (2008) are two recent examples that provide an alternative, Bayesian account. These results demonstrate how Bayes' formula can be used to relate the conditional probabilities of causes given consequences to the probabilities of consequences given causes. Such Bayesian models show the power of this approach and can be used, for example, to compute the probability of a target's location from probability distributions for relevant features, locations, knowledge about how likely particular features are for particular targets, and the statistics of natural images. These types of models provide specific predictions for eye movement trajectories that can be directly compared to the human eye movement data from inspecting identical images. As these sorts of models give a good account of where we look and why we look there in terms of the properties of our environment and our experiences, why do we need the term attention?

The incorporation of a cost function into a Bayesian model to turn it into a BDP is less common, but greatly enriches the potential of a Bayesian formulation to account for attentional phenomena. The cost function is the part of the BDP terminology that captures and quantifies the metric for behavioral relevance that is at the heart of the idea of biased competition. Where there is an objective measure of cost or gain, attentional metrics can be predicted to align with actions that prioritize the detection of relevant events: those at the extremes of benefit and risk.

Another strength of the BDP formalism is that it describes how to integrate these different components. If you manipulate the prior probability of stimuli (e.g., by manipulating the predictive validity of antecedent cues), and you offer different rewards for different detections, then you can predict the order of your response measures to each manipulation alone and in combination (Liston and Stone, 2008). The combination of these factors can explain data better than any one factor. Milstein and Dorris (2007) demonstrated this idea when they had people saccade to one of two targets while varying the probability of the target side. They also varied the reward associated with targets on each side and found that the combination of monetary value and probability correlated negatively with response time (greater value faster response). Directly linking the idea to attentional phenomena, the authors occasionally displayed a distractor on the screen at varying distances from one of the target locations and found that the distractor was more likely to capture a saccade, in this task an error, when it was near a high value x high probability target. The combination of these factors was a better explanation than any one of them (Milstein and Dorris, 2007).

The possibility of a circular argument, introduced into the biased competition account by reference to the concept of relevance, is also a risk when using the BDP terminology. It would be tempting to say that if all options were equally likely and equally discriminable, then the class of response chosen most quickly or most often must, therefore, be the one with the greatest benefit or highest relevance. The way to avoid this problem is to construct an objective metric for quantifying the basis for such a preference

and then to test this idea in an independent setting. As an example, it has been suggested that novelty contributes to relevance. Therefore, novel items in displays of familiar distractors pop-out (Strayer and Johnston, 2000; Johnston and Strayer, 2001). But what does it mean to be novel, how novel is "novel?" Itti and Baldi (2009) have developed an objective Bayesian approach to quantification. Itti and Baldi (2009) compare our estimate of the probability of a hypothesis before and after we observe some data. If our estimate changes a lot then the data were probably surprising. The updating mechanism for converting a prior probability into a posterior uses Bayes' formula applied to the prior and the likelihood distributions. The Kullback–Leibler distance, a way of measuring how far apart two probability distributions are, provides the metric for comparing the old prior and the new posterior to give a *wow*, the measure of how "surprising" were the data. Itti and Baldi (2009) find a good concordance between where people look and this metric of surprise.

Again, an advantage of the BDP approach, something that is difficult to do with the reified view of attention as a cause, is to combine components cleanly. If the prior probability of informative areas of a visual display were varied, and if spatially inhomogeneous noise were applied to the display, then all three factors (prior, likelihood, and surprise) could be combined and compared in a single experiment. The BDP formula provides the mathematical relationship; since all the components are probability distributions they are all commensurate, the combination yields an unambiguous prediction.

One of the attractions of the biased competition account of attention is the natural way it connect psychological ideas with neural implementations. While I have argued that neural data are not directly relevant to the debate about what conception of attention is the right one, cause or effect, a valuable element of the BDP terminology is that it has the qualities of being a good theoretical guide for neuroscientific investigations of attentional effects. Neuronal tuning curves can be easily imagined as measures of stimulus likelihood (Ma et al., 2006), neuronal populations can be said to fire in proportion to estimates of probability (Janssen and Shadlen, 2005) and other neuronal populations give an index

of the expected value of an action (Padoa-Schioppa and Assad, 2006). Furthermore, the mechanisms of synaptic plasticity provide the basis for learning these distributions. More probable stimuli lead to greater synaptic modifications. A BDP account can provide a psychological explanation for attention that emerges from the activities of neurons. A psychological account of attention in the language of the BDP offers an easy route to extending the ideas to computational modeling (Yu et al., 2009; Chikkerur et al., 2010). While these examples point a way forward, Bayesian formulations of attentional phenomena cannot be regarded as mainstream. The word "Bayes" does not appear in either a recent review of visual attention (Carrasco, 2011) or an elaboration of one of the more long lived and well regarded models of attention (Bundesen et al., 2011).

CONCLUSION

Early on I referenced Newell (1973) and his famous 20 questions essay. In that essay Newell develops prescriptions for what psychological explanations of psychological phenomena should look like. Newell suggests that we should build complete processing models rather than partial ones, second, we should either choose for analysis one big complex task or show that our single processing model accounts for numerous smaller tasks. A BDP is a complex model for a complex task that can be addressed by application to numerous smaller, more focused, experiments.

My resort to an attack on our terminology might be classed as an example of the poor workman blaming his tools, but as I have sampled the literature from the last 100 years, I have been struck by the creativity and intelligence of attention researchers. If their attack has not been successful, it is not for a failure of skill or ingenuity, rather it is because we need new tools. Refashioning the old tools simply will not do. To the man with a hammer not only does everything look like a nail, but such a man is equally likely to use his hammer as a shoehorn, door stop, backscratcher, and pounding square pegs into round holes. We need to discard our causal conceptions of attention and adopt effect accounts. A process model, like a BDP, is one reasonable starting point for us to begin anew.

REFERENCES

Allport, A. (1993). *Attention and Control: Have We Been Asking the Wrong Questions? A Critical Review of Twenty-Five Years. Attention and Performance XIV: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience.* Cambridge, MA: MIT Press, 183–218.

Anderson, B., and Sheinberg, D. L. (2008). Effects of temporal context and temporal expectancy on neural activity in inferior temporal cortex. *Neuropsychologia* 46, 947–957.

Angell, J. R., and Pierce, A. H. (1892). Experimental research upon the phenomena of attention. *Am. J. Psychol.* 4, 528–541.

Balcetis, E., and Dunning, D. (2006). See what you want to see: motivational

influences on visual perception. *J. Pers. Soc. Psychol.* 91, 612.

Brignani, D., Guzzon, D., Marzi, C., and Miniussi, C. (2009). Attentional orienting induced by arrows and eye-gaze compared with an endogenous cue. *Neuropsychologia* 47, 370–381.

Bundesen, C., Habekost, T., and Kyllingsbaek, S. (2011). A neural theory of visual attention and short-term memory (NTVA). *Neuropsychologia* 49, 1446–1457.

Carrasco, M. (2009). Cross-modal attention enhances perceived contrast. *Proc. Natl. Acad. Sci. U.S.A.* 106, 22039.

Carrasco, M. (2011). Visual attention. *Vision Res.* 51, 1484–1525.

Cattell, J., and Farrand, L. (1896). Physical and mental measurements of

the students of Columbia University. *Psychol. Rev.* 3, 618–648.

Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *J. Acoust. Soc. Am.* 25, 975–979.

Chikkerur, S., Serre, T., Tan, C., and Poggio, T. (2010). What and where: a Bayesian inference theory of attention. *Vision Res.* 50, 2233–2247.

Chun, M. M. (2000). Contextual cueing of visual attention. *Trends Cogn. Sci. (Regul. Ed.)* 4, 170–178.

Colby, C. L., and Goldberg, M. E. (1999). Space and attention in parietal cortex. *Annu. Rev. Neurosci.* 22, 319–349.

Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222.

Dodd, M. D., and Wilson, D. (2008). Training attention: interactions between central cues and reflexive attention. *Vis. Cogn.* 17, 736–754.

Drew, T., McCollough, A., Horowitz, T., and Vogel, E. (2009). Attentional enhancement during multiple-object tracking. *Psychon. Bull. Rev.* 16, 411–417.

Duncan, J. (2006). Brain mechanisms of attention. *Q. J. Exp. Psychol.* 59, 2–27.

Duncan, J., and Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychol. Rev.* 96, 433–458.

Eckstein, M. P., Drescher, B. A., and Shimozaki, S. S. (2006). Attentional cues in real scenes, saccadic targeting, and Bayesian priors. *Psychol. Sci.* 17, 973–980.

Fernandez-Duque, D., and Johnson, M. L. (1999). Attention metaphors: how metaphors guide the cognitive psychology of attention. *Cogn. Sci.* 23, 83–116.

Hatfield, G. (1998). "Attention in early scientific psychology," in *Visual Attention*, ed. R. D. Wright (New York: Oxford University Press), 3–25.

Hebb, D. (1949). *The Organization of Behavior: A Neuropsychological Theory*. New York, NY: Wiley.

Helmholtz, H. V. (1881). *Popular Lectures on Scientific Subjects*. New York: D. Appleton.

Huang, L., and Pashler, H. (2007). A boolean map theory of visual attention. *Psychol. Rev.* 114, 599–631.

Itti, L., and Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Res.* 49, 1295–1306.

James, W. (1890/1950). *The Principles of Psychology*, Mineola, NY: Dover Publications.

Janssen, P., and Shadlen, M. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. *Nat. Neurosci.* 8, 234–241.

Johnston, W., and Strayer, D. (2001). "Is a dynamic, evolutionary perspective on attention capture1," in *Attraction, Distraction and Action: Multiple Perspectives on Attentional Capture*, eds F. Charles and G. Bradley (Amsterdam: Elsevier), 375.

Johnston, W. A., and Dark, V. J. (1986). Selective attention. *Annu. Rev. Psychol.* 37, 43–75.

Jones, M., and Love, B. (2010). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behav. Brain Sci.* 34, 169–188.

Koopman, B. O. (1956). The theory of search. II. Target detection. *Oper. Res.* 4, 503–531.

Lasley, D. J., and Cohn, T. (1981). Detection of a luminance increment: effect of temporal uncertainty. *J. Opt. Soc. Am.* 71, 845–850.

Liston, D. B., and Stone, L. S. (2008). Effects of prior information and reward on oculomotor and perceptual choices. *J. Neurosci.* 28, 13866–13875.

Logan, G. D. (1992). Attention and preattention in theories of automaticity. *Am. J. Psychol.* 105, 317–339.

Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438.

Milstein, D. M., and Dorris, M. C. (2007). The influence of expected value on saccadic preparation. *J. Neurosci.* 27, 4810–4818.

Mole, C. (2011). *Attention is Cognitive Unison: An Essay in Philosophical Psychology*. NewYork, NY: Oxford University Press.

Najemnik, J., and Geisler, W. (2005). Optimal eye movement strategies in visual search. *Nature* 434, 387–391.

Newell, A. (1973). "You can't play 20 questions with nature and win: projective comments on the papers of this symposium," in *Visual Information Processing*, ed. W. G. Chase (New York: Academic Press), 283–308.

Padoa-Schioppa, C., and Assad, J. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature* 441, 223–226.

Pillsbury, W. (1922). *The Fundamentals of Psychology (Revised Edition)*. New York, NY: MacMillan Co.

Prinzmetal, W., Presti, D. E., and Posner, M. I. (1986). Does attention affect visual feature integration? *J. Exp. Psychol. Hum. Percept. Perform.* 12, 361–369.

Reynolds, J., Pasternak, T., and Desimone, R. (2000). Attention increases sensitivity of v4 neurons. *Neuron* 26, 703–714.

Ristic, J., Friesen, C., and Kingstone, A. (2002). Are eyes special? it depends on how you look at it. *Psychon. Bull. Rev.* 9, 507–513.

Shaw, M. L., and Shaw, P. (1977). Optimal allocation of cognitive resources to spatial locations. *J. Exp. Psychol. Hum. Percept. Perform.* 3, 201–211.

Skinner, B. F. (1987). Whatever happened to psychology as the science of behavior? *Am. Psychol.* 42, 780–786.

Spitzer, H., Desimone, R., and Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science* 240, 338–340.

Stewart, D. (1792/1866). *Elements of the Philosophy of the Human Mind*. Boston: William H. Dennet.

Strayer, D., and Johnston, W. (2000). Novel popout is an attention-based phenomenon: an ERP analysis. *Atten. Percept. Psychophys.* 62, 459–470.

Theeuwes, J., Kramer, A. F., and Atchley, P. (1999). Attentional effects on preattentive vision: spatial precues affect the detection of simple features. *J. Exp. Psychol. Hum. Percept. Perform.* 25, 341–347.

Titchener, E. B. (1908). *Lectures on the Elementary Psychology of Feeling and Attention*. New York, NY: MacMillan Co.

Treisman, A. (1985). Preattentive processing in vision. *Comput. Vis. Graph. Image Process.* 31, 156–177.

Treisman, A., and Gelade, G. (1980). A feature-integration theory of attention. *Cogn. Psychol.* 12, 97136.

Wolfe, J. M. (1998). What can 1 million trials tell us about visual search? *Psychol. Sci.* 9, 33–39.

Yantis, S., and Jonides, J. (1990). Abrupt visual onsets and selective attention: voluntary versus automatic allocation. *J. Exp. Psychol. Hum. Percept. Perform.* 16, 121–134.

Yu, A., Dayan, P., and Cohen, J. (2009). Dynamics of attentional selection under conflict: toward a rational Bayesian account. *J. Exp. Psychol. Hum. Percept. Perform.* 35, 700.

Zhang, L., Tong, M., Marks, T., Shan, H., and Cottrell, G. (2008). Sun: a Bayesian framework for saliency using natural statistics. *J. Vis.* 8, 1–20.

Zhang, Y., Meyers, E., Bichot, N., Serre, T., Poggio, T., and Desimone, R. (2011). Object decoding with attention in inferior temporal cortex. *Proc. Natl. Acad. Sci. U.S.A.* 108, 8850.