

Top-down versus bottom-up attentional control: a failed theoretical dichotomy

Edward Awh¹, Artem V. Belopolsky² and Jan Theeuwes²

¹ Department of Psychology, University of Oregon, Eugene, OR 97403, USA

² Cognitive Psychology, Vrije Universiteit, Van der Boechorststraat 1, 1081 BT Amsterdam, The Netherlands

Prominent models of attentional control assert a dichotomy between top-down and bottom-up control, with the former determined by current selection goals and the latter determined by physical salience. This theoretical dichotomy, however, fails to explain a growing number of cases in which neither current goals nor physical salience can account for strong selection biases. For example, equally salient stimuli associated with reward can capture attention, even when this contradicts current selection goals. Thus, although ‘top-down’ sources of bias are sometimes defined as those that are not due to physical salience, this conception conflates distinct – and sometimes contradictory – sources of selection bias. We describe an alternative framework, in which past selection history is integrated with current goals and physical salience to shape an integrated priority map.

The ‘top-down’ versus ‘bottom-up’ dichotomy is inadequate

A typical environment contains far more information than we can process at one time. Thus, goal-driven perception and action depend on attention to direct limited resources towards a subset of relevant items. The most prominent models of attentional control have described attentional control in terms of a dichotomy between top-down and bottom-up control (sometimes referred to as ‘endogenous’ and ‘exogenous’ control, respectively – see [Glossary](#)), with the former determined by the current goals of the observer and the latter determined by the physical characteristics of the scene in question [1–10]. Thus, whereas attention can be voluntarily allocated towards items in a goal-driven manner (e.g., when you are searching for your friend’s red jacket in a crowd), attention can also be captured involuntarily by physically salient stimuli (e.g., a loud noise or sudden movement).

Nevertheless, despite the dominance of this theoretical dichotomy, it has become increasingly clear that the presumed equivalence between ‘top-down’ attentional control and ‘goal-driven’ selection yields a significant explanatory gap. Specifically, a growing body of evidence documents cases in which strong selection biases cannot be explained by current selection goals or the physical salience of potential targets. We highlight two broad classes of selection phenomena where this problem is evident. First, we

consider studies that have shown how the recent history of attentional deployments can elicit lingering selection biases, even though such ‘selection history’ effects may be unrelated to current goals or the physical salience of items competing for selection. Second, we consider the fact that visual selection is biased towards items associated with previous reward, even though ‘reward history’ can also be dissociated from current goals and stimulus salience. Thus, although it is tempting to presume a ‘top-down’ form of attentional control when differences in physical salience can be ruled out, this approach conflates distinct and sometimes contradictory sources of selection bias. Thus, the top-down versus bottom-up dichotomy is an inadequate taxonomy of attentional control.

To provide a more comprehensive taxonomy, we propose a theoretical framework that extends the well-known construct of a priority map that integrates multiple selection

Glossary

Bottom-up attentional control: attentional control that is driven by factors external to the observer, such as stimulus salience (e.g., ‘pop-out’ stimuli that contrast strongly with surrounding items based on a simple feature value, sudden flashes of light, or loud noises in an otherwise quiet environment). We view this as the same concept as ‘exogenous attentional control’.

Goal-driven selection: the imposition of a selection bias based only on the current selection goals of the observer. This definition of goal-driven selection excludes selection biases that are a lingering consequence of past selection episodes or goals, because it is possible for such influences to countermand current selection goals.

Selection bias: an early perceptual bias towards a specific defining feature, such as color or location, such that stimuli with that feature are prioritized over other stimuli during initial encoding. This can be distinguished from the ability to render a stimulus-specific response in the absence of a selection bias. For example, without imposing a selection bias, an observer could search for a specific target in a search array by evaluating each item in turn until the target-defining properties are encountered.

Selection history: the bias to prioritize items that have been previously attended in a given context. Because such selection history effects may contradict current selection goals, we argue that selection history and current goals should be viewed as distinct categories of control.

Stimulus salience: the degree to which a stimulus is likely to attract attention based on its low-level properties and independently of the internal mental state of the observer. This is the driving force behind bottom-up or exogenous attentional control.

Top-down attentional control: attentional control that is driven by factors that are ‘internal’ to the observer. We view this as the same concept as ‘endogenous attentional control’. The key problem we highlight with this construct is that grouping together control signals that are ‘internal’ (i.e., control signals unrelated to stimulus salience) conflates the effects of current selection goals and selection history. Because current goals and selection history may generate conflicting selection biases, we argue that they should be viewed as distinct categories of control.

Corresponding author: Awh, E. (awh@uoregon.edu).

influences. In this case, however, we emphasize that priority is determined not just by goal- and stimulus-driven selection, but also by the lingering effects of past selection episodes (e.g., reward and selection history). This framework may provide a more productive platform for categorizing and elucidating diverse forms of attentional control.

Selection phenomena that defy the top-down versus bottom-up dichotomy

Below, we review two broad classes of selection phenomena that are not explained by physical stimulus properties or voluntary goal-driven selection. We note, however, that a full account of this class of phenomena is beyond the scope of this review. Instead, our goal is to provide concrete examples that underscore the limitations in the standard theoretical dichotomy for attentional control, and set the stage for a more comprehensive framework to guide future research.

Selection history

The modern study of attention was inspired by studies of spatial orienting that demonstrated faster processing at attended relative to unattended locations [1,11]. Soon afterwards, the distinction between voluntary and automatic orienting of attention was highlighted by Jonides [2]. He contrasted the impact of symbolic cues (thought to elicit voluntary, goal-driven shifts of attention) and peripheral cues (thought to elicit automatic, stimulus-driven shifts of attention), and found that symbolic cues were more susceptible to interference from a secondary task, took longer to elicit changes in selection, and were more readily controlled by changes in voluntary selection goals. These data demonstrated the critical distinction between voluntary and automatic control over attention, a perspective that has remained dominant through the present day [1–10].

That said, whereas the contrast between goal-driven and stimulus-driven modes of attentional control has the virtue of highlighting the fundamental difference between voluntary and involuntary control over attention, this dichotomy may have the undesired side effect of obscuring a third category of selection biases that are unrelated to current goals and the physical salience of the items competing for selection. In turn, this explanatory gap can obscure the interpretation of empirical patterns that have been attributed to top-down attentional control. For example, consider a study by Wolfe and colleagues [12] that compared performance when the target-defining feature in a search task varied from trial to trial (i.e., between color and orientation) with performance in pure blocks in which a single feature value defined the target for the entire block (e.g., target was always ‘red’ or ‘vertical’) [12]. The results showed that participants were about 80 ms slower in the mixed condition. Because the individual displays were identical between the pure and mixed conditions, Wolfe *et al.* concluded that ‘top-down information makes a substantial contribution to [reaction times] even for the simplest of feature searches’ ([12], p. 485). The typical explanation of this empirical pattern is that goal-driven selection is enhanced when top-down resources can be focused on only one dimension instead of two; thus, the

strength of the top-down signal is higher in the pure blocks than in the mixed blocks and search is guided more efficiently. This idea of top-down modulation of weights is at the heart of many visual search models and accounts, such as Guided Search [5], Dimensional Weighting [13], Theory of Visual Attention (TVA) [14], Feature-Based Attention [15] and Attentional Engagement Theory [16].

Nevertheless, although the Wolfe *et al.* findings are consistent with the claim that goal-driven control signals are stronger in the pure blocks, it is also possible that advantages during pure blocks were explained by the involuntary consequences of selection history. That is, if the selection of a given feature value (e.g., ‘vertical’) is primed by the recent selection of the same value, selection might be stronger in pure blocks because of a cumulative effect of inter-trial priming rather than because of enhanced goal-driven orienting. A classic demonstration of such a priming effect was provided by Maljkovic and Nakayama [17], who showed that, when a search target was defined by a given feature (e.g., color or spatial frequency), search for the same feature was more efficient on subsequent trials, an effect that has been labeled ‘priming of pop-out’ [18–20]. Importantly, this priming effect was robust even when it contradicted the observer’s voluntary goal to select a different feature value. In line with this finding, Theeuwes and Van der Burg [21] measured search reaction time (RT) with displays that contained two equally salient color singletons (e.g., a red and a green circle) and found that participants could not use advance cues to voluntarily override the interference from the irrelevant singleton. Instead, interference from the irrelevant singleton was eliminated only on trials where the selected feature value matched the value from the previous trial. These findings dovetail with the priming of pop-out phenomenon to show that passive priming effects can yield strong selection biases that are unrelated to the observer’s current selection goals [22–25]. Thus, pure versus mixed block designs may confound the putative effects of goal-driven and history-driven selection.

Indeed, although designs employing trial-to-trial variations in the cued position have established that spatial selection can be purely goal-driven rather than based on selection history (e.g., [2]), some recent research challenges the efficacy of goal-driven selection for non-spatial features [26]. For example, Theeuwes *et al.* gave participants informative cues regarding the likely defining feature of an upcoming target singleton. For example, if the word ‘red’ was presented as a cue, observers knew with 80% certainty that the target would be a red circle. For the purposes of the present discussion, the key point was that the cued feature varied randomly from trial to trial, such that goal-driven effects could be distinguished from the effects of bottom-up priming. Critically, neither reaction time nor perceptual sensitivity [21,27] benefitted from the advance cue. Thus, goal-driven selection effects can be elusive when inter-trial priming effects are eliminated. This suggests that future efforts to demonstrate goal-driven selection of non-spatial features should focus on isolating putative goal-driven effects from the known consequences of selection history (Box 1). These comments are also relevant for the interpretation of various neural

Box 1. Contingent capture

The distinction between voluntary and automatic control over attention was acknowledged only shortly after the advent of modern laboratory studies of attention. Jonides [2] used central arrow cues to elicit voluntary shifts of spatial attention, and peripheral cues with abrupt visual onsets to elicit involuntary shifts of attention. However, Folk *et al.* [4] challenged the claim that attention could be captured in a fully bottom-up fashion. In their ‘contingent capture’ paradigm, observers responded – in separate blocks of trials – to a target that was defined either by its color or by its status as an abrupt visual onset. Just prior to the onset of the target display, a cue display was presented in which a single location was marked by either a color singleton or an abrupt visual onset. The critical finding was that the color cues only captured attention when the target was defined by color and abrupt onsets only captured attention when observers were expecting onset targets. Thus, Folk *et al.* concluded that attentional capture is not wholly automatic, but is contingent on the observer’s top-down attentional settings. This seminal finding highlights the fact that attentional capture is determined by an interplay between the physical display and the internal control settings of the observer [78,79].

A key goal is to understand the nature of these ‘internal control settings’. The modal interpretation of contingent capture effects is that observers’ voluntary goals determine whether or not a given stimulus will capture attention. However, the blocked manipulation of the target-defining property also allowed for strong effects of selection history. For example, red items may have captured attention because of priming from recent trials in which observers had selected and responded to red items. Belopolsky *et al.* [80] provided proof for this explanation by showing that selection history alone was capable of generating the contingent capture empirical pattern, independently of the observer’s voluntary control settings. This suggests that future efforts to determine the role of volitional control in visual selection should focus on isolating the effects of current goals and selection history. We suggest two design elements that may help achieve this goal. First, trial-by-trial variations in the cued feature value enable a direct estimation of sequential priming effects, and thereby an opportunity to test whether goal-driven selection induces further biases in selection. Second, we recommend the use of symbolic cues that do not contain the voluntarily attended feature value to rule out bottom-up priming by the cue itself [30].

studies of feature-based attention that have found amplified stimulus-evoked responses for items that match the ‘top-down’ set for a non-spatial feature [15,28,29]. Because these studies have typically employed blocked manipulations of the target defining feature, it is unclear whether the observed selection biases were due to goal-driven selection or the lingering effects of selection history (Box 2).

Our central point about selection history is that it cannot be explained by the traditional dichotomy in which top-down and goal-driven control are synonymous. Some have claimed that such priming effects are top-down because they rely ‘on what the observer has learned about the prior trial and [not] solely on the state of the stimulus’ [12]. Others have maintained that these priming effects are completely bottom-up ‘because the effects cannot be counteracted by volitional top-down control’ [30]. Yet others have argued that bottom-up priming and top-down attentional set act upon different aspects of visual search [31]. Our position is that selection history effects call for a third category within the taxonomy of attentional control, because history effects can influence selection priority in a manner that is disconnected from both the observer’s current goals and the physical salience of the stimulus.

Reward history

In 1911, Thorndike described his ‘law of effect’: ‘Of several responses made to the same situation, those which are accompanied or closely followed by satisfaction to the animal will . . . be more likely to recur’ [32]. Although Thorndike was focused primarily on explaining the frequency of overt behaviors, the same principle is remarkably effective at predicting the orienting of selective attention. Across a broad range of behavioral and neural approaches, there is mounting evidence of selection biases towards previously rewarded items, even in the absence of explicit instructions [33–44]. Indeed, in the case of animal studies, the exclusive use of reward to control attentional orienting means that the putative effects of ‘attention’ in these studies cannot typically be distinguished from the effects of reward alone [45]. Studies explicitly focused on the neural substrates of reward processing further emphasize the parallel nature of reward and attention; activity in the lateral intraparietal sulcus – a brain region that has been presumed to play a primary role in the orienting of attention [46–50] – is directly modulated by reward contingencies [51–53]. In fact, although in this review we have chosen to respect the traditional boundaries in the literature between ‘selection’ and ‘reward’, these phenomena may be more intertwined than it appears at first glance. Specifically, the selection bias towards previously attended features or positions may be grounded in the reward that observers experience when they achieve their task goals [54]. Thus, a complete taxonomy of attentional control needs to account for reward-induced selection biases.

At first glance, it is natural to presume that the influence of reward is to shape the voluntary attentional state of the observer. According to the notion of incentive motivation [55], monetary gains can enhance perceptual and executive control processes to achieve more efficient goal-directed behavior. In this view, reward provides motivational significance, and motivational significance elicits all the known consequences of voluntary attentional orienting to the rewarded items. Recent evidence, however, has shown that the effect of reward on subsequent selection can also run counter to the voluntary selection goals of the observer [56–59]. For example, in a study by Hickey *et al.* [56] (Figure 1), participants were required to search for a popout diamond shape and ignore an irrelevant color singleton in the same display (see [60] for a review of this ‘additional singleton paradigm’). When observers responded correctly, they received either a high or a low monetary reward. The key finding was that these rewards determined whether selection was biased towards the target or the irrelevant singleton on subsequent trials. If the target color remained the same across trials, then target responses were faster following a high reward (the ‘no color swap’ condition); by contrast, if the distractor had the same color as the last target (the ‘color swap’ condition), reaction times were slower following a high reward, suggesting that the rewarded color captured attention regardless of whether it indicated a target or a distractor.

Crucially, Hickey *et al.* also examined the role of voluntary orienting by including another condition in which participants were rewarded for switching the attended

Box 2. Neural substrates of attentional control

Lesion studies, animal studies, and human neuroimaging studies have converged on the conclusion that top-down control over attention is mediated by a distributed network of regions that include frontal and parietal cortex [3,6,9]. Single-unit recording and human imaging studies have shown that neural activity in sensory regions is altered after attention-directing cues, such that processing is biased in favor of stimuli that contain the attended stimulus attributes (e.g., the cued position, color, orientation, etc.). In line with the goal-driven versus stimulus-driven dichotomy of attentional control, the modal interpretation is that this frontal-parietal network implements the voluntary selection goals of the observer by providing modulatory signals that bias processing in the relevant sensory areas. We note, however, that a large proportion of past studies cannot distinguish whether the observed sensory modulations were due to goal-driven selection *per se* or due to the lingering effects of reward and selection history. The typical animal study, for example, includes highly trained animals whose behavior has been shaped by months of reward feedback; in these cases there is often a confound between the sensory biases that would be imposed by current selection goals and selection history. This raises the intriguing possibility that the putative behavioral consequences of goal-driven selection – and the frontal-parietal activity presumed to generate those effects – may have been documented in some cases where voluntary selection played little or no role. Rossi and colleagues [81] demonstrated this in a powerful way by measuring attentional function in macaque monkeys following a unilateral resection of the prefrontal cortex. Strikingly, they found that ‘top-down’ selection remained almost as good in the affected hemifield as in the unaffected one. Deficits in behavioral performance only emerged when the animals were required to switch their attentional focus to a different target color. Although these data show that the prefrontal cortex is essential for reconfiguring attention based on changing task demands, they might appear to challenge the role of the PFC in goal-driven selection *per se*. We suggest that intact performance in the affected hemifield could have been entirely driven by selection and reward history, because these forces were completely congruent with the putative goals of the animal (until the target color switched). Although this explanation does not rule out the hypothesis that the PFC is critical for goal-driven selective attention, this study highlights the point that understanding the neural basis of goal-driven selection requires experimental designs that can distinguish between the effects of current selection goals and selection history.

color (e.g., switching the current selection goal from red to green) to receive a high reward on the next trial. Nevertheless, observers were slow on the trials in which the cued color changed relative to the last target color, but fast on those few trials where the color stayed the same. The authors concluded that reward automatically triggered a selection bias towards the rewarded color, even when the current goals of the participant were diametrically opposed.

Anderson *et al.* [58,59] also demonstrated a disconnect between reward-induced selection biases and the current selection goals of the observer. Rather than relying on trial-by-trial rewards, they used a pre-training procedure to associate high and low monetary rewards with a particular target color (see also [34,35]). Subsequently, a colored distractor associated with a high monetary reward slowed target processing more than one that was associated with a low monetary reward. Importantly, this result was obtained even though the reward contingency from the training phase had been eliminated, and the target never appeared in the previously rewarded color. Thus, learned reward value directly biases selection towards reward-associated stimulus features in an involuntary manner.

These dissociations between the effects of reward history and the observer’s current goals pose the same problem for the goal-driven versus stimulus-driven dichotomy as the selection history effects described above. That is, multiple studies have shown that selection is biased towards previously rewarded feature values, but that this effect cannot be explained by either the voluntary selection goals of the observer or the physical characteristics of the selected items. Thus, both selection and reward history effects highlight the point that equating ‘top-down’ control with goal-driven selection yields an incomplete taxonomy of attentional control.

Although we have focused on selection and reward history effects, it should be emphasized again that there are other classes of selection phenomena in which the observer’s past experience shapes subsequent selection without affecting voluntary selection goals [61–64]. Thus, although selection and reward history provide examples of the problem with the traditional dichotomy for attentional control, the implications for our understanding of attentional control are broader.

An integrative framework

Our proposal is that a modified taxonomy of attentional control may provide a more productive platform for understanding diverse sources of selection bias. Specifically, we highlight the well-known construct of a ‘priority map’ [5,48,65,66] that integrates three distinct categories of selection bias (Figure 2):

- (i) *Current goals.* This category of attentional control acknowledges the critical role of voluntary selection. Our conception of such goal-driven effects is in line with that of prominent models of attentional control [1–10], such that moment-to-moment changes in selection goals can elicit changes in selection bias. However, by explicitly distinguishing current goals from selection history effects, we believe that the expanded framework we describe here will make it clearer which empirical patterns can be unambiguously attributed to voluntary attentional control. This, in turn, will facilitate our understanding of the boundary conditions and specific effects of voluntary selection.
- (ii) *Selection history.* This category of control is intended to represent both classes of ‘history’ effect highlighted above (i.e., selection and reward history). In addition, there are multiple other examples in which the lingering effects of past experience shape the overall landscape of the observer’s selection biases. For example, likely targets are more likely to be encoded into working memory, even when observers lack explicit knowledge of this contingency [64]. Likewise, when specific configurations of distractors are associated with specific target positions, implicit knowledge of this association drives more efficient orienting to the target position [62]. Indeed, the full range of selection history effects is too large to cover comprehensively here, but we believe that they share one core feature: in each case, past selection episodes are recapitulated in subsequent trials when the relevant context is encountered again. As such, the selection




Figure 1. General paradigm and results from [56]. (a) The additional singleton paradigm of [82]. The target was a shape singleton (the circle) and a salient color singleton served as a distractor. The target and distractor colors switched ('Color Swap') or stayed the same ('No Color Swap') on subsequent trials. Following a correct discrimination of the target's orientation, participants either received 1 or 10 points (worth €0.2 or €2, respectively). (b) Reaction times for the target discrimination: when there was no color swap, participants were fast following a high reward; when there was a color swap, participants were slow indicating that the previously rewarded color now strongly captured attention even though it indicated a distractor. Adapted from [56].

history concept dovetails with existing frameworks that emphasize the role of associative learning and memory in the deployment of attention [67,68], except that we emphasize the numerous cases in which selection history effects countermand goal-driven selection.

- (iii) *Physical salience*. This category represents the well-documented fact that selection is sometimes biased in a manner that depends only on the properties of the stimulus display itself, rather than on the internal mental state of the observer [7,8,69]. Examples of physically salient stimulus events include visual pop-out (e.g., a red item amongst a field of green ones), sudden motion in an otherwise motionless visual field, or a loud noise in a quiet environment. However, although physical salience may appear to be the most straightforward of the three categories of attentional

control, there are interesting cases that highlight the grey areas at the border of this category. For instance, some evidence suggests that stimuli with high emotional valence capture attention [70–74]. Given the hypothesis that this is an evolutionarily adaptive response to ecologically important stimuli [73] and given that it can be observed using words composed of familiar characters [70], this kind of attentional capture may not be easily described in terms of low-level physical stimulus attributes *per se*; nevertheless, it has typically been described as a 'bottom-up' process. In fact, even stimuli that pop out due to low-level feature differences with surrounding stimuli reveal an interaction with the internal state of the observer (e.g., color or orientation singletons); for instance, visual deprivation can mute orientation-dependent responses [75] and would presumably alter the strength of pop-out for orientation singletons. These caveats notwithstanding, we maintain that 'physical salience' is a productive category for a taxonomy of attentional control, given that numerous low-level stimulus qualities induce strong selection biases with virtually no exceptions.




Figure 2. A schematic representation of a priority map that integrates three sources of selection bias: the observer's current selection goals, selection history, and the physical salience of the items competing for attention. Although these three effects could in principle operate in a coordinated fashion, various studies have demonstrated that they can also work in direct opposition to one another. This suggests that they are distinct sources of selection bias.

This modification of the dominant theoretical dichotomy between 'top-down' and 'bottom-up' attentional control may help to clarify ongoing debates regarding the role of voluntary control over selective attention. For example, attention research is often focused on whether an observed selection bias is due to 'top-down' or 'bottom-up' factors (i.e., goal-driven versus stimulus-driven factors, respectively). From our perspective, this may be an ill-formed question because the 'top-down' label does not acknowledge the crucial distinction between the active effects of current goals and the lingering effects of selection history. Indeed, a common strategy in studies of 'top-down' control is to compare the observers' response to a given stimulus display when it appears in different contexts (e.g., a block of trials in which the target has a high probability of being red versus one in which the target is only rarely red [12,76,77]). At first glance, this design nicely controls for the effects of the physical stimulus display by measuring performance

Box 3. Outstanding questions

- Do current selection goals and selection history affect target processing in the same way? A broad range of evidence has shown that attention can influence early sensory encoding, memory consolidation, and later decision stages of processing. Given that many of these studies have confounded goal-driven and history-driven sources of selection bias, future work would need to examine whether different types of 'top-down' selection affect target processing in different ways.
- Do similar capacity limits constrain different sources of selection bias? Studies of attentional tracking and rapid enumeration suggest that there may be strong limits in the number of items or positions that can be simultaneously attended. Are there common constraints on the simultaneous selection of items via goal-driven, stimulus-driven, and history-driven sources?
- Previous reward has been shown to shape selection even multiple days after the initial association between stimulus and reward [34]. What factors determine the rate of extinction for such reward-driven selection effects?

with the same display across contexts that elicit distinct selection goals. However, given that both voluntary goals and recent selection history are strongly influenced by the blocked variable, this design conflates the effects of current selection goals and selection history.

Concluding remarks

The boundary conditions of volitional control represent one of the fundamental questions regarding human cognition. To what extent are we in command of where the mind's eye is directed? Although a major purpose of the top-down versus bottom-up dichotomy has been to highlight the boundary between voluntary and automatic control over attention, we have emphasized how the tendency to equate top-down and goal-driven control over attention results in a taxonomy that cannot account for large swaths of selection phenomena that are unrelated to current selection goals and physical salience. This, in turns, confuses ongoing efforts to determine the prevalence and impact of voluntary selection. Thus, our hope is that acknowledging selection history as a third category of control will aid in the classification and elucidation of a broader range of empirical findings within this core area of research (Box 3).

Acknowledgments

This work was supported by NIH-R01MH077105 to E.A. We thank Andrew Leber and Howard Egeth for helpful discussions and comments.

References

- 1 Posner, M.I. (1980) Orienting of attention. *Q. J. Exp. Psychol.* 32, 3–25
- 2 Jonides, J. (1981) Voluntary versus automatic control over the mind's eye's movement. In *Attention and Performance IX* (Long, J.B. and Baddeley, A.D., eds), pp. 187–203, Lawrence Erlbaum Associates
- 3 Posner, M.I. and Petersen, S.E. (1990) The attention system of the human brain. *Annu. Rev. Neurosci.* 13, 25–42
- 4 Folk, C.L. et al. (1992) Involuntary covert orienting is contingent on attentional control settings. *J. Exp. Psychol. Hum. Percept. Perform.* 18, 1030–1044
- 5 Wolfe, J.M. et al. (1989) Guided search: an alternative to the feature integration model for visual search. *J. Exp. Psychol. Hum. Percept. Perform.* 15, 419–433
- 6 Desimone, R. and Duncan, J. (1995) Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222
- 7 Egeth, H.E. and Yantis, S. (1997) Visual attention: control, representation, and time course. *Annu. Rev. Psychol.* 48, 269–297
- 8 Itti, L. and Koch, C. (2000) A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Res.* 40, 1489–1506
- 9 Kastner, S. and Ungerleider, L.G. (2000) Mechanisms of visual attention in the human cortex. *Annu. Rev. Neurosci.* 23, 315–341
- 10 Corbetta, M. and Shulman, G.L. (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nat. Rev. Neurosci.* 3, 201–215
- 11 Eriksen, C. and Hoffman, J. (1973) The extent of processing of noise elements during selective encoding from visual displays. *Atten. Percept. Psychophys.* 14, 155–160
- 12 Wolfe, J.M. et al. (2003) Changing your mind: on the contributions of top-down and bottom-up guidance in visual search for feature singletons. *J. Exp. Psychol. Hum. Percept. Perform.* 29, 483–502
- 13 Found, A. and Müller, H. (1996) Searching for unknown feature targets on more than one dimension: Investigating a 'dimension-weighting' account. *Atten. Percept. Psychophys.* 58, 88–101
- 14 Bundesen, C. (1990) A theory of visual attention. *Psychol. Rev.* 97, 523–547
- 15 Serences, J.T. and Boynton, G.M. (2007) Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron* 55, 301–312
- 16 Duncan, J. and Humphreys, G.W. (1989) Visual search and stimulus similarity. *Psychol. Rev.* 96, 433–458
- 17 Maljkovic, V. and Nakayama, K. (1994) Priming of pop-out: I. Role of features. *Mem. Cogn.* 22, 657–672
- 18 Eimer, M. et al. (2010) Priming of pop-out modulates attentional target selection in visual search: behavioural and electrophysiological evidence. *Vision Res.* 50, 1353–1361
- 19 Kristjansson, A. and Campana, G. (2010) Where perception meets memory: a review of repetition priming in visual search tasks. *Atten. Percept. Psychophys.* 72, 5–18
- 20 Brascamp, J.W. et al. (2011) Deciding where to attend: priming of pop-out drives target selection. *J. Exp. Psychol. Hum. Percept. Perform.* 37, 1700–1707
- 21 Theeuwes, J. and Van der Burg, E. (2011) On the limits of top-down control of visual selection. *Atten. Percept. Psychophys.* 73, 2092–2103
- 22 Hillstrom, A.P. (2000) Repetition effects in visual search. *Percept. Psychophys.* 62, 800–817
- 23 Kristjansson, A. et al. (2002) The role of priming in conjunctive visual search. *Cognition* 85, 37–52
- 24 Maljkovic, V. and Nakayama, K. (1996) Priming of pop-out: II. The role of position. *Percept. Psychophys.* 58, 977–991
- 25 Wang, D. et al. (2005) Efficient visual search without top-down or bottom-up guidance. *Atten. Percept. Psychophys.* 67, 239–253
- 26 Van der Stigchel, S. et al. (2009) The limits of top-down control of visual attention. *Acta Psychol. (Amst.)* 132, 201–212
- 27 Theeuwes, J. and Van der Burg, E. (2007) The role of spatial and nonspatial information in visual selection. *J. Exp. Psychol. Hum. Percept. Perform.* 33, 1335–1351
- 28 Zhang, W. and Luck, S.J. (2009) Feature-based attention modulates feedforward visual processing. *Nat. Neurosci.* 12, 24–25
- 29 Saenz, M. et al. (2002) Global effects of feature-based attention in human visual cortex. *Nat. Neurosci.* 5, 631–632
- 30 Theeuwes, J. et al. (2006) Visual search for featural singletons: No top-down modulation, only bottom-up priming. *Vis. Cogn.* 14, 466–489
- 31 Leonard, C.J. and Egeth, H.E. (2008) Attentional guidance in singleton search: An examination of top-down, bottom-up, and intertrial factors. *Vis. Cogn.* 16, 1078–1091
- 32 Thorndike, E.L. (1911) *Animal Intelligence*, Macmillan
- 33 Shuler, M.G. and Bear, M.F. (2006) Reward timing in the primary visual cortex. *Science* 311, 1606–1609
- 34 Della Libera, C. and Chelazzi, L. (2009) Learning to attend and to ignore is a matter of gains and losses. *Psychol. Sci.* 20, 778–784
- 35 Della Libera, C. and Chelazzi, L. (2006) Visual selective attention and the effects of monetary rewards. *Psychol. Sci.* 17, 222–227
- 36 Dayan, P. and Balleine, B.W. (2002) Reward, motivation, and reinforcement learning. *Neuron* 36, 285–298
- 37 Seitz, A.R. et al. (2009) Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron* 61, 700–707
- 38 Raymond, J.E. and O'Brien, J.L. (2009) Selective visual attention and motivation: the consequences of value learning in an attentional blink task. *Psychol. Sci.* 20, 981–988
- 39 Ikeda, T. and Hikosaka, O. (2003) Reward-dependent gain and bias of visual responses in primate superior colliculus. *Neuron* 39, 693–700

- 40 Kiss, M. *et al.* (2009) Reward priority of visual target singletons modulates event-related potential signatures of attentional selection. *Psychol. Sci.* 20, 245–251
- 41 Peck, C.J. *et al.* (2009) Reward modulates attention independently of action value in posterior parietal cortex. *J. Neurosci.* 29, 11182–11191
- 42 Small, D.M. *et al.* (2005) Monetary incentives enhance processing in brain regions mediating top-down control of attention. *Cereb. Cortex* 15, 1855–1865
- 43 Kristjánsson, Á. *et al.* (2010) Fortune and reversals of fortune in visual search: reward contingencies for pop-out targets affect search efficiency and target repetition effects. *Atten. Percept. Psychophys.* 72, 1229–1236
- 44 Navalpakkam, V. *et al.* (2010) Optimal reward harvesting in complex perceptual environments. *Proc. Natl. Acad. Sci. U.S.A.* 107, 5232–5237
- 45 Maunsell, J.H. (2004) Neuronal representations of cognitive state: reward or attention? *Trends Cogn. Sci.* 8, 261–265
- 46 Colby, C.L. and Goldberg, M.E. (1999) Space and attention in parietal cortex. *Annu. Rev. Neurosci.* 22, 319–349
- 47 Bisley, J.W. (2011) The neural basis of visual attention. *J. Physiol.* 589, 49–57
- 48 Bisley, J.W. and Goldberg, M.E. (2010) Attention, intention, and priority in the parietal lobe. *Annu. Rev. Neurosci.* 33, 1–21
- 49 Goldberg, M.E. *et al.* (2006) Saccades, salience and attention: the role of the lateral intraparietal area in visual behavior. *Prog. Brain Res.* 155, 157–175
- 50 Gottlieb, J.P. *et al.* (1998) The representation of visual salience in monkey parietal cortex. *Nature* 391, 481–484
- 51 Dorris, M.C. and Glimcher, P.W. (2004) Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 44, 365–378
- 52 Louie, K. *et al.* (2011) Reward value-based gain control: divisive normalization in parietal cortex. *J. Neurosci.* 31, 10627–10639
- 53 Sugrue, L.P. *et al.* (2004) Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782–1787
- 54 Nakayama, K. and Martini, P. (2011) Situating visual search. *Vision Res.* 51, 1526–1537
- 55 Pessoa, L. and Engelmann, J.B. (2010) Embedding reward signals into perception and cognition. *Front. Neurosci.* 4, 17
- 56 Hickey, C. *et al.* (2010) Reward changes salience in human vision via the anterior cingulate. *J. Neurosci.* 30, 11096–11103
- 57 Hickey, C. *et al.* (2010) Reward guides vision when it's your thing: trait reward-seeking in reward-mediated visual priming. *PLoS ONE* 5, e14087
- 58 Anderson, B.A. *et al.* (2011) Learned value magnifies salience-based attentional capture. *PLoS ONE* 6, e27926
- 59 Anderson, B.A. *et al.* (2011) Value-driven attentional capture. *Proc. Natl. Acad. Sci. U.S.A.* 108, 10367–10371
- 60 Theeuwes, J. (2010) Top-down and bottom-up control of visual selection. *Acta Psychol. (Amst.)* 135, 77–99
- 61 Awh, E. *et al.* (2005) Resolving visual interference during covert spatial orienting: online attentional control through static records of prior visual experience. *J. Exp. Psychol. Gen.* 134, 192–206
- 62 Chun, M.M. and Jiang, Y. (1998) Contextual cueing: implicit learning and memory of visual context guides spatial attention. *Cogn. Psychol.* 36, 28–71
- 63 Crump, M.J. *et al.* (2008) Context-specific learning and control: the roles of awareness, task relevance, and relative salience. *Conscious. Cogn.* 17, 22–36
- 64 Umemoto, A. *et al.* (2010) Statistical learning induces discrete shifts in the allocation of working memory resources. *J. Exp. Psychol. Hum. Percept. Perform.* 36, 1419–1429
- 65 Itti, L. and Koch, C. (2001) Computational modelling of visual attention. *Nat. Rev. Neurosci.* 2, 194–203
- 66 Fecteau, J.H. and Munoz, D.P. (2006) Salience, relevance, and firing: a priority map for target selection. *Trends Cogn. Sci.* 10, 382–390
- 67 Logan, G.D. (2002) An instance theory of attention and memory. *Psychol. Rev.* 109, 376–400
- 68 Logan, G.D. (1988) Toward an instance theory of automatization. *Psychol. Rev.* 95, 492–527
- 69 Nothdurft, H.C. (1993) Saliency effects across dimensions in visual search. *Vision Res.* 33, 839–844
- 70 Anderson, A.K. and Phelps, E.A. (2001) Lesions of the human amygdala impair enhanced perception of emotionally salient events. *Nature* 411, 305–309
- 71 Most, S.B. and Wang, L. (2011) Dissociating spatial attention and awareness in emotion-induced blindness. *Psychol. Sci.* 22, 300–305
- 72 Williams, M. *et al.* (2005) Look at me, I'm smiling: Visual search for threatening and nonthreatening facial expressions. *Vis. cogn.* 12, 29–50
- 73 Hodsoll, S. *et al.* (2011) Attentional capture by irrelevant emotional distractor faces. *Emotion* 11, 346–353
- 74 Jiang, Y. *et al.* (2006) A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proc. Natl. Acad. Sci. U.S.A.* 103, 17048–17052
- 75 Chapman, B. and Stryker, M.P. (1993) Development of orientation selectivity in ferret visual cortex and effects of deprivation. *J. Neurosci.* 13, 5251–5262
- 76 Muller, H.J. *et al.* (2009) Attentional capture by salient color singleton distractors is modulated by top-down dimensional set. *J. Exp. Psychol. Hum. Percept. Perform.* 35, 1–16
- 77 Geyer, T. *et al.* (2008) Expectancies modulate attentional capture by salient color singletons. *Vision Res.* 48, 1315–1326
- 78 Anderson, B.A. and Folk, C.L. (2010) Variations in the magnitude of attentional capture: testing a two-process model. *Atten. Percept. Psychophys.* 72, 342–352
- 79 Eimer, M. and Kiss, M. (2008) Involuntary attentional capture is determined by task set: evidence from event-related brain potentials. *J. Cogn. Neurosci.* 20, 1423–1433
- 80 Belopolsky, A.V. *et al.* (2010) What is top-down about contingent capture? *Atten. Percept. Psychophys.* 72, 326–341
- 81 Rossi, A.F. *et al.* (2007) Top down attentional deficits in macaques with lesions of lateral prefrontal cortex. *J. Neurosci.* 27, 11306–11314
- 82 Theeuwes, J. (1991) Exogenous and endogenous control of attention: the effect of visual onsets and offsets. *Percept. Psychophys.* 49, 83–90