

This manuscript was published as Liesefeld, H.R., Liesefeld, A.M., Pollmann, S., & Müller, H.J. (2019). Biasing allocations of attention via selective weighting of saliency signals: behavioral and neuroimaging evidence for the dimension-weighting account. In T. Hodgson (Ed.), *Current Topics in Behavioral Neurosciences: Processes of Visuo-spatial Attention and Working Memory*. Berlin/Heidelberg, Germany: Springer. doi:10.1007/7854\_2018\_75

© 2018 Springer. This article may not exactly replicate the final version published. It is not the copy of record. The final publication is available at Springer via [http://dx.doi.org/10.1007/7854\\_2018\\_75](http://dx.doi.org/10.1007/7854_2018_75)

## **Biasing allocations of attention via selective weighting of saliency signals: behavioral and neuroimaging evidence for the Dimension-Weighting Account**

Heinrich René Liesefeld<sup>1,2,\*</sup>, Anna M. Liesefeld<sup>1</sup>, Stefan Pollmann<sup>3</sup>, and Hermann J. Müller<sup>1,4</sup>

<sup>1</sup>Department Psychologie, Ludwig-Maximilians-Universität, München, Germany; <sup>2</sup>Graduate School of Systemic Neurosciences, Ludwig-Maximilians-Universität, München, Germany; <sup>3</sup>Institute of Psychology and Center for Behavioral Brain Sciences, Otto von Guericke University, Magdeburg, Germany;

<sup>4</sup>Department of Psychological Sciences, Birkbeck College, University of London, UK

\* e-mail: Heinrich.Liesefeld@psy.lmu.de

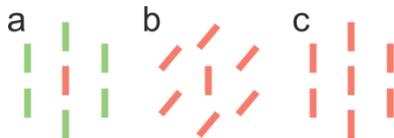
**Author notes** This work was supported by German Research Foundation (DFG) grants MU 773/14-1 and MU 773/16-1 (to H.J.M.), Open Research Area grant DFG PO 548/16-1, and DFG grant CRC779/A4 (to S.P.), by LMU Munich's Institutional Strategy LMUexcellent within the framework of the German Excellence Initiative (to H.J.M. and H.R.L.), and by the Graduate School of Systemic Neurosciences, Munich Center for Neurosciences – Brain & Mind (to H.R.L.).

**Abstract** Objects that stand out from the environment tend to be of behavioral relevance and the visual system is tuned to preferably process these salient objects by allocating focused attention. However, attention is not just passively (bottom-up) driven by stimulus features, but previous experiences and task goals exert strong biases towards attending or actively ignoring salient objects. The core and eponymous assumption of the Dimension-Weighting Account (DWA) is that these top-down biases are not as flexible as one would like them to be; rather, they are subject to dimensional constraints. In particular, DWA assumes that people can often not search for objects that have a particular feature, but only for objects that stand out from the environment (i.e., that are salient) in a particular *feature dimension*. We review behavioral and neuroimaging evidence for such dimensional constraints in three areas: search history, voluntary target enhancement, and distractor handling. The first two have been the focus of research on DWA since its inception and the latter the subject of our more recent research. Additionally, we discuss various challenges to the DWA and its relation to other prominent theories on top-down influences in visual search.

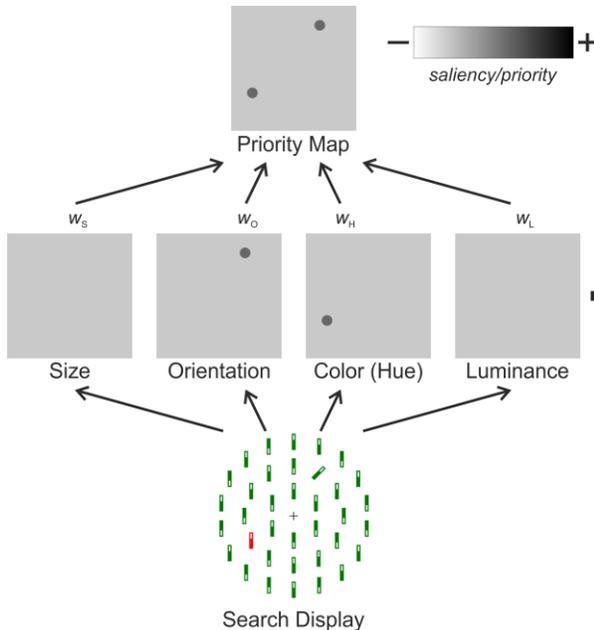
## 1 Saliency Computations for Guiding Visual Search

The visual environment provides a wealth of information of which only a tiny fraction is processed. A crucial factor determining the information uptake is *focused attention*. Major theories of visual attention assume that only one object is highlighted by focused attention at a time, thereby obtaining a competitive advantage in influencing perception and action. Objects standing out from their surround, such as a black sheep in a flock of white sheep, often receive this competitive advantage. This drawing of attention by ‘special’ stimuli is referred to as stimulus-driven or bottom-up guidance of attention and is explained via the workings of saliency maps. Saliency maps are spatial representations of the visual scene coding for the conspicuity of each object in terms of a single saliency value at each (occupied) location. A common assumption is that for each feature dimension, the distance in feature-space between each object and its (immediate) surround (its *saliency*) is calculated before spatial attention is allocated and that these dimension-specific saliency maps are then integrated into a (pre-attentive) priority map that influences the allocation of attentional resources (e.g., Bundesen et al. 2011; also termed ‘activation map’, Wolfe 2007; ‘saliency map’, Itti and Koch 2001; Li 1999, or ‘master map’, Treisman 1988). This means that instead of the object’s features proper, the relation of an object’s features to the respective surrounding features influence attention allocations. The more an object differs from its immediate surround, the more salient and the easier to find it is (Duncan and Humphreys 1989; Liesefeld et al. 2016; Nothdurft 1993; Töllner et al. 2011).

To illustrate the concept of saliency more firmly, take a red, vertical bar, for example. This object is salient due to its color only when presented among objects of a different color (e.g., among green objects; Fig. 1A), but not among homogeneous objects with the same or a very similar shade of red (Fig. 1C). In the latter case, the object can still be salient due to other features, such as its orientation, if presented among red objects of another orientation (e.g., tilted to the right, Fig. 1B). In this review, we will refer to objects that stand out in a particular dimension (mostly targets, but also salient distractors) as *singletons* in that dimension; for example, Fig. 1A features a *color singleton* and Fig. 1B an *orientation singleton*. Thus, it is not any feature of the object per se that is salient, but saliency is defined by the relationship to other features from the same dimension. Saliency signals from the various dimensions are integrated on the priority map. Due to this integration, the priority map does not represent any information on *how* a salient object differs featurally from its surround, but only on *how much* it differs: the priority map is feature- and dimension-less (see Fig. 2). One consequence of this is that a strong activation at the priority map can be induced by any singleton, independently of its specific featural relation to its environment, so that knowledge of target and distractor features is not required to guide attention to a singleton target (Müller et al. 2017).



**Fig. 1** Saliency depends on local feature contrast. The same object (red, vertical bar) is either salient **a** because of its color or **b** because of its orientation or **c** not salient at all. In other words, the very same red, vertical bar is a color singleton in **a**, an orientation singleton in **b**, and not a singleton in **c**



**Fig. 2** Simplified sketch of saliency computations in the visual system. From the search display, saliency values are extracted for each feature dimension, reflecting how strongly each object differs from its surround. As there are only color (hue) and orientation singletons in the display, only these two maps contain any (differential) activation. These activations are weighted and integrated at the superordinate priority map, which in turn guides the allocation of focal attention. Note that the saliency maps at the intermediate layer do not preserve any information on feature values and the priority map is even blind to feature dimensions. In other words, for the intermediate saliency maps, it does not matter how an object differs from its surround, but only whether and by how much it differs in the coded dimension; the priority map is even more abstract in that it cares only about how much an object differs from its surround in

any dimension. The influences of the various saliency maps on the priority map depend on the weight settings ( $w_s$ ,  $w_o$ ,  $w_M$ , and  $w_L$ ; see text); this example assumes that all weights are set to the same positive value and that the bottom-up saliency of the red and the tilted object are identical, so that their priority map activations are identical as well.

However, several decades of research have firmly established that spatial attention is not only reflexively guided by bottom-up saliency, but also by how the observer is ‘tuned’ prior to the start of the search (i.e., before the search display is presented in a laboratory visual-search task). This tuning proceeds either voluntarily, because of what the observer intends to find (*search goals*), or involuntarily, because of what the observer has done before (*search history*; sometimes also termed ‘selection history’; see Awh et al. 2012, or Wolfe and Horowitz 2017, for reviews on these and other potential influences on search guidance). One specific idea of how this tuning might work is that incoming saliency signals are weighted before they are integrated on the priority map (see Fig. 1 and Box 1). In this conceptual framework, ‘tuning’ means changing the weights to the advantage of one type of saliency signal (e.g., target enhancement) or at the disadvantage of another type of saliency signal (e.g., distractor suppression). Biased in this way, the priority map then ‘guides’ attention, ensuring that objects with a high priority are processed first (Wolfe 2007), which higher likelihood (Liesefeld et al. 2016; Moran et al. 2013,2017), and/or more efficiently (Bundesen et al. 2011; Moran et al. 2016).

Given that such tuning heavily influences perception via controlling attention allocations (e.g., Treisman 2006), one of the most pressing questions is: what are the properties the visual system is tuned to in various situations? One intuitively obvious answer is that the system is tuned towards certain feature values. For example, when searching for a book and knowing that its binder is red, the observer might up-weight the feature ‘red’ during search, thereby conferring an advantage to all red objects in the competition for focal attention. The aim of the present review is to summarize the extensive existing evidence that this view might be too simplistic: the dimension-weighting account (DWA) postulates that people in many different situations do not (and sometimes cannot) tune to specific features, but to singletons in specific dimensions.

### **Box 1: Core Assumptions of the DWA**

The Dimension-Weighting Account (DWA) can be summarized by a few core assumptions, all of which have received considerable empirical support, as detailed in the text.

**Integration of saliency signals on and search guidance by the priority map.** The DWA assumes that several saliency maps constitute a subordinate level that inputs to a superordinate (master-saliency or) priority map. Activations at that superordinate map, in turn, guide the allocation of focal attention. Accordingly, DWA is a version/specification of Guided Search (Wolfe 2007).

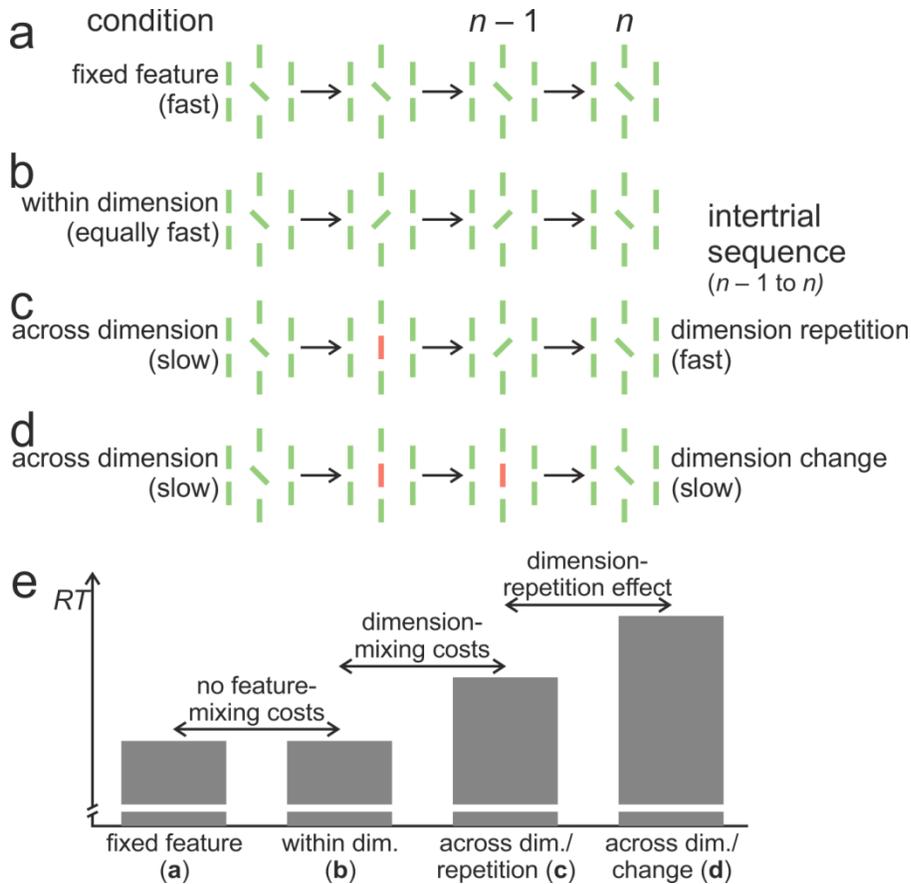
**Dimensional constraints on weighting.** Similar to most Guided-Search-type theories, the DWA assumes that saliency signals are not just summed at the priority map, but that – depending on the current context – some signals are amplified (up-weighted) and others attenuated (down-weighted). The eponymous (and the core distinguishing) feature of the DWA is its emphasis on dimensional constraints of this weighting.

**Involuntary changes in weights.** Whereas most theories focus on goal-driven voluntary control of saliency weighting, the DWA – since its inception – emphasizes involuntary mechanisms (due to search history) that influence these weights and that actually are often irrelevant to the task (because they do not improve performance) or even run counter to search goals (because a different set of weights would help perform the task more efficiently). Of note, both voluntary control and search history are top-down influences (see Gaspelin and Luck 2018c) that are assumed to work via the same mechanism (dimension weighting), but for different reasons (search goals vs. experience).

**Pre-attentive locus of weighting.** Another strong emphasis of the DWA lies on the idea that weights are set before the search display comes up and influence saliency extraction from the display– either voluntarily or involuntarily. Much work has shown that processing a target involuntarily changes dimensional weight settings that then influence saliency computations on the next trial. As saliency processing, in turn, influences attention allocations, this also means that dimension weighting takes effect at a pre-attentive stage (with respect to the current trial) of stimulus processing.

DWA started off with an intriguing observation: search responses to the exact same target item were slower when the target could be either an orientation, size, or color singleton, with the odd-one-out target item changing unpredictably across trials (across-dimension condition), compared to when it was always a

color singleton and the specific color changed across trials (within-dimension condition; Egeth 1977, p. 300; Treisman 1988, pp. 207-210). In both cases, it is by design not possible to prepare for a specific target feature, but it is possible to prepare for a specific target dimension in the latter, easier condition (namely, the color dimension). A follow-up study replicated the effect and additionally showed that keeping the target constant at a specific feature value did not improve performance beyond keeping the dimension constant (see Fig. 3; Müller et al. 1995; but see Wolfe et al. 2003). Thus, it seems that observers can prepare for (i.e., up-weight) a specific target dimension, but not a specific target feature (hence Dimension-Weighting Account).



**Fig. 3** Empirical findings that have inspired and supported the DWA. Search is faster when the physically identical target (on trial  $n$ ) is defined in a constant dimension within a block (**a** and **b** vs. **c** and **d**), whereas it does not matter whether the target is always defined by the same feature (**a** vs. **b**). With variable targets (across-dimension condition, **c** and **d**), search is (relatively) fast when the target dimension repeats and slow when the dimension changes across trials (**c** vs. **d**). **a-d** only illustrate local contrast; target position is unpredictable across trials. **e** illustrates the various effects on RTs which support the DWA.

## 2 The impact of the preceding trial: Search-history effects on dimensional weights

The studies reviewed above (Egeth 1977; Müller et al. 1995; Treisman 1988; Wolfe et al. 2003) showed that observers cannot effectively prepare for a target if they do not know in advance in which dimension the target will differ from its surround in the across-dimension condition (Fig. 3c-d); in other words, no goal-driven, voluntary control is possible when the target's singleton dimension is not predictable. Interestingly, however, attention allocations in this condition are not merely bottom-up driven either:

examining how the target dimension on the previous trial (trial  $n - 1$ ) influences performance on the current trial (trial  $n$ ), Müller et al. (1995; see Fig. 3c-e) found that responses are faster when the dimension repeats (e.g., orientation [on trial  $n - 1$ ]  $\rightarrow$  orientation [on trial  $n$ ]) compared to when it changes (color  $\rightarrow$  orientation). This effect has been taken to indicate that processing of the target (dimension) on trial  $n - 1$  somehow changes the dimensional weight settings and that these changes persist into trial  $n$ . Persistence of the weights is likely involuntary, given that the dimension on trial  $n - 1$  did not predict the dimension on trial  $n$  (i.e., there was no incentive to keep these settings). This indicated a third influence beyond bottom-up saliency and voluntary control that was dubbed *intertrial effects* and constitutes a particularly powerful form of (immediate) *search history* (Allemark et al. 2018; see also Awh et al. 2012; Theeuwes 2018; Wolfe and Horowitz 2017). Found and Müller (1996) went on to confirm the weight-shifting account and showed that, just like the mixing costs reviewed above (Fig. 3a, b, and e), this effect is largely dimension-specific rather than feature-specific in nature: repeating the exact target feature (e.g., left-tilted  $\rightarrow$  left-tilted) did not (typically) improve performance beyond repeating the dimension (e.g., right-tilted  $\rightarrow$  left-tilted). Notably, this *dimension-repetition effect* emerges even though observers are not usually aware of the target dimension on a given trial (Müller et al. 2004), indicating that it reflects a largely implicit memory effect.

A strong claim of the DWA is that dimensional weighting influences (via the priority map) the allocation of attention (Box 1). This means that – with respect to the search display on trial  $n$  – dimension-repetition effects must emerge at a pre-attentive stage of processing. Töllner et al. (2008; see also Gramann et al. 2010; Töllner et al. 2010) used ERPs to directly test this claim: Observers searched for a singleton target object that was either a blue square or a red circle among blue circles and responded to the orientation of its object-internal grating (vertical vs. horizontal) while their EEG was recorded. The most important finding was that the N2pc component (see Box 2 and Fig. 4) emerged earlier for dimension repetitions than for dimension changes, indicating that dimension-repetition effects influence the very first allocation of covert attention (for a review of converging behavioral evidence, see Krummenacher and Müller 2012).

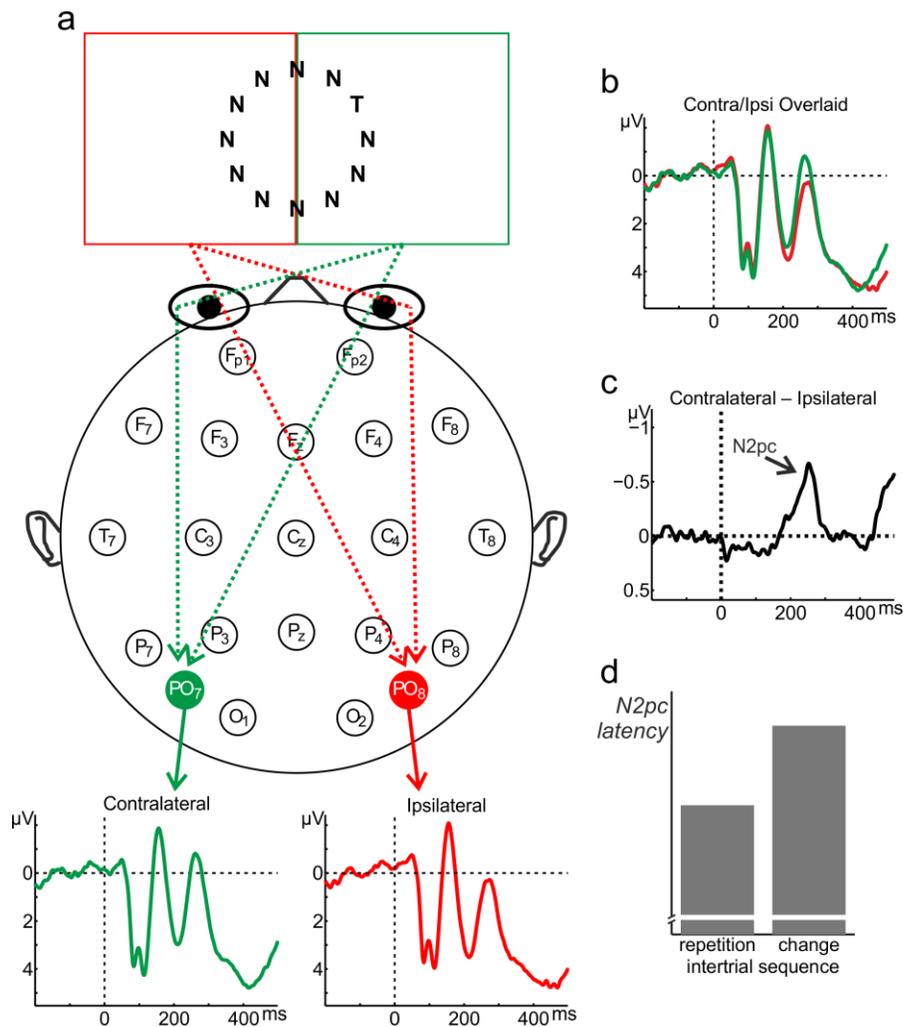
## Box 2: Obtaining neuroscientific evidence for the DWA

Various neuroscientific techniques have helped to elucidate the nature of dimension weighting:

**EEG.** The high temporal resolution of event-related potentials (ERPs) extracted from the electroencephalogram (EEG) allows the examination of cognitive processes that occur before an overt response is executed. This is of considerable importance, because the DWA assumes that the dimensional weights are set before the search display is presented, so that they influence processing already at a pre-attentive stage (see Box 1). Response times are, however, measured only at the end of a trial when attention has already been allocated to the target. The N2pc component (posterior-contralateral negativity in the N2 range; sometimes referred to simply as posterior-contralateral negativity, PCN) is a validated marker of spatial allocation of *covert attention* and therefore particularly useful for examinations of attentional dynamics (Eimer 1996; Luck and Hillyard 1994a,b). In particular, the N2pc can be used to determine (a) whether attention was allocated to a particular object and (b) when in time attention was allocated. Similar (though not identical) information can be gained from eye tracking by measuring *overt attention* allocations (fixations; e.g., Geyer et al. 2008).

**fMRI.** The high spatial resolution of functional magnetic resonance imaging (fMRI) permits identification of brain areas that are selectively activated by experimental manipulations. This has led to the mapping of a large network of brain areas, including the fronto-parietal attention network (e.g., Corbetta and Shulman 2002), that are active during the dimension-weighting process. The consequence of this process (up-weighting of certain dimensions), in turn, is reflected by increased activation in dimension-specific brain areas.

**Lesions and TMS.** Existing or induced changes of local neuronal function by lesions or transcranial magnetic stimulation (TMS), respectively, have been used to test the functional contribution of brain areas whose activation was observed with fMRI. Both disruption (by lesions) and facilitation (by TMS) of supported a functional contribution of various cortical regions to dimension weighting.



**Fig. 4** Generation of the N2pc component. **a** The left (red) and right (green) hemifield are initially processed by both eyes, but information on the left hemifield is projected mainly to right cortical regions and vice versa. Consequently, processing of a target (T) placed in the right hemifield is more strongly reflected in EEG activity measured at posterior electrodes contralateral to the target (i.e., left) than at electrodes ipsilateral to the target (i.e., right). Lateralized activity measured at electrode pair PO7/8 is typically used to extract the N2pc component. **b** Some 200 ms after stimulus onset, activity is more negative at contralateral than ipsilateral electrodes. This can be seen by overlaying both traces. **c** Target-related activity, including the N2pc component, is extracted by subtracting

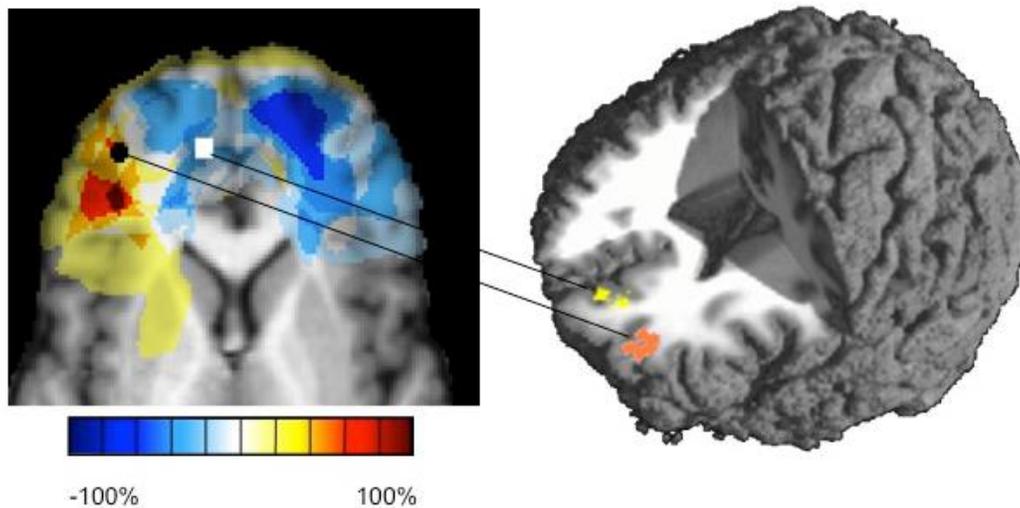
ipsilateral from contralateral activity. **d** The relative time at which the N2pc emerges in different experimental conditions is informative with respect to the timing of attention allocations (Töllner et al., 2011).

To identify the neuronal substrate of dimension weighting, Pollmann et al. (2000) had participants search for a target either defined by color or by movement. In one condition, target identity was kept constant within a block (akin to Egeth 1977, Müller et al. 1995, and Treisman 1988). They observed target-dimension-specific activations in areas involved in the processing of the respective aspects of visual information (see Box 2): in the color block, bilateral areas in the fusiform gyrus (concordant with previous reports of V4) were selectively activated, whereas in the motion block locations in the vicinity of hMT+ were selectively activated. Thus, V4, known to be involved in color processing (e.g., Bartels and Zeki 2000), exhibits increased activation as long as color is up-weighted, and hMT+, known to support motion processing (e.g., Beauchamp et al. 1997), shows increased activation when motion is up-weighted (see also Pollmann et al. 2006b; see Schledde et al. 2017, for converging single-cell evidence).

While dimensional biases would likely be implemented in posterior brain areas that are involved in sensory processing of the respective dimension, (top-down induced) weight changes were expected to activate frontal and parietal areas (e.g., Corbetta and Shulman 2002; Box 2). Indeed, dimension changes elicited widespread activation of the dorsal attention network (Pollmann et al. 2000). In addition to this network, anterior prefrontal cortex, including, most notably, the left frontopolar cortex, was activated following dimension changes (Pollmann et al. 2000; Weidner et al. 2002). At the time, this was quite unexpected, because anterior prefrontal cortex was seen as a high-level executive control structure (e.g., Koechlin et al. 1999), which would likely be insensitive to non-voluntary processes. However, the target-dimension changes in a simple feature-search task that elicited frontopolar activation shared an element of ambiguity with other tasks that involved frontopolar cortex (Burgess et al. 2005). More specifically, while attentional weighting of the new target dimension was by no means required by the rules of the search task, it is an adaptive process – facilitating processing of the next target in the new dimension – that might be supported by frontopolar cortex (Pollmann 2004).

The causal role of lateral frontopolar cortex was confirmed in a lesion study (Pollmann et al. 2007; Box 2): During across-dimension search, patients with a lesion centered on left lateral frontopolar cortex (Fig. 5) showed a selective increase of reaction times on dimension-change trials – but not on (within-dimension) feature-change trials, thus indicating a functional contribution of left lateral frontopolar cortex to dimension weighting. This selective increase of dimension-change effects was not observed in a control group with frontomedial lesions. Thus, dimension-change-related frontopolar activations had not been epiphenomenal; rather, left lateral frontopolar cortex indeed supports dimension-weighting processes in across-dimension search. This early evidence for frontopolar involvement in attention weighting fits well with the developing concept of a critical role of anterior prefrontal cortex in resource exploration versus exploitation, which is based on recent findings across paradigms and across human as well as non-human primates (e.g., Pollmann 2016; Raja Beharelle et al. 2015).

The causal role of right inferior parietal cortex was confirmed as well. Sorting patients according to whether they showed dimension-repetition effects (lesion-symptom mapping) revealed that lesions of right inferior parietal cortex – one of the structures with dimension-change related fMRI activation (Pollmann et al. 2000) – prevented dimension weighting (Utz et al. 2013). In turn, repetitive TMS over right angular gyrus facilitated subsequent stimulus processing selectively on same-dimension, same-response trials, thus indicating amplified dimension weighting (Bocca et al. 2015).



**Fig. 5** The involvement of frontal cortex in dimension weighting. The right side shows fMRI data by Weidner et al. (2002). Left lateral frontopolar cortex (marked red) showed selective BOLD increase for dimension changes, but not for (within-dimension) feature changes during feature search. In contrast, left frontomedial cortex (yellow) showed the same activation pattern during conjunction search. Left side: The causal involvement of left lateral frontopolar cortex in dimension weighting during feature search was shown by selectively increased dimension-change effects in patients with left lateral frontopolar lesions (red) as compared to patients with frontomedian lesions (blue) and matched controls. The color scale indicates lesion density in the patient sample. The black disk indicates the center of the left lateral frontopolar activation and the white square the center of the frontomedial activation observed by Weidner et al. (2002).

Furthermore, Pollmann et al. (2006b) found that dimension-change-related activity precedes response-change-related activity, thus providing converging evidence for an early stage of dimension-repetition/change effects. Taken together these findings indicate that prefrontal areas control the weight shifting between visual dimensions and that areas in parietal and temporal cortices implement these control signals by inducing dimension-specific pre-activations in visual input areas.

That across-dimension effects (Egeth 1977; Müller et al. 1995; Treisman 1988) and dimension-repetition effects (Found and Müller 1996) were observed for pop-out searches actually poses an explanatory challenge for theories of saliency computations, such as the DWA (e.g., Fig. 2). This is because pop-out is the subjective phenomenon that the target stands (or pops) out of the display so that it is found *immediately*, without any need to scrutinize any of the non-target objects (like the black sheep in a flock of white sheep). Thus, one might assume that pop-out search is already maximally efficient, so that no

further speed-up is possible. Search efficiency is typically measured by varying the number of non-target objects in the display and determining the slope of the function relating response times to number of objects (search slope). It has been shown that there is a continuum of search slopes (i.e., search efficiencies) across different visual-search tasks (Wolfe 1998) and within the same task dependent on a variation of target saliency (Duncan and Humphreys 1989; Liesefeld et al. 2016; Nothdurft 1993). Some might consider pop-out to be one pole of this continuum. Liesefeld et al. (2016), however, showed that search can speed up beyond the point where the target reliably pops out. Their observers searched for a tilted target bar among vertical non-target bars. Saliency was manipulated by varying the tilt of the target in fine steps from 1.5° to 45°; the resulting search slopes ranged from 70 ms/object (very inefficient) to 0 ms/object (very efficient). Furthermore, and of critical importance for the present point, absolute response times decreased even in the range of very efficient searches (flat search slopes, i.e., ~0 ms/object) with a further increase in saliency. These decreases ranged from 615 ms to 440 ms (6° and 45° targets, respectively), showing that there is indeed quite some leeway (~175 ms) for further speed-up even in the efficient (pop-out) range.

Töllner et al. (2011) showed that (some of) this leeway emerges at a pre-attentive stage (see Box 1): They manipulated target saliency while making sure (via pilot studies) that all targets were still within the pop-out range. Measured N2pc latency indeed decreased with increasing saliency, indicating that attention allocations were speeded. Notwithstanding the general speed-up with increases in saliency, the Liesefeld et al. (2016) data also showed that the effect of saliency is not linear, but approaches an asymptote with higher saliencies, indicating that the probability of finding any effect on search times decreases with the bottom-up saliency of the target object: the more salient a target is, the less leeway there is for further speed-ups. Indeed, Zehetleitner et al. (2011) observed smaller dimension-repetition effects for very salient objects compared to medium-salient objects (all within the pop-out range), for which an additional increase in priority due to favorable dimensional weight settings would yield only marginal extra benefits. A general take-home message from the Zehetleitner-et-al. and the Liesefeld-et-al. studies is that the absence (or non-significance) of effects on search times might sometimes simply reflect that target priority is already maximal (at ceiling) in the less efficient condition (e.g., dimension-change condition), a situation that should be avoided if sensitive statistical tests are desired. However, the Liesefeld et al. study also shows that this ceiling is not yet reached when search slopes level off (when the target pops out), thus supporting the interpretation that dimension repetitions do speed up search proper even in this very efficient search range.

Dimension-repetition effects do not only occur for selected dimensions, but are observed for almost all dimensions that have so far unambiguously been shown to guide search (Wolfe and Horowitz 2004, 2017). In Guided Search (Wolfe 2007), attributes are considered ‘basic’ or ‘guiding’ when targets defined by that attribute are found efficiently (pop out; see above). According to Wolfe and Horowitz (2017; see their Box 1), ‘undoubted’ guiding attributes are: color, motion, orientation, and size. Indeed, dimension-repetition effects were observed for color (e.g., Found and Müller 1996; Müller et al. 2003; Krummenacher et al. 2009; Töllner et al. 2008; Pollmann et al. 2000; Weidner and Müller 2013; Weidner et al. 2002; Zehetleitner et al. 2011), motion (e.g., Pollmann et al. 2000; Weidner and Müller 2013; Weidner et al. 2002), and orientation (e.g., Found and Müller 1996; Müller et al. 2003; Krummenacher et al. 2009; Zehetleitner et al. 2011). Furthermore, dimension-repetition effects were shown for ‘probable’ (according to Wolfe and Horowitz 2017) guiding attributes such as shape (Töllner et al. 2008) and luminance polarity (Zehetleitner et al. 2011). This renders the DWA a general account of visual search, not limited to a single dimension (see also Section 5).

Notwithstanding the strong focus on pop-out search, dimension-repetition effects are assumed to influence saliency computations in general (Box 1). Therefore, they should also occur for inefficient searches. Indeed, Weidner and Müller (2013) observed dimension-repetition effects in a conjunction-search task. In this type of task, the target is defined by a conjunction of two features (e.g., a red, vertical bar among green, vertical and red, horizontal bars) and search is inefficient even if the individual features are very salient (e.g., Wolfe 1998; Wolfe et al. 2010). Weidner and Müller's participants searched for a target that was (unpredictably across trials) defined by either size (large vs. small) *and* motion direction (diagonal vs. horizontal oscillatory motion) or by size (large vs. small) *and* color (red/blue vs. green). That is, one of the two target-defining dimensions stayed constant across trials (size; primary dimension) whereas the other dimension varied (motion direction vs. color; secondary dimension). Replicating findings from efficient feature-search tasks, responses were speeded when the secondary dimension repeated (e.g., size and color → size and color) relative to when it changed (e.g., size and motion direction → size and color).

The neuronal sources of dimension weighting might be somewhat different for conjunction compared to feature search (Fig. 5). In contrast to the Pollmann et al. (2000) study reviewed above, Weidner et al. (2002) employed a conjunction-search task comparable to that of Weidner and Müller (2013) and found dimension-change-related activation in pregenual frontomedian cortex (instead of left lateral frontopolar cortex found by Pollmann 2000 during feature search). In a further experiment, Weidner et al. replicated this dissociation between conjunction and feature search within the same participants. In more posterior brain areas, in contrast, dimension-repetition/change-related activation was comparable across studies, matching the known structural specialization of the visual system (see Box 2).

Two or more dimensions can be up-weighted concurrently, but at a cost for each individual dimension. In the Weidner and Müller (2013) study just reviewed, orientation (arguably) had to be up-weighted throughout and still dimension-repetition effects emerged for the 'secondary' dimension (size vs. motion direction). Krummenacher et al. (2001, 2002a,b) examined situations in which the target was a feature singleton on each trial (feature search), but differed either in one or in two features from its surround (because each feature on its own already sufficiently defines the target [in contrast to the Weidner and Müller 2013, conjunction task], this is referred to as *redundant-signals paradigm*). Most important for the present review, Krummenacher et al. found that up-weighting one dimension comes at a cost for other dimensions: when a single-feature trial was preceded by a redundant-feature trial (e.g., orientation *and* color → color), responses were slower compared to when the single-feature trial was preceded by a single-feature trial from the same dimension (e.g., color → color). Responses on a single-feature trial were, however, faster when preceded by a redundant-feature trial (e.g., orientation *and* color → color) compared to when preceded by a single-feature trial from the other dimension (e.g., orientation → color). This indicates that processing a redundantly defined target on trial  $n - 1$  increases the weights for both feature dimensions, but not as much as a single-feature trial increases the weight for its particular dimension. This pattern of results can be explained by either of two ideas: (a) a limit to the total attentional weight that is shared between dimensions if multiple dimensions are relevant for the search, or (b) it is not the absolute weight assigned to a given dimension that determines target priority, but the relative weight with respect to all other feature dimensions (a form of divisive normalization; e.g., Carandini and Heeger 2012). Differentiating these two ideas is an interesting avenue for future research.

The strength of behavioral dimension-repetition effects is somewhat modulated by the type of task. When the task is not only to detect whether a target is present, but to classify the target according to some (other) property not critical for basic detection (e.g., find a tilted bar among vertical distractors and decide

whether that bar has a gap on its top or bottom; see Fig. 2), or to localize the target (e.g., left vs. right side of the display), dimension-repetition effects manifest in behavioral, response-time measures are reduced or abolished (Chan and Hayward 2009; Müller and Krummenacher 2006; Kumada 2001; Krummenacher et al., 2009; Krummenacher et al., 2002b; Pollmann et al. 2000, 2006b; Theeuwes et al. 2006; Zehetleitner et al. 2011).

To determine the temporal locus of the dimension-repetition effect in the classification task, Töllner et al. (2008) directly measured attention allocations via the N2pc (see Box 2). As predicted by the DWA, allocation of attention was speeded for dimension repetitions compared with dimension changes. Also dimension-repetition/change-related fMRI activation in dimension-specific cortical areas comparable to that found in a detection task (Pollmann et al. 2000; see Box 2) was observed in a classification task as well (Pollmann et al. 2006a). This indicates that dimension-repetition effects emerge at a pre-attentive stage, independently of the search task.

Töllner et al. (2012a) went on to show that the timing of the N2pc is unaffected by the task demands, whether the task requires simple target detection or localization or classification. However, task demands influence a later, post-selective stage at which response-relevant information is extracted from the attended stimulus to decide on the appropriate motor action (stimulus-response, S-R, mapping). Töllner et al. (2008) provided electrophysiological evidence (based on comparing the timing of the N2pc with the *stimulus*-locked lateralized response potential) that this stage is subject to implicit ‘linked-expectancy’ effects: repetition of the target-defining dimension biases S-R mappings towards the same response issued on the previous trial (whereas a target-dimension change may induce a bias towards a different response). Accordingly, response selection is fast when both the target-defining dimension and the response-critical target property repeat; but response selection is slow when one property repeats and the other changes (partial-repetition costs), because this requires the overcoming of a dimension-induced S-R expectancy. Such partial-repetition costs (Hommel 1998; Hommel et al. 2001) seem to cancel out dimension-repetition effects on response-change trials, thus yielding an underestimation of the ‘true’ dimension-repetition (vs. -change) effect in averaged classification-search data (Müller and Krummenacher 2006; Pollman et al., 2006a).

A complementary approach to measuring response times is to present search displays only briefly followed by a mask and then measure accuracies. The mask serves to erase iconic memory of the display and thereby effectively abort search at an experimentally controlled point in time. The minimum time between search display and mask onset that is needed to achieve a certain accuracy level is a measure of the time taken for visual search proper, excluding processes occurring at the response-selection stage, so that any effects can be less ambiguously attributed to a perceptual stage of processing (see previous paragraph and Box 1). Using this approach, Zehetleitner et al. (2011, Exp. 1) presented search displays for about 65 ms, immediately followed by a mask, and participants had to indicate whether a target bar (luminance or orientation singleton) was presented on the right or the left of the display. Replicating RT results, localization accuracy was indeed higher on dimension-repetition than on dimension-change trials, thus providing additional evidence for a perceptual stage of dimension-repetition effects.

A similar, yet in a crucial aspect distinct, intertrial effect was discovered at about the same time as dimension-repetition effects and was termed ‘priming of pop-out’ (PoP; Lamy et al. 2011; Maljkovic and Nakayama 1994): Similar to most of the experiments reviewed above, the target in a typical PoP task is defined as the only object of a particular feature (a color singleton in most experiments). Crucially, however, only very few other (non-target) objects are typically present and target and non-target features

swap randomly across trials. In contrast to the studies reviewed above (where the non-targets were kept constant and numerous), there was a huge effect of whether the target (and therefore also non-target) *feature* was (a) predictable (constant across all trials; see also Bravo and Nakayama 1992) and (b) repeated or swapped relative to the previous trial, even though the target *dimension* was fixed. That is, in contrast to the dimension-specific effects reviewed above, PoP is a feature-specific effect. The different outcomes of the two designs were explained by Rangelov et al. (2013; see also Krummenacher et al. 2010; Rangelov et al. 2011a,b; Zehetleitner et al. 2012): The feature-specific PoP vanished when more non-targets were presented. They argued (and later showed more directly; see Rangelov et al. 2017) that with a (for PoP studies typical) set size of 3 objects (1 target and 2 non-targets), targets often do actually not pop out, that is, often distractors are attended before the target (see also Becker 2008). One interpretation of these findings is that the target is not particularly salient in sparse displays (it lacks local contrast; see Liesefeld et al. 2016; Nothdurft 1993). If the target is not sufficiently salient to reliably draw attention on every trial, more elaborate processing might become necessary before a response is issued. This would be needed to ensure that the attended object is indeed the target, so as to avoid erroneous responses to a non-target, and (logically) this may require a process of feature comparison of some attended object against the other items. PoP might emerge because repeatedly processing the same target feature speeds up this validation process. On this interpretation, ‘priming of pop-out’ (in contrast to dimension-repetition effects) is not, or only to a limited degree, driven by saliency-computation mechanisms and would be better described as priming of (post-selective) target-feature validation (see Huang et al. 2004).

### **3 Partial control over dimensional weights: Interactions of search history and voluntary control**

In the preceding section, we summarized evidence that dimensional weights do shift without the observer’s explicit intention in situations where changing the weight settings confers no particular advantage. We will now review how far dimensional weighting can occur when the observer knows that changing the weights is advantageous for an upcoming search, so that weight changes are intended by the observer or at least in line with search goals. In other words, we will address voluntary control over dimensional weights based on task goals.

Also starting from the inception of DWA, it was clear that dimensional weights are, indeed, subject to voluntary control: Müller et al. (1995, Exp. 3) found that when observers knew that one target dimension appeared more likely in an across-dimension search task, performance for these targets improved dramatically. The flexibility of this voluntary preparation was later examined directly by, on each trial, informing participants about the upcoming target in an across-dimension search task (Müller et al. 2003; see also Wolfe et al. 2003; Zehetleitner et al. 2011): The target could be either one of two orientation singletons (20° tilted to the left or right) or one of two color singletons (red or blue), and a symbolic (written) cue was given before onset of the search display that informed participants about the singleton dimension of the upcoming target with 80% validity. For example, when observers saw the word “color”, the upcoming target was a color singleton in 80% of cases and an orientation singleton in 20%. Compared to a neutral baseline (the cue word “neutral” was shown), responses were speeded for valid and slowed for invalid trials, indicating that participants were able to voluntarily prepare for the upcoming target. In a second experiment, the specific target feature was cued with a 79% validity using the words “red”, “blue”, “left”, and “right”. As predicted by the DWA, cueing a specific feature also expedited responses to another feature from the same dimension on invalid trials (e.g., blue targets were detected faster after the

cue “red”), even though it was equally likely to appear as any feature from the respective other dimension (each 7%).

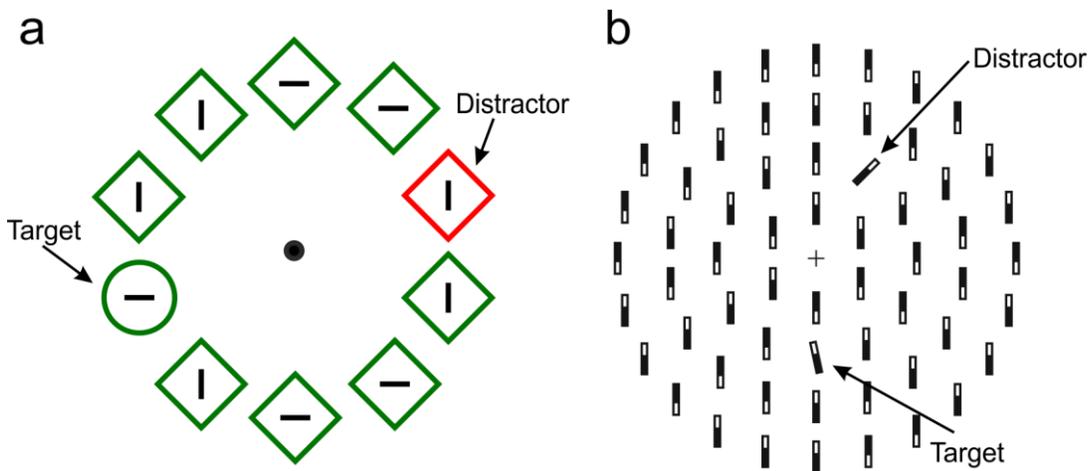
Just like dimension-repetition effects, the strength of the cueing advantage depends on the type of task. Cueing effects are sometimes abolished when the task is not just to detect whether a target is present, but to classify one of its features (classification task) or to localize it. This response-dependence of the effect might indicate that cueing effects occur at a post-selective, response-related stage (Mortier et al. 2010; Theeuwes et al. 2006). However, in line with a pre-attentive origin (Box 1), cueing effects were re-established by increasing task difficulty (Zehetleitner et al. 2011) and by increasing the incentive to use the cue (Müller and Krummenacher 2006). Again, ERPs provided a crucial piece of evidence in this debate by demonstrating that the cueing advantage due to dimensional weighting has a pre-attentive origin, even in a classification task (Töllner et al. 2010). Similar to Müller et al. (2003), Töllner et al. (2010) presented a cue word (color, shape, or neutral), followed by a short inter-stimulus interval and a search display featuring either a color- or an orientation-singleton target. This time, the task was to identify the orientation of the target (horizontal or vertical grating), instead of simply detecting the presence of a target (as in Müller et al. 2003). Besides replicating the Müller-et-al. (2003) cueing effect on RTs with this classification task (cf. Theeuwes et al. 2006; see also Müller and Krummenacher 2006), they also showed that valid cues speeded attention allocations as indicated by expedited (and increased) N2pc relative to trials with invalid cues and, thus, confirmed a pre-attentive origin of the cueing effect.

These findings clearly show that dimensional weight settings are susceptible to voluntary control. If there are two influences on these weight settings, voluntary control and search history, one might wonder which of the two is more potent when pitted against each other. To examine this question, Müller et al. (2003; see also Zehetleitner et al. 2011) compared the dimension-repetition effects on trials with neutral and with informative cues. As expected, strong dimension-repetition effects occurred when cues were uninformative (neutral) and could therefore not influence the weight settings. These effects were reduced, but not completely abolished when cues were informative, both when the informative cue was valid- and when it was invalid, indicating that search-history effects persist, even if observers try to voluntarily control their weight settings. In fact, dimension-repetition effects remained even when the cue was 100% valid, thus removing any incentive to up-weight the preceding target dimension instead of the cued dimension.

Weidner et al. (2009) examined the neuronal correlates of voluntary weight shifts in a paradigm similar to Müller et al. (2003) and added a manipulation of target saliency. Decreases in saliency (and therefore in search efficiency) yielded activity increases in the fronto-parietal attention network and the right anterior middle frontal gyrus and activity decreases in the medial anterior prefrontal cortex. Having prepared for the wrong target dimension (because of an invalid cue) caused activity increases in the left lateral frontopolar cortex, the left supramarginal gyrus and the cerebellum as well as bilaterally in the posterior orbital gyrus, the inferior frontal gyrus, and the pre-SMA and decreases (with respect to validly cued trials) in the superior frontal gyrus, anterior to the junction with the precentral gyrus (frontal eye fields, FEF). These areas might be involved in re-organizing the weight settings when the current settings are invalid. Furthermore, an interaction between saliency and cue validity was observed in the left temporo-parietal junction, rendering this a probable site for integration of saliency signals and dimensional weight settings as would be expected to occur at the priority map (see also Serences et al. 2005).

#### 4 Handling distraction by setting dimensional weights

The top-down setting of dimensional weights also allows successful handling of salient-but-irrelevant distractors. In his classical studies, Theeuwes (1991, 1992) showed that search for a singleton pop-out target is hampered by the presence of a salient distractor that is a singleton in a different dimension (e.g., a color distractor during search for a shape target; see Fig. 6a). This was interpreted as evidence for attentional capture: attention is involuntarily allocated first towards the salient distractor, before it can be re-allocated to the target. The first misallocation (capture) of attention causes a delay of target processing and therefore produces response-time costs. However, more recent evidence suggests that such a distractor does not typically capture attention, but is suppressed (at least when target and distractor are constant across trials, so that observers can prepare effectively; Burra and Kerzel 2013; Hickey, McDonald, & Theeuwes, 2006; Kerzel and Barras 2016; Kiss, Grubert, Petersen, & Eimer, 2012; McDonald et al., 2013). Jannati et al. (2013), for example, showed that a color distractor during search for a shape target (as displayed in Fig. 5a) does not elicit an N2pc, but a  $P_D$  indicating suppression of the distractor instead of attentional capture (Hickey et al. 2009; Sawaki et al., 2012; Toffanin et al. 2011).



**Fig. 6** Search-displays used to examine attentional capture. **a** An example of a typical additional-singleton task (as used by Jannati et al. 2013). Observers had to determine the orientation of the line within the circle (vertical vs. horizontal). **b** The additional-singleton task used by Liesefeld et al. (2017). Observers had to determine the position of the notch in the bar tilted  $12^\circ$  to the left (top vs. bottom). Non-targets (green diamonds and vertical bars, respectively) and singleton-distractor (red diamond and bar tilted  $45^\circ$  to the right, respectively) were completely irrelevant in both tasks

That the distractor in Jannati et al. (2013, and many similar studies) does not have a strong potential to capture attention follows from the DWA: In preparation for a search, observers either up-weight the target dimension (shape) so that the target is the object achieving the highest value on the priority map, and/or they down-weight the distractor dimension (color) so that the distractor does not stand out on the priority map. The observed evidence for suppression (Jannati et al. 2013; Gaspar and McDonald 2014; Gaspelin and Luck 2018a,b) might indicate that this weighting is not perfect and the distractor produces some residual activation on the priority map that must be suppressed to pave the way for an attention allocation to the target (i.e., implementing a winner-takes-all mechanism; Desimone and Duncan 1995).

Assigning a low priority to a salient-but-irrelevant distractor would be impossible or counter-productive if distractor and target are coupled to the same weight. According to the DWA, this is the case when both singletons are defined in the same dimension. Any attempt to down-weight the distractor (dimension)

would also down-weight the target (dimension) and would therefore obstruct search. As a consequence, such a same-dimension distractor should reliably capture attention. Liesefeld et al. (2017) directly tested this prediction with a 45° distractor and a 12° target, tilted into opposite directions (see Fig. 6b). The notion of attentional capture implies that (a) attention is first allocated to the distractor and only afterwards to the target, and (b) that attention allocation towards the target is delayed with respect to a search display without a distractor (where attention is allocated directly to the target). Indeed, Liesefeld et al. observed an N2pc to the distractor that preceded an N2pc to the target, and the target N2pc on distractor-present trials was delayed with respect to the target N2pc on distractor-absent trials. Furthermore, distractor presence delayed responses by more than 200 ms! Thus, when the target is constant across trials, a different-dimension distractor is unlikely to capture attention, but a same-dimension distractor reliably does (see also Schubö and Müller 2009; Töllner et al. 2012b; for behavioral evidence, see Liesefeld et al. under review; Sauter et al. 2018, in revision; for a focused review on dimension weighting as a mechanism for distractor handling, see Liesefeld and Müller, under review).

## 5 Is color special?

Even the first studies on DWA already observed an apparent idiosyncrasy for one particular feature dimension: color. As for other dimensions, responses were speeded for a color target preceded by a color target compared to a color target preceded by, say, an orientation target. In contrast to other dimensions such as orientation, however, it mattered whether the exact same color was repeated: There was a further speed-up when a red target was preceded by a red target relative to when a red target was preceded by a blue target (Found and Müller 1996). This peculiarity of color was also observed for cueing studies: Whereas cueing the exact feature did not matter for orientation, cueing a specific color speeded up responses when the cue matched the subsequent target (Müller et al. 2003). Furthermore, in contrast to orientation distractors during search for orientation targets (Liesefeld et al. 2017), salient color distractors do not (reliably) capture attention during search for color targets (Gaspar and McDonald 2014; Gaspar et al. 2016; but see Weichselbaum and Ansorge 2018).

Much of the confusion probably arises from treating ‘color’ as one dimension whereas it is in fact a multi-dimensional property. Already at early stages of color processing, retinal ganglion cells (as transmitted to the early visual cortical areas via the thalamic lateral geniculate nucleus) represent color in a three-dimensional space (De Valois et al. 1966; Derrington et al. 1984). A similar dimensional structure is implemented in the psychologically meaningful CIE *Lab* color space, where color is defined along the dimensions luminance (*L*), red-green (*a*) and blue-yellow (*b*). It is thus far unknown how this multi-dimensional space is structured for saliency computations; and this is further complicated by the possibility that it has more than three dimensions (see D’Zmura 1991; Lindsey et al. 2010; but see Martinovic et al. 2018). But let us, for sake of the argument, speculate using the three dimensions of *Lab* space: even if, in most experiments, one dimension is kept constant (by using equiluminant stimuli), the employed colors still typically vary on (at least) two dimensions. Gaspar and McDonald (2014) for example had observers search for a yellow target among green distractors and found that a red distractor does not capture attention. It is quite evident that in our hypothetical example search for a yellow target would be supported by up-weighting yellow-blue (*b*) and a red distractor is best ignored by down-weighting red-green (*a*). Thus, red distractors can be sufficiently down-weighted to avoid attentional capture during search for yellow targets among green non-targets (Gaspar and McDonald 2014; for a related, though somewhat different, explanation, see the literature on linear separability; Bauer et al. 1996a,b; Daoutis et al. 2006; Kong et al. 2016).

Notwithstanding the above, the various dimensions of color space are certainly not treated fully independently. Separate but non-independent dimensions are typically referred to as ‘integral’ (Algom and Fitousi 2016; Garner 1974). The empirical criterion for integrality is that during classification according to one dimension, irrelevant variability in the other dimension cannot be (fully) ignored (thus hampering classification). It appears likely that dimension weighting typically spills over to integral dimensions, but can be restricted to single dimensions if necessary. This would explain both dimension-repetition and feature-repetition effects (Found and Müller 1996) as well as dimension-cueing and feature-cueing (Müller et al. 2003) effects observed for color.

## 6 The scope of the DWA and its relation to other ideas

The DWA makes the positive claim that the transfer from the dimension-specific saliency maps to the priority map can be biased by dimensional weighting. However, it does not make the negative claim that there are no other nodes in the priority-calculation system that are (under certain conditions) influenced by search history and/or voluntary control, including the possibility of feature weighting. Weighting mechanisms specific to a certain feature, to a certain dimension or even to a certain modality (visual vs. auditory vs. haptic; e.g., Töllner et al. 2009) as well as spatial mechanisms might work in concert, because each mechanism in isolation is less effective and/or less efficient (i.e., would cost more cognitive resources) and because a combination of mechanisms increases the adaptability of the visual system.

*Contingent capture* (Folk and Remington 1998; Folk et al. 1992) in a DWA interpretation means that not only distractors with target *features* capture attention, but also distractors that are singletons in the same *dimension* as the target. Regarding the discussion on search modes (singleton detection vs. feature search; Bacon and Egeth 1994; Leber and Egeth, 2006), we would contend that observers can (also) take on a dimension-search mode. There is, indeed, evidence of contingent capture in a general color-search mode (Folk and Anderson 2010; Folk and Remington 1998; Harris et al. 2015). *Relational coding* (Becker 2010; Becker et al. 2010) assumes that observers search for a singleton that differs from its surround in a particular direction (e.g., redder or larger) instead of searching for any singleton in a given dimension (e.g., red-green singleton or size singleton, without a negative or positive sign). It appears plausible that relational coding is possible for some dimensions, while other dimensions are restricted to general dimension weighting (such as orientation; Liesefeld et al. 2017).

Given these numerous interrelations, future research should identify situations in which one and/or the other mechanism takes effect and how the various mechanisms are related, instead of treating these accounts as competitors. Any attempt to disentangle dimension-based and feature-based mechanisms should rather focus on dimensions that are easily characterized such as orientation and luminance and avoid dimensions that are less so, such as color or shape. Notwithstanding the general compatibility of these approaches, such a research agenda might turn out to unify some (or all) of these ideas.

## 7 References

- Algom D, Fitousi D (2016) Half a century of research on Garner interference and the separability-integrality distinction. *Psychol Bull* 142:1352-1383. doi:10.1037/bul0000072
- Allenmark F, Müller HJ, Shi Z. (2018). Inter-trial effects in visual search: Factorial comparison of Bayesian updating models. *PLoS Comput. Biol.* 14: e1006328. doi:10.1371/journal.pcbi.1006328

- Awh E, Belopolsky AV, Theeuwes J (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends Cogn Sci* 16:437-443. doi:10.1016/j.tics.2012.06.010
- Bacon WF, Egeth HE (1994). Overriding stimulus-driven attentional capture. *Percept Psychophys* 55:485-496. doi:10.3758/BF03205306
- Bartels A, Zeki S (2000). The architecture of the colour centre in the human visual brain: New results and a review. *Eur J Neurosci* 12:172-193. doi:10.1046/j.1460-9568.2000.00905.x
- Bauer B, Jolicoeur P, Cowan WB (1996). Distractor heterogeneity versus linear separability in colour visual search. *Perception* 25:1281-1294. doi:10.1068/p251281
- Bauer B, Jolicoeur P, Cowan WB (1996). Visual search for colour targets that are or are not linearly separable from distractors. *Vision Res* 36:1439-1465. doi:10.1016/0042-6989(95)00207-3
- Becker SI (2008). Can intertrial effects of features and dimensions be explained by a single theory?. *J Exp Psychol Hum Percept Perform* 34:1417-1440. doi:10.1037/a001138
- Becker SI (2010). The role of target-distractor relationships in guiding attention and the eyes in visual search. *J Exp Psychol Gen* 139:247-265. doi:10.1037/a0018808
- Becker SI, Folk CL, Remington RW (2010). The role of relational information in contingent capture. *J Exp Psychol Hum Percept Perform* 36:1460-1476. doi:10.1037/a0020370
- Beauchamp MS, Cox RW, DeYoe EA (1997). Graded effects of spatial and featural attention on human area MT and associated motion processing areas. *J Neurophysiol* 78:516-520. doi:10.1152/jn.1997.78.1.516
- Bocca F, Töllner T, Müller HJ, Taylor PC (2015). The right angular gyrus combines perceptual and response-related expectancies in visual search: TMS-EEG evidence. *Brain Stimul* 8:816-822 doi:10.1016/j.brs.2015.02.001
- Bravo MJ, Nakayama K (1992). The role of attention in different visual-search tasks. *Percept Psychophys* 51:465-472. doi:10.3758/BF03211642
- Bundesen C, Habekost T, Kyllingsbæk S (2011). A neural theory of visual attention and short-term memory (NTVA). *Neuropsychologia* 49:1446-1457. doi:10.1016/j.neuropsychologia.2010.12.006
- Burra N, Kerzel D (2013). Attentional capture during visual search is attenuated by target predictability: Evidence from the N2pc, Pd, and topographic segmentation. *Psychophysiology* 50:422-430. doi:10.1111/psyp.12019
- Burgess PW, Simoins, JS, Dumontheil I, Gilbert SJ (2005). The gateway hypothesis of rostral prefrontal cortex (area 10) function. In: Duncan J, Phillips L, McLeod P (eds) *Measuring the mind: speed, control, and age*. Oxford University Press, Oxford, p 217-248
- Carandini M, Heeger DJ (2012). Normalization as a canonical neural computation. *Nat Rev Neurosci* 13:51-62. doi:10.1038/nrn3136
- Chan LH, Hayward WG (2009). Feature integration theory revisited: Dissociating feature detection and attentional guidance in visual search. *J Exp Psychol: Hum Percept Perform* 35:119-132. doi:10.1037/0096-1523.35.1.119
- Corbetta M, Shulman GL (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* 3:201-215. doi:10.1038/nrn755
- Daoutis CA, Pilling M, Davies IL (2006). Categorical effects in visual search for colour. *Vis Cogn* 14:217-240. doi:10.1080/13506280500158670
- Desimone R, Duncan J (1995). Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193-222. doi:10.1146/annurev.ne.18.030195.001205
- De Valois RL, Abramov I, Jacobs GH (1966). Analysis of response patterns of LGN cells. *J Opt Soc Am* 56:966-977. doi:10.1364/JOSA.56.000966

- Derrington AM, Krauskopf J, Lennie P (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *J Physiol* 357:241-265. doi:10.1113/jphysiol.1984.sp015499
- Duncan J, Humphreys GW (1989). Visual search and stimulus similarity. *Psychol Rev* 96:433-458. doi:10.1037/0033-295X.96.3.433
- D'Zmura, M. (1991). Color in visual search. *Vision Res* 31:951-966. doi:10.1016/0042-6989(91)90203-H
- Egeth H (1977). Attention and preattention. In: Bower GH (ed) *The psychology of learning and motivation*, vol 11. Academic Press, New York, pp. 277–320
- Eimer M. (1996). The N2pc component as an indicator of attentional selectivity. *Electroencephalogr Clin Neurophysiol* 99:225-234. doi:10.1016/0013-4694(96)95711-9
- Folk CL, Anderson BA (2010). Target-uncertainty effects in attentional capture: Color-singleton set or multiple attentional control settings?. *Psychon Bull Rev* 17:421-426. doi:10.3758/PBR.17.3.421
- Folk CL, Remington R (1998). Selectivity in distraction by irrelevant featural singletons: Evidence for two forms of attentional capture. *J Exp Psychol Hum Percept Perform* 24:847-858. doi:10.1037/0096-1523.24.3.847
- Folk CL, Remington RW, Johnston JC (1992). Involuntary covert orienting is contingent on attentional control settings. *J Exp Psychol Hum Percept Perform* 18:1030-1044. doi:10.1037/0096-1523.18.4.1030
- Found A, Müller HJ (1996). Searching for unknown feature targets on more than one dimension: Investigating a 'dimension-weighting' account. *Percept Psychophys* 58:88-101. doi:10.3758/BF03205479
- Garner WR (1974). *The processing of information and structure*. Erlbaum, Potomac
- Gaspar JM, Christie GJ, Prime DJ, Jolicoeur P, McDonald JJ (2016). Inability to suppress salient distractors predicts low visual working memory capacity. *Proc Natl Acad Sci U S A* 113:3693-3698. doi:10.1073/pnas.1523471113
- Gaspar JM, McDonald JJ (2014). Suppression of salient objects prevents distraction in visual search. *J Neurosci* 34:5658-5666. doi:10.1523/JNEUROSCI.4161-13.2014
- Gaspelin N, Luck SJ (2018a). Combined electrophysiological and behavioral evidence for the suppression of salient distractors. *J Cogn Neurosci* 30:1265-1280. doi:10.1162/jocn\_a\_01279
- Gaspelin, N, Luck SJ (2018b). The role of inhibition in avoiding distraction by salient stimuli. *Trends Cogn Sci* 22:79-92. doi:10.1016/j.tics.2017.11.001
- Gaspelin N, Luck SJ (2018c). “Top-down” does not mean “voluntary”. *J Cogn* 1:25. doi:10.5334/joc.28
- Geyer T, Müller HJ, Krummenacher J (2008). Expectancies modulate attentional capture by salient color singletons. *Vision Res* 48:1315-1326. doi:10.1016/j.visres.2008.02.006
- Gramann K, Töllner T, Müller HJ (2010). Dimension-based attention modulates early visual processing. *Psychophysiology* 47:968-978. doi:10.1111/j.1469-8986.2010.00998.x
- Harris AM, Becker SI, Remington RW (2015). Capture by colour: Evidence for dimension-specific singleton capture. *Atten Percept Psychophys* 77:2305-2321. doi:10.3758/s13414-015-0927-0
- Hickey C, McDonald JJ, Theeuwes, J (2006). Electrophysiological evidence of the capture of visual attention. *J Cogn Neurosci* 18:604-613. doi:10.1162/jocn.2006.18.4.604
- Hommel B (1998). Event files: Evidence for automatic integration of stimulus-response episodes. *Vis Cogn* 5:183-216. doi:10.1080/713756773
- Hommel B, Müsseler J, Aschersleben G, Prinz W (2001). The Theory of Event Coding (TEC): A framework for perception and action planning. *Behav Brain Sci* 24:849-937. doi:10.1017/S0140525X01000103
- Huang L, Holcombe AO, Pashler H (2004). Repetition priming in visual search: Episodic retrieval, not feature priming. *Mem Cognit* 32:12-20. doi:10.3758/BF03195816

- Itti L, Koch C (2001). Computational modelling of visual attention. *Nat Rev. Neurosci* 2:194-203. doi:10.1038/35058500
- Jannati A, Gaspar JM, McDonald JJ (2013). Tracking target and distractor processing in fixed-feature visual search: Evidence from human electrophysiology. *J Exp Psychol Hum Percept Perform* 39:1713-1730. doi:10.1037/a0032251
- Kerzel D, Barras C (2016). Distractor rejection in visual search breaks down with more than a single distractor feature. *J Exp Psychol Hum Percept Perform* 42:648-657. doi:10.1037/xhp0000180
- Kiss M, Grubert A, Petersen A, Eimer M (2012). Attentional capture by salient distractors during visual search is determined by temporal task demands. *J Cogn Neurosci* 24:749-759. doi:10.1162/jocn\_a\_00127
- Koechlin E, Basso G, Pietrini P et al (1999). The role of the anterior prefrontal cortex in human cognition. *Nature* 399:148–151. doi:10.1038/20178
- Kong G, Alais D, van der Burg E (2016). An investigation of linear separability in visual search for color suggests a role of recognizability. *J Exp Psychol Hum Percept Perform* 42:1724-1738. doi:10.1037/xhp0000249
- Krummenacher J, Grubert A, Müller HJ (2010). Inter-trial and redundant-signals effects in visual search and discrimination tasks: Separable pre-attentive and post-selective effects. *Vision Res* 50:1382-1395. doi:10.1016/j.visres.2010.04.006
- Krummenacher J, Müller HJ (2012). Dynamic weighting of feature dimensions in visual search: behavioral and psychophysiological evidence. *Front Psychol* 3:221. doi:10.3389/fpsyg.2012.00221
- Krummenacher J, Müller HJ, Heller D (2001). Visual search for dimensionally redundant pop-out targets: Evidence for parallel-coactive processing of dimensions. *Percept Psychophys* 63:901-917. doi:10.3758/BF03194446
- Krummenacher J, Müller HJ, Heller D (2002a). Visual search for dimensionally redundant pop-out targets: Parallel-coactive processing of dimensions is location specific. *J Exp Psychol Hum Percept Perform* 28:1303-1322. doi:10.1037/0096-1523.28.6.1303
- Krummenacher J, Müller HJ, Heller D (2002b). Visual search for dimensionally redundant pop-out targets: Redundancy gains in compound tasks. *Vis Cogn* 9:801-837. doi:10.1080/13506280143000269
- Krummenacher J, Müller HJ, Zehetleitner M, Geyer T (2009). Dimension- and space-based intertrial effects in visual pop-out search: Modulation by task demands for focal-attentional processing. *Psychol Res* 73:186-197. doi:10.1007/s00426-008-0206-y
- Kumada T (2001). Feature-based control of attention: Evidence for two forms of dimension weighting. *Percept Psychophys* 63 :698-708. doi:10.3758/BF03194430
- Lamy D, Zivony A, Yashar A (2011). The role of search difficulty in intertrial feature priming. *Vision Res* 51:2099-2109. doi:10.1016/j.visres.2011.07.010
- Leber AB, Egeth HE (2006). It's under control: Top-down search strategies can override attentional capture. *Psychon Bull Rev* 13:132-138. doi:10.3758/BF03193824
- Li Z (1999). Contextual influences in V1 as a basis for pop out and asymmetry in visual search. *Proc Natl Acad Sci U S A* 96:10530-10535. doi:10.1073/pnas.96.18.10530
- Liesefeld HR, Moran R, Usher M, Müller HJ, Zehetleitner M (2016). Search efficiency as a function of target saliency: The transition from inefficient to efficient search and beyond. *J Exp Psychol Hum Percept Perform* 42:821-836. doi:10.1037/xhp0000156
- Liesefeld HR, Liesefeld AM, Müller HJ (under review). Distractor-interference reduction is dimensionally constrained

- Liesefeld HR, Liesefeld AM, Töllner T, Müller HJ (2017). Attentional capture in visual search: Capture and post-capture dynamics revealed by EEG. *Neuroimage* 156:166-173. doi:10.1016/j.neuroimage.2017.05.016
- Liesefeld HR, Müller HJ (under review). Distractor handling via dimension weighting. *Curr Opin Psychol*.
- Lindsey DT, Brown AM, Reijnen E et al. (2010). Color channels, not color appearance or color categories, guide visual search for desaturated color targets. *Psychol Sci* 21:1208-1214. doi:10.1177/0956797610379861
- Luck SJ, Hillyard SA (1994a). Electrophysiological correlates of feature analysis during visual search. *Psychophysiology* 31:291-308. doi:10.1111/j.1469-8986.1994.tb02218.x
- Luck SJ, Hillyard SA (1994b). Spatial filtering during visual search: Evidence from human electrophysiology. *J Exp Psychol Hum Percept Perform* 20:1000-1014. doi:10.1037/0096-1523.20.5.1000
- Maljkovic V, Nakayama K (1994). Priming of pop-out: I. Role of features. *Mem Cognit* 22:657-672. doi:10.3758/BF03209251
- Martinovic J, Wuerger SM, Hillyard SA, Müller MM, Andersen SK (2018). Neural mechanisms of divided feature-selective attention to color. *Neuroimage* 181:670-682. doi:10.1016/j.neuroimage.2018.07.033
- McDonald JJ, Green JJ, Jannati A, Di Lollo V (2013). On the electrophysiological evidence for the capture of visual attention. *J Exp Psychol Hum Percept Perform* 39:849-860. doi:10.1037/a0030510
- Moran R, Liesefeld HR, Usher M, Müller HJ (2017). An appeal against the item's death sentence: Accounting for diagnostic data patterns with an item-based model of visual search. *Behav Brain Sci* 40:e148. doi:10.1017/S0140525X16000182
- Moran R, Zehetleitner M, Liesefeld HR, Müller HJ, Usher M (2016). Serial vs. parallel models of attention in visual search: accounting for benchmark RT-distributions. *Psychon Bull Rev* 23:1300-1315. doi:10.3758/s13423-015-0978-1
- Moran R, Zehetleitner M, Müller HJ, Usher M (2013). Competitive guided search: Meeting the challenge of benchmark RT distributions. *J Vis* 13(8):24. doi:10.1167/13.8.24
- Mortier K, van Zoest W, Meeter M, Theeuwes J (2010). Word cues affect detection but not localization responses. *Atten Percept Psychophys* 72:65-75. doi:10.3758/APP.72.1.65
- Müller HJ, Heller D, Ziegler J (1995). Visual search for singleton feature targets within and across feature dimensions. *Percept Psychophys* 57:1-17. doi:10.3758/BF03211845
- Müller HJ, Krummenacher J (2006). Locus of dimension weighting: preattentive or postselective?. *Vis Cogn* 14, 490-513. doi:10.1080/13506280500194154
- Müller HJ, Krummenacher J, Heller D (2004). Dimension-specific intertrial facilitation in visual search for pop-out targets: Evidence for a top-down modulable visual short-term memory effect. *Vis Cogn* 11:577-602. doi:10.1080/13506280344000419
- Müller HJ, Liesefeld HR, Moran R, Usher M (2017). Parallel attentive processing and pre-attentive guidance. *Behav Brain Sci* 40:e149. doi:10.1017/S0140525X16000194
- Müller HJ, Reimann B, Krummenacher J (2003). Visual search for singleton feature targets across dimensions: Stimulus- and expectancy-driven effects in dimensional weighting. *J Exp Psychol Hum Percept Perform* 29:1021-1035. doi:10.1037/0096-1523.29.5.1021
- Nothdurft H (1993). Saliency effects across dimensions in visual search. *Vis Res* 33:839-844. doi:10.1016/0042-6989(93)90202-8
- Pollmann S (2004). Anterior prefrontal cortex contributions to attention control. *Exp Psychol* 51:270-278. doi:10.1027/1618-3169.51.4.270

- Pollmann S. (2016). Frontopolar resource allocation in human and nonhuman primates. *Trends Cogn Sci* 20:84-86. doi:10.1016/j.tics.2015.11.006
- Pollmann S, Mahn K, Reimann B et al. (2007). Selective visual dimension weighting deficit after left lateral frontopolar lesions. *J Cogn Neurosci* 19:365-375. doi:10.1162/jocn.2007.19.3.365
- Pollmann S, Weidner R, Müller HJ et al. (2006a). Selective and interactive neural correlates of visual dimension changes and response changes. *Neuroimage* 30, 254-265. doi:10.1016/j.neuroimage.2005.09.013
- Pollmann S, Weidner R, Müller HJ, von Cramon DY (2006b). Neural correlates of visual dimension weighting. *Vis Cogn* 14:877-897. doi:10.1080/13506280500196142
- Pollmann S, Weidner R, Müller HJ, von Cramon DY (2000). A fronto-posterior network involved in visual dimension changes. *J Cogn Neurosci* 12:480-494. doi:10.1162/089892900562156
- Raja Beharelle A, Polanía R, Hare TA, Ruff CC (2015). Transcranial stimulation over frontopolar cortex elucidates the choice attributes and neural mechanisms used to resolve exploration–exploitation trade-offs. *J Neurosci* 35:14544-14556. doi:10.1523/JNEUROSCI.2322-15.2015
- Rangelov D, Müller HJ, Zehetleitner M (2011a). Dimension-specific intertrial priming effects are task-specific: Evidence for multiple weighting systems. *J Exp Psychol Hum Percept Perform* 37:100-114. doi:10.1037/a0020364
- Rangelov D, Müller HJ, Zehetleitner M (2011b). Independent dimension-weighting mechanisms for visual selection and stimulus identification. *J Exp Psychol Hum Percept Perform* 37:1369-1382. doi:10.1037/a0024265
- Rangelov D, Müller HJ, Zehetleitner M (2013). Visual search for feature singletons: Multiple mechanisms produce sequence effects in visual search. *J Vis* 13(3):22. doi:10.1167/13.3.22
- Rangelov D, Müller HJ, Zehetleitner M (2017). Failure to pop out: Feature singletons do not capture attention under low signal-to-noise ratio conditions. *J Exp Psychol Gen* 146:651-671. doi:10.1037/xge0000284
- Sauter M, Liesefeld HR, Zehetleitner M, Müller HJ (2018). Region-based shielding of visual search from salient distractors: Target detection is impaired with same- but not different-dimension distractors. *Atten Percept Psychophys* 80:622-642. doi:10.3758/s13414-017-1477-4
- Sauter M, Liesefeld HR, Müller HJ (in revision). Learning to suppress salient distractors in the target dimension: Region-based inhibition is persistent and transfers to distractors in a non-target dimension.
- Schledde B, Galashan FO, Przybyla M et al. (2017): Task-specific, dimension-based attentional shaping of motion processing in monkey area MT. *J Neurophysiol* 118:1542-1555. doi:10.1152/jn.00183.2017
- Schubö A, Müller HJ (2009). Selecting and ignoring salient objects within and across dimensions in visual search. *Brain Res* 1283:84-101. doi:10.1016/j.brainres.2009.05.077
- Serences JT, Shomstein S, Leber AB et al. (2005). Coordination of voluntary and stimulus-driven attentional control in human cortex. *Psychol Sci* 16:114-122. doi:10.1111/j.0956-7976.2005.00791.x
- Theeuwes J (1991). Cross-dimensional perceptual selectivity. *Percept Psychophys* 50:184-193. doi:10.3758/BF03212219
- Theeuwes J (1992). Perceptual selectivity for color and form. *Percept Psychophys* 51:599-606. doi:10.3758/BF03211656
- Theeuwes J (2018). Visual selection: usually fast and automatic; seldom slow and volitional. *J Cogn* 1:29. doi:10.5334/joc.13
- Theeuwes J, Reimann B, Mortier K (2006). Visual search for featural singletons: No top-down modulation, only bottom-up priming. *Vis Cogn* 14:466-489. doi:10.1080/13506280500195110

- Töllner T, Gramann K, Müller HJ, Eimer M (2009). The anterior N1 component as an index of modality shifting. *J Cogn Neurosci* 21:1653-1669. doi:10.1162/jocn.2009.21108.
- Töllner T, Gramann K, Müller HJ et al. (2008). Electrophysiological markers of visual dimension changes and response changes. *J Exp Psychol Hum Percept Perform* 34:531-542. doi:10.1037/0096-1523.34.3.531
- Töllner T, Rangelov D, Müller HJ (2012a). How the speed of motor-response decisions, but not focal-attentional selection, differs as a function of task set and target prevalence. *Proc Natl Acad Sci U S A* 109:e1990-e1999. doi:10.1073/pnas.1206382109
- Töllner T, Müller HJ, Zehetleitner M (2012b). Top-down dimensional weight set determines the capture of visual attention: Evidence from the PCN component. *Cereb Cortex* 22:1554-1563. doi:10.1093/cercor/bhr231
- Töllner T, Zehetleitner M, Gramann K, Müller HJ (2010). Top-down weighting of visual dimensions: Behavioral and electrophysiological evidence. *Vision Res* 50:1372-1381. doi:10.1016/j.visres.2009.11.009
- Töllner T, Zehetleitner M, Gramann K, Müller HJ (2011). Stimulus saliency modulates pre-attentive processing speed in human visual cortex. *PLoS ONE* 6:e16276. doi:10.1371/journal.pone.0016276
- Treisman A (1988). Features and objects: The Fourteenth Bartlett Memorial Lecture. *Q J Exp Psychol A* 40:201-237. doi:10.1080/02724988843000104
- Treisman A (2006). How the deployment of attention determines what we see. *Vis Cogn* 14:411-443. doi:10.1080/13506280500195250
- Utz S, Humphreys GW, Chechlacz M (2013). Parietal substrates for dimensional effects in visual search: Evidence from lesion-symptom mapping. *Brain* 136:751-760. doi:10.1093/brain/awt003
- Weidner R, Müller HJ (2013). Dimensional weighting in cross-dimensional singleton conjunction search. *J Vis* 13(3):25. doi:10.1167/13.3.25
- Weidner R, Krummenacher J, Reimann B et al. (2009). Sources of top-down control in visual search. *J Cogn Neurosci* 21:2100-2113. doi:10.1162/jocn.2008.21173
- Weidner R, Pollmann S, Müller HJ, von Cramon DY (2002). Top-down controlled visual dimension weighting: An event-related fMRI study. *Cereb Cortex* 12:318-328. doi:10.1093/cercor/12.3.318
- Wolfe JM (1998). What can 1 million trials tell us about visual search?. *Psychol Sci* 9:33-39. doi:10.1111/1467-9280.00006
- Wolfe JM (2007). Guided search 4.0: Current progress with a model of visual search. In: Gray W (ed) *Integrated models of cognitive systems*. Oxford: New York, p. 99-119
- Wolfe JM, Butcher SJ, Lee C, Hyle M (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. doi:10.1037/0096-1523.29.2.483
- Wolfe JM, Horowitz TS (2004). What attributes guide the deployment of visual attention and how do they do it?. *Nat Rev Neurosci* 5:495-501. doi:10.1038/nrn1411
- Wolfe JM, Horowitz TS (2017). Five factors that guide attention in visual search. *Nat Hum Behav* 1:58. doi:10.1038/s41562-017-0058
- Wolfe JM, Palmer EM, Horowitz TS (2010). Reaction time distributions constrain models of visual search. *Vision Res* 50:1304-1311. doi:10.1016/j.visres.2009.11.002
- Zehetleitner M, Krummenacher J, Geyer T et al. (2011). Dimension intertrial and cueing effects in localization: Support for pre-attentively weighted one-route models of saliency. *Atten Percept Psychophys* 73:349-363. doi:10.3758/s13414-010-0035-0

Zehetleitner M, Rangelov D, Müller HJ (2012). Partial repetition costs persist in nonsearch compound tasks: Evidence for multiple-weighting-systems hypothesis. *Atten Percept Psychophys* 74:879-890. doi:10.3758/s13414-012-0287-y