# Guidance of Attention by Feature Relationships: The End of the Road for Feature Map Theories?

**Stefanie I. Becker**

## 1 Introduction

It is well known that conscious perception is severely capacity limited: At any moment in time, only a few objects can be consciously perceived. Attention is needed to select items from cluttered visual scenes for further in-depth processing. In the past, much effort has been devoted to identify the factors that guide attention and determine which item will be selected first.

The currently dominant view is that attention can be guided by two independent attentional systems: First, a *stimulus-driven system* guides attention to the most salient locations in the visual field, such as suddenly appearing items ('onsets'; e.g. Yantis 2000), or items with a high-feature contrast (e.g. Theeuwes 1994, 2010; Wolfe 1994). Importantly, attention is allocated to these items in a purely stimulus-driven fashion, that is, without or even against the goals and intentions of the observers to perform a certain task (e.g. Yantis 1993). A second attentional system is *goal-dependent* and guides attention to items that match the observer's goals and intentions to find a sought-after item (e.g. Folk et al. 1992; Wolfe 1994). For instance, when we search for a red item, such as a woman with a red skirt, attention can be involuntarily captured by other red items. Importantly, capture by target-similar items is usually much stronger than capture by salient irrelevant items that do not match our top-down settings (e.g. Folk and Remington 1998). This *similarity effect* (defined as stronger capture by target-similar than target-dissimilar distractors) has been found in numerous studies with the visual search paradigm and has usually been taken to show that top-down tuning to specific feature values can override effects of bottom-up saliency (e.g. Ansorge and Heumann 2003; Becker et al. 2009; Eimer et al. 2009; Folk and Remington 1998; Ludwig and Gilchrist 2002).

Most models of attentional guidance assume that attention is tuned to particular *feature values*. For example, in the colour dimension, attention can be biased to

S. I. Becker (✉)
The University of Queensland, Brisbane, Australia
e-mail: s.becker@psy.uq.edu.au

select either red, green, yellow or blue items; in the orientation dimension, attention can be biased for horizontal, vertical or differently tilted orientations and in the size dimension, attention can be biased to select items of different sizes (e.g. Treisman and Gelade 1980; Treisman and Sato 1990; Wolfe 1998; Wolfe and Horowitz 2004).

Multiple different mechanisms have been proposed to describe feature-based tuning of attention. For instance, *feature similarity views* assume that attention is tuned to the target feature value (e.g. Martinez-Trujillo and Treue 2004). According to the *attentional engagement theory,* attention can be additionally biased against selecting the feature value of the nontarget(s) ('nontarget inhibition'; see; e.g. Duncan and Humphreys 1989). Although accounts of top-down selection differ with respect to the feature value that will be prioritized in a given instance, they uniformly assume that top-down selection is achieved by activating or inhibiting specific feature maps (e.g. red, green; Duncan and Humphreys 1989; Folk and Remington 1998; Koch and Ullman 1985; Maunsell and Treue 2006; Navalpakkam and Itti 2007; Treisman and Gelade 1980; Treisman and Sato 1990; Wolfe 1994). Feature maps are populations of sensory neurons that are topographically organized and respond to specific feature values ('feature detectors'). Most theories assume that visual selection is achieved by modulating the response gain of feature-specific neurons, increasing the response gain of feature detectors responding to the target feature, and/or decreasing the response gain of feature detectors responding to the nontarget features (e.g. Koch and Ullman 1985; Maunsell and Treue 2006; Navalpakkam and Itti 2007; Spitzer et al. 1988).

Contrary to these feature-based theories, it has recently been proposed that attention is guided by target–nontarget relationships, that is, information that specifies how the target differs from irrelevant items (Becker 2010). According to this new *relational view,* the target and irrelevant context are not evaluated separately—with activation applied to the target feature and inhibition to the nontarget feature. Instead, the feature relationship between the target and context is evaluated and attention is guided towards items sharing the target–context relation (e.g. redder). For example, when searching for the orange shirt of a goalkeeper, it would depend on the context how attention is top-down tuned: When the goalkeeper is embedded in a team wearing all yellow shirts, as in the left panel of Fig. 1, attention would be tuned towards all *redder* items. By contrast, if the goalkeeper is surrounded by a team wearing all red shirts, attention would be tuned to all *yellower* items (see right panel of Fig. 1).

## 1.1   Is There a Similarity Effect?

The relational theory is overall quite similar to the classical feature-based views, but the two theories make different predictions with regard to stimuli that can capture our attention. For example, if observers have to search for an orange target, and attention is top-down tuned to the target feature value, then only orange items should capture attention. Items of different colours should capture to the degree that they are similar to orange (feature similarity view; e.g. Folk and Remington 1998). By contrast, the relational account predicts that attention should be tuned to target–

**Fig. 1** Example of searching for orange in different contexts: According to the relational account, colours are encoded relative to the context, and visual selection is biased to the relative attribute that distinguishes the target from the context. Hence, when searching for the orange-shirted goalkeeper, visual selection would be biased to *redder* when the goalkeeper is among a team clad in yellow (left), whereas attention would be biased to *yellower*, or the *yellowest* item when the goalkeeper is among red-shirted players (right).

nontarget relationships. Thus, if the orange target is redder than the nontargets (e.g. yellow), attention should be tuned to all redder items. A consequence of tuning to redder is that items that are redder than the target itself should be able to attract attention. This holds because a relational top-down setting specifies only the direction in which an item differs from the context, and does not contain information about the exact feature values of the target or the context. Hence, the item with the most extreme feature in the specified direction (e.g. reddest item) should always capture attention most strongly (e.g. Becker 2008, 2010).

Studies testing the relational account against the feature similarity account confirmed the predictions of the relational account (e.g. Becker 2010; Becker et al. 2010): When observers had to search for an orange target among irrelevant yellow nontargets, a red distractor captured attention more strongly than an orange distractor—despite the fact that the red distractor was more dissimilar from the target. When the orange target was embedded among red irrelevant nontargets, a yellow distractor captured more strongly than a target-similar orange distractor, consistent with a top-down setting for yellower (e.g. Becker et al. 2010). Across all conditions, the colour contrasts of all stimuli were controlled, so that stronger capture by the red or yellow distractors could not be explained by bottom-up factors such as feature contrast or visual saliency (Becker et al. 2010; Becker and Horstmann 2011).

These results were diametrically opposite to the often-reported similarity effect, that items can only attract attention when they are similar to the sought-after target feature (e.g. Folk and Remington 1998), and demonstrated, for the first time, that a target-dissimilar distractor can capture attention more strongly than a target-similar distractor. This is an important finding, as the similarity effect can be regarded as the most direct evidence for feature-based theories of top-down guidance. Although these first results ruled out similarity theories of attentional guidance, they are still consistent with a less well-known group of feature-based theories, viz., optimal tuning accounts.

## 2  The Relational Account Versus Optimal Tuning Accounts

More capture by target-dissimilar distractors may still be consistent with feature-based accounts, when we assume that attention can be top-tuned tuned towards a different feature value than the target feature value (e.g. tuning to red in search for orange). *Optimal tuning accounts* assume that attention is always tuned towards the feature value that optimally distinguishes the target from the nontarget features, and would allow for such a shift in top-down tuning. Especially when the feature value distributions for the target and nontarget features overlap to a large extent, tuning attention to a more extreme feature that is shifted away from the nontarget features can enhance the signal-to-noise ratio and provide a more optimal setting than tuning to the target feature itself (e.g. Lee et al. 1999; Navalpakkam and Itti 2007; Scolari and Serences 2010). With this, optimal tuning accounts could provide an alternative explanation for the finding that a red distractor captured more than a target-similar orange distractor, simply by assuming that attention was tuned more towards red than orange.

The optimal tuning account and the relational account both predict that a distractor with an 'exaggerated' target colour (i.e. a colour shifted away from the nontarget colour) should capture more strongly than a target-similar distractor. However, optimal tuning accounts assume that a distractor with the *nontarget colour* could *not* capture attention, because attention has to be biased away from the nontarget feature (to provide a better signal-to-noise ratio). By contrast, the relational account would *allow* capture by a nontarget colour (provided that the set-up allows this colour to differ from all other colours and to be the most extreme in the target-defining direction). This holds because visual selection of the target is thought to depend only on its relationship to other items, not a specific feature value, so any item with the same relationship(s) as the target should capture.

In the visual search paradigm, the target and distractors are always presented in the same display and hence it is impossible to create targets and distractors that have different feature values and yet both have the most extreme feature value (e.g. reddest item in the display). To critically test whether a nontarget-coloured distractor can capture attention, it is necessary to present the distractors ('cues') prior to the target in a separate cueing display, so that the distractor (cue) features can be manipulated independently of the target and nontarget features. In a recent study (Becker et al. 2013), observers were asked to search for an orange target among three yellow-orange nontargets (target redder condition), and to ignore differently coloured cues that were briefly flashed (100 ms) prior to the target display. The cueing displays consisted of four cues, three of which constituted the context for the differently coloured singleton cue. In one condition, the singleton cue had the same colour as the nontargets (yellow-orange), and the other three cues were yellow, rendering the singleton cue redder than the cue context (same relative colour as the target). The results showed that the yellow-orange cue with the nontarget colour captured attention (Becker et al. 2013), as reflected in faster response times when it was presented at the target location (valid cue) than when it was presented at a

nontarget location (invalid cue; e.g. Posner 1980). These results demonstrate that, contrary to the optimal tuning account, a distractor with the nontarget colour can still capture attention, provided that the cue–context relations match the relationship between the target and the nontargets.

The study of Becker et al. (2013) included multiple control conditions to ensure that capture by a nontarget-coloured cue is not due to bottom-up factors (e.g. the specific colours used in one condition). In one experimental block, the target and nontarget colours were also reversed, so that observers now had to search for a yellow-orange target among orange nontargets (target yellower condition). In this condition, the yellow-orange cue (among yellow-other cues) failed to capture attention—despite the fact that it had the same colour as the target. This outcome was predicted by the relational theory: As the task required tuning to a yellower target, a target-similar cue should not capture if it is itself redder than the cue context. These findings strongly support the relational account that capture is largely independent of the absolute feature values of target and nontargets, and instead depends on the relationships of target and distractor.

Taken together, the current evidence invalidates the prevalent doctrine that capture necessarily depends on similarity to the target feature (i.e. 'similarity effect'; e.g. Folk and Remington 1998; Becker et al. 2008), or similarity to the exaggerated target feature (i.e. 'optimal feature'; e.g., Lee et al., 1999, Scolari & Serences, 2010) and suggests that capture is instead determined by feature relationships. Previous studies did not vary the similarity of distractors independently of their relative features: Hence, it is possible that the often-reported similarity effect in previous studies was due to the distractor matching the relative, not absolute, colour of the target. In fact, one of the strengths of the relational theory is that it seems consistent with all results that were previously interpreted in support of a feature-specific top-down setting (cf. Becker 2010).

## 3   Can Guidance by Relationships Be Explained by Feature-Based Theories of Attention?

Is guidance of attention by feature relationships really inconsistent with common feature-based accounts of attention, or can feature based-accounts somehow account for top-down tuning to relative features? This question is unfortunately not easy to answer, as the top-down tuning component as well as the feature maps, 'channels' or 'filters' are not clearly specified in the mainstream models of visual attention (e.g. Mozer and Baldwin 2008). The view most commonly found in the literature seems to be that attention is guided by separate, disconnected feature maps that basically act like feature detectors and signal the location of items with particular feature values (e.g. Itti and Koch 2000; Koch and Ullman 1985; Treisman and Gelade 1980; Treisman and Sato 1990; Wolfe 1994). With this, feature map theories do not seem to be able to account for guidance by feature relationships, as it is impossible

to specify feature relationships within feature maps. For example, an orange feature detector cannot signal whether the item in its receptive field is redder or yellower than other items in the context. To obtain information about feature relationships, we would need additional feature detectors that signal, for example, the presence of red and yellow items. Then, feature relations could be computed by another layer of neurons (i.e. 'comparator maps') that receive input from the orange, yellow and red feature detectors. However, to account for guidance by feature relationships, attention would have to be guided by comparator maps—not feature maps or feature detectors, as is proposed by current feature-based models of visual attention.

Adding another layer of neurons to feature map models may also not seem the most parsimonious way to account for guidance by feature relationships. In fact, one rather problematic aspect of feature-based theories is that they have to propose a feature map or population of feature-specific neurons for each feature that can be top-down selected. As noted by Maunsell and Treue (2006), the number of neurons required by feature map theories may very well exceed the number of neurons in the brain that are actually involved in the guidance of attention. Adding comparator maps to the already proposed feature maps would aggravate this problem and hence does not seem advisable.

An alternative solution would be to propose a small number of categorical channels that respond to an entire range of feature values (e.g., all colour values from yellow to red), with a red detector responding maximally to red (e.g., Wolfe, 1994). Such broadly tuned feature detectors could indeed explain many relational effects (e.g., Becker, 2010). However, such an account could not explain our ability to select intermediate colours (e.g., orange among red and yellow). If the account is equipped with an extra channel to explain fine-grained selectivity (e.g., an extra orange channel), it loses the ability to explain relational effects such as capture by nontarget-coloured items.

In sum, the only practical and parsimonious approach seems to be to abandon the idea of separate feature detectors, and to ponder the idea that sensory neurons can directly encode the direction in which features differ from one another, and that this information can be used to direct attention towards items that have the same feature relationship(s) (e.g. Becker 2010; Becker et al. 2013). In the following, it will be argued that the current neurophysiological evidence also supports the possibility of 'relational feature detectors'.

## 4   Are Sensory Neurons Feature Detectors or Relational?

The idea that attention is guided by independently working feature maps or feature detectors has dominated the theoretical landscape for a long time, and the reason that the concept has been so very successful is probably because feature detectors seem to have the most basic response characteristics, and we can easily imagine how they work (e.g. Nakayama and Martini 2011). However, the idea that atten-

tion is guided by feature detectors is not strongly supported by neurobiological evidence. For feature detectors or 'channels' to be operating *independently* of one another, it appears that different colours, for example, must be encoded by entirely different populations of neurons. Contrary to this claim, neurophysiological studies have found many neurons in the visual cortex of the monkey that respond to two colours, in the fashion of an opponent-colour mechanism (e.g. Hubel and Wiesel 1967; De Valois et al. 2000; Gouras 1974). For example, some opponent-colour cells increase their firing rate in response to input from L-cones (e.g. red) but significantly decrease their firing rate in response to input from M-cones (e.g. green). These neurons cannot be regarded as feature detectors, because they do not respond to a specific colour. Opponent cells also often show different responses depending on whether they receive input from the same cone in the centre or in the surround of their receptive fields (RFs). L–M opponent-colour cells, for example, fire in response to an L-cone increment in the centre but also to an M-cone decrement in the surround of the RF (e.g. Conway 2001; De Valois et al. 2000). These cells and their counterparts (e.g. M–L cells) are predominantly found in the lateral geniculate nucleus (LGN) and could be regarded as 'purely relational cells', because they signal that the centre of the receptive field is occupied by something 'redder' than the surround—without specifying whether the centre contains red or whether the surround contains green. Also at later cortical stages, there are still 'relational neurons' that relate inputs from different cone types to one another, in line with the hypothesis that neurons can signal feature relationships (e.g. Conway 2001; De Valois et al. 2000).

Admittedly, it is at present unclear what proportions of colour-sensitive neurons should be classified as relational cells vs. feature detectors, and whether and to what extent these different classes of neurons are involved in the guidance of attention. Of note, studies investigating the neurophysiological underpinnings of attention also cannot distinguish between a relational and feature-specific account of guidance: Several studies reported that the response gain of some feature-specific neurons increased in expectation of the target (e.g. Martinez-Trujillo and Treue 2004; Motter 1994). However, none of the studies systematically varied the target–background relations and/or the distractor-background relations (e.g. Atiani et al. 2009; David et al. 2008; Kastner et al. 1999; Luck et al. 1997; Motter 1994; Scolari and Serences 2010; Spitzer et al. 1988). Hence, the available neurophysiological evidence which has always been interpreted in support of a feature-based account also seems consistent with a relational account.

In sum, the current state of evidence does not seem to provide good reasons to claim that relational neurons could not exist or that they could not guide attention. Naturally, the existence of neurons signalling feature relationships also does not preclude that feature detectors exist or that they can guide attention as well (e.g. Conway and Tsao 2009). However, at least in the monkey, early processing of colour (e.g. in the LGN) seems to initially rely mostly on L–M type of cells that are context dependent (e.g. De Valois et al. 2000). All information that can be used at

later, cortical stages has therefore to be extracted from relational, context-dependent information.

It seems possible that colour processing initially proceeds relational and becomes feature specific only at later stages of visual processing. This is also in line with the observation that colour perception shows both relational and absolute characteristics. Relational or context-dependent characteristics are reflected in the fact that the visual system remains susceptible to different surrounding colours, so that, for example, a grey patch can look slightly green when surrounded by a large green area ('simultaneous colour contrast'; e.g. Conway et al. 2002). Absolute characteristics are, for instance, reflected in our ability to categorize colours, and to recognize a specific colour in many different lighting conditions ('colour constancy'). Given that all colour processing has to be based on the initially relational information, it is an interesting question to what extent colour perception may still be relational at very late stages of visual processing. In fact, Foster (2003) argued that there is insufficient evidence for colour constancy, as colour constancy could itself be relational.

## 5   Can the Relational Theory Explain Feature-Specific Tuning?

Above, it was argued that the current neurophysiological evidence does not appear to support a feature detector account more strongly than a relational account of attention. In addition, it could be asked whether there is any more direct evidence that a relational theory can account for tuning to specific feature values. A first problem with this question is that it is not entirely clear what would qualify as an instance for 'clearly feature-specific tuning'.

As outlined above, the perhaps most compelling evidence for feature-specific tuning was the similarity effect, that is, the finding that target-similar distractors capture attention most strongly. This finding has been called into question by findings demonstrating that target-dissimilar distractors can capture as well, provided that relative attributes as the target (e.g. Becker et al. 2013, in press). Another finding that can be interpreted as evidence for feature-specific tuning is our ability to select items that have intermediate feature values, that is, features that are directly sandwiched between more extreme features. For example, visual search studies have shown that attention can be tuned to a medium-sized target when half of the nontargets are larger and half are smaller than the target (e.g. Hodsoll and Humphreys 2005). Similarly, we can also select an orange item when half of the nontargets are red and the other half is yellow (e.g. Bauer et al. 1995; Navalpakkam and Itti 2006). In these examples, the target feature does not differ in a single direction from the nontarget features; rather, as half of the nontargets is smaller (or redder) than the target and the other half is larger (or non-redder), the target differs in two opposing directions from the nontargets ('non-linearly separable';

e.g. Bauer et al. 1995). Previous studies show that search efficiency is decreased for such intermediate features; however, search is not random, indicating that attention can still be top-down guided to intermediate features (Bauer et al. 1995; Hodsoll and Humphreys 2005; Navalpakkam and Itti 2006).

To explain these results from a relational perspective, it has to be assumed that attention can be biased to two opposing relationships simultaneously, to select the item that is 'neither the reddest nor the yellowest item' in the display. Such a two-fold relational setting would result in an attentional bias for items intermediate between two more extreme feature values (i.e. reddest and yellowest) and, thus, would mimic a feature-specific setting. However, deviating from a feature-specific account, selection of the intermediate item should still strongly depend on the context; that is, in search for an orange target among red and yellow nontargets, an orange distractor should only capture attention when it is similarly embedded in a context containing both redder and yellower items. Orange distractors should, however, fail to capture when they are either the reddest or the yellowest items in their context(s).

This hypothesis was tested in a spatial cueing study, where observers had to search for an orange target among two red and two yellow nontargets (Becker et al. 2013). Prior to the target frame, a cueing frame was presented that contained five cues, four of which constituted the cue context. Capture was tested by a cue that could have either the same colour as the target (orange) or a different colour (yellow-orange). Critically, the context cues were coloured such that the uniquely coloured singleton cue sometimes had an intermediate colour relative to the cue context and sometimes an extreme colour (e.g. orange the reddest colour in the cueing display). The results showed that an orange cue captured attention only when it was embedded among redder and yellower cues, but not when the cue context rendered it the reddest cue. In addition, a target-dissimilar yellow-orange cue captured attention to the same extent as the target-similar (orange) cue when it was embedded among redder and yellower cues, but failed to capture when it was the yellowest cue in the cueing display (Becker et al. 2013).

These results indicate that attention was not top-down tuned to the exact feature value of the target (orange). Rather, capture depended on whether the context rendered the singleton cue colour intermediate between more extreme colours, indicating that attention had been tuned simultaneously towards the target feature and the features of the context. This finding has two important implications: First, it indicates that the target and nontarget colours (red, yellow and orange) were not processed independently and separately of each other—contrary to the major claims of current feature-based theories, that visual selection is achieved by a number of independent feature detectors (e.g. Treisman and Sato 1990; Wolfe 1994). Second, and even more importantly, the results demonstrate that information from different target–nontarget relations (redder/yellower) can be used to bias attention to intermediate features. With this, a relational top-down setting can mimic the effects of a feature-specific top-down setting, as it allows selection of a range of intermediate feature values that are bounded by more extreme features. In fact, it is possible that

evidence previously interpreted in favour of feature-specific tuning was really due to relational tuning. Of note, the target and nontarget features were always held constant in previous studies (e.g. Bauer et al. 1995; Hodsoll and Humphreys 2005; Navalpakkam and Itti 2006). Hence, it is possible that observers did not apply a feature-specific setting, but used information provided by the nontarget context to tune attention to the target in a relational manner.

That said, it should also be noted that the most recent studies showed that we are also able to tune attention to a specific target size or target colour, when the nontargets varied randomly and tuning attention to relative features could not help localising the target (Becker et al., in press; Harris et al., 2013). Interestingly, performance was poorer in this feature-specific search mode than when observers engaged in relational search (Becker et al., in press). Taken together, there is evidence that attention can be tuned to the target attributes in a context-dependent manner, as well as to specific features in a context-independent manner. The challenge for future research will be to find out how the two search modes can be explained in a unified theory of attention, and what their respective neural substrates are.

## 6   Conclusion

The present chapter introduced a new relational account of attention and argued that it can provide an alternative to current feature-based theories of attention. Among the findings in favour of the relational account were results showing that feature relationships account for (1) capture by irrelevant distractors and (2) selection of a target with an intermediate feature, among heterogeneous nontargets. Some of the findings in favour of the relational account were clearly inconsistent with feature-specific views of guidance, amongst them the finding that a nontarget-coloured distractor can capture attention. Moreover, it has been argued that the current neurophysiological evidence also does not unequivocally support the feature detector concept. From a theoretical stance, top-down tuning to feature relationships seems to be able to account for the majority of findings that were previously interpreted as evidence for feature-specific tuning of attention, including neurophysiological evidence (but see Becker et al., in press; Harris et al., 2013). Whether a relational account can eventually explain all the instances of feature-specific guidance of attention, and may eventually replace feature-based accounts of attention remains to be determined by future research. What is clear from the present review is that the features in the context play an important role in top-down tuning of attention. Thus, studies focussing on attention shifts to single stimuli in isolation are unlikely to provide the insights necessary to advance current theories of visual search.

# References

Ansorge, U., & Heumann, M. (2003). Top-down contingencies in peripheral cueing: The roles of color and location. *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 937–948.

Atiani, S., Elhilali, M., David, S. V., Fritz, J. B., & Shamma, S. A. (2009). Task difficulty and performance induce diverse adaptive patterns in gain and shape of primary auditory cortical receptive fields. *Neuron*, *61*, 467–480.

Bauer, B., Jolicoeur, P., & Cowan, W. B. (1995). Visual search for colour targets that are or are not linearly separable from distractors. *Vision Research*, *36*, 1439–1465.

Becker, S. I. (2010). The role of target-distractor relationships in guiding attention and the eyes in visual search. *Journal of Experimental Psychology: General*, *139*, 247–265.

Becker, S. I. (2008). Can intertial effects of features and dimensions be explained by a single theory? *Journal of Experimental Psychology: Human Perception and Performance*, *34*, 1417–1440.

Becker, S. I., & Horstmann, G. (2011). Novelty and saliency in attentional capture by unannounced motion singletons. *Acta Psychologica*, *136*, 290–299.

Becker, S. I., Ansorge, U., & Horstmann, G. (2009). Can intertial priming account for the similarity effect in visual search? *Vision Research*, *49*, 1738–1756.

Becker, S. I., Folk, C. L., & Remington, R. W. (2010). The role of relational information in contingent capture. *Journal of Experimental Psychology: Human Perception and Performance*, *36*, 1460–1476.

Becker, S. I., Folk, C. L., & Remington, R. W. (2013). Attentional capture does not depend 2 feature similarity, but on target-nontarget relations. *Psychological Science*. 24, 634–647.

Becker, S.I., Harris, A.M., Venini, D., & Retell, J.D. (in press). Visual search for colour and shape: When is the gaze guided by feature relationships, when by feature values? Journal of Experimental Psychology: Human Perception and Performance.

Conway, B. R. (2001). Spatial structure of cone inputs to color cells in alert macaque primary visual cortex (V1). *The Journal of Neuroscience*, *21*, 2768–2783.

Conway, B. R., & Tsao, D. Y. (2009). Color-tuned neurons are spatially clustered according to color preference within alert macaque posterior inferior temporal cortex. *PNAS*, *106*, 18034–18039.

Conway, B. R., Hubel, D. H., & Livingstone, M. S. (2002). Color contrast in macaque V1. *Cerebral Cortex*, *12*, 915–925.

David, S. V., Hayden, B. Y., Mazer, J. A., & Gallant, J. L. (2008). Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. *Neuron*, *59*, 509–521.

De Valois, R. L., Cottaris, N. P., Elfar, S., Mahon, L. E., & Wilson, J. A. (2000). Some transformations of color information from lateral geniculate nucleus to striate cortex, *PNAS*, *97*, 4997–5002.

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

D'Zmura, M. (1991). Color in visual search. *Vision Research*, *31*, 951–966.

Eimer, M., Kiss, M., Press, C., & Sauter, D. (2009). The roles of feature-specific task set and bottom-up salience in attentional capture: An ERP study. *Journal of Experimental Psychology: Human Perception and Performance*, *35*, 1316–1328.

Folk, C. L., & Remington, R. (1998). Selectivity in distraction by irrelevant featural singletons: Evidence for two forms of attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 847–858.

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 1030–1044.

Foster, D. H. (2003). Does color constancy exist? *TICS*, *7*, 439–443.

Gouras, P. (1974). Opponent-color cells in different layers of foveal striate cortex. *Journal of Physiology*, *238*, 583–602.

Harris, A.M., Remington, R.W., & Becker, S.I. (2013). Feature specificity in attentional capture by size and colour. Journal of Vision, 13:12, 1–15.

Hodsoll, J. P., & Humphreys, G. W. (2005). The effect of target foreknowledge on visual search for categorically separable orientation targets. *Vision Research*, *45*, 2346–2351.

Hubel, D. H., & Wiesel, T. N. (1969). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, *195*, 215–243.

Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, *40*, 1489–1506.

Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., Ungerleider, L. G. (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron*, *22*, 751–761.

Koch, C., & Ullman, S. (1985). Shifts in visual selective attention: Towards the underlying neural circuitry. *Human Neurobiology*, *4*, 219–227.

Lee, D. K., Itti, L., Koch, C., & Braun, J. (1999). Attention activates winner-take-all competition among visual filters. *Nature Neuroscience*, *2*, 375–381.

Luck, S. J., Chelazzi, L., Hillyard, S. A., Desimone, R. (1997) Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, *77*, 24–42.

Ludwig, C. J. H., & Gilchrist, I. D. (2002). Stimulus-driven and goal-driven control over visual selection. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 902–912.

Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology*, *14*, 744–751.

Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *TICS*, *29*(6):317–322.

Motter, B. C. (1994). Neural correlates of attentive selection for color and luminance in extrastriate area V4. *The Journal of Neuroscience*, *14*, 2178–2189.

Mozer, M. C., & Baldwin, D. S. (2008). Experience-guided search: A theory of attentional control. In J. Platt, D. Koller, & Y. Singer (Eds.), *Advances in Neural Information Processing Systems 20* (pp. 1033–1040). Cambridge: MIT Press.

Navalpakkam, V., & Itti, L. (2006). An integrated model of top-down and bottom-up attention for optimizing detection speed. Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2049–2056.

Navalpakkam, V., & Itti, L. (2007). Search goals tunes visual features optimally. *Neuron*, *53*, 605–617.

Nakayama, K., & Martini, P. (2011). Situating visual search. *Vision Research*, *51*, 1526–1537.

Posner, M. I. (1980) Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*, 3–25.

Scolari, M., & Serences, J. T. (2010). Basing perceptual decisions on the most informative sensory neurons. *Journal of Neurophysiology*, *104*, 2266–2273.

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

Theeuwes, J. (1994). Stimulus-driven capture and attentional set: Selective search for color and visual abrupt onsets. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 799–806.

Theeuwes, J. (2010). Top-down and bottom-up control of visual selection. *Acta Psychologica*, *135*, 97–135.

Treisman, A., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, *12*, 97–136.

Treisman, A., & Sato, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 459–478.

Wolfe, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, *1*, 202–238.

Wolfe, J. M. (1998). Visual search. In: H. Pashler (Ed.), *Attention* (pp. 30–73). London: University College London Press.

Wolfe, J. M., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neurosience*, *5*, 495–501.

Yantis, S. (1993). Stimulus-driven attentional capture and attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 676–681.

Yantis, S. (2000). Goal-directed and stimulus-driven determinants of attentional control. In S. Monsell & J. Driver (Eds.), *Attention & performance*, *XVIII* (pp. 73–103). Cambridge: MIT Press.