



Ten simple rules to study distractor suppression

Malte Wöstmann^{a,b,*}, Viola S. Störmer^{c,**}, Jonas Obleser^{a,b}, Douglas A. Addleman^{c,1}, Søren K. Andersen^{d,e,1}, Nicholas Gaspelin^{f,1}, Joy J. Geng^{g,1}, Steven J. Luck^{g,1}, MaryAnn P. Noonan^{h,1}, Heleen A. Slagter^{i,j,1}, Jan Theeuwes^{i,j,1}

^a Department of Psychology, University of Lübeck, Lübeck, Germany

^b Center of Brain, Behavior and Metabolism (CBBM), University of Lübeck, Lübeck, Germany

^c Department of Psychological and Brain Sciences, Dartmouth College, USA

^d School of Psychology, University of Aberdeen, UK

^e Department of Psychology, University of Southern Denmark, Denmark

^f Department of Psychology and Department of Integrative Neuroscience, Binghamton University, State University of New York, USA

^g Center for Mind and Brain and Department of Psychology, University of California, Davis, USA

^h Department of Experimental Psychology, University of Oxford, UK

ⁱ Department of Experimental and Applied Psychology, Vrije Universiteit Amsterdam, Amsterdam, the Netherlands

^j Institute for Brain and Behavior, Vrije Universiteit Amsterdam, Amsterdam, the Netherlands

ARTICLE INFO

Keywords:

attention
distractor
suppression
inhibition
target
enhancement

ABSTRACT

Distractor suppression refers to the ability to filter out distracting and task-irrelevant information. Distractor suppression is essential for survival and considered a key aspect of selective attention. Despite the recent and rapidly evolving literature on distractor suppression, we still know little about how the brain suppresses distracting information. What limits progress is that we lack mutually agreed upon principles of how to study the neural basis of distractor suppression and its manifestation in behavior. Here, we offer ten simple rules that we believe are fundamental when investigating distractor suppression. We provide guidelines on how to design conclusive experiments on distractor suppression (Rules 1–3), discuss different types of distractor suppression that need to be distinguished (Rules 4–6), and provide an overview of models of distractor suppression and considerations of how to evaluate distractor suppression statistically (Rules 7–10). Together, these rules provide a concise and comprehensive synopsis of promising advances in the field of distractor suppression. Following these rules will propel research on distractor suppression in important ways, not only by highlighting prominent issues to both new and more advanced researchers in the field, but also by facilitating communication between sub-disciplines.

1. Introduction

Selective attention operates by enhancing task-relevant sensory input and suppressing irrelevant and distracting information. Selective attention research has long focused on understanding the psychological and neural processes of target enhancement. Less well understood are two factors: what makes task-irrelevant stimuli more or less distracting? And, what are the mechanisms by which the neural system suppresses such distraction?

In recent years, the study of whether and how irrelevant and

potentially distracting information is suppressed has attracted renewed interest (for recent reviews, see e.g., Chelazzi et al., 2019; de Vries et al., 2020; Geng, 2014; Luck et al., 2021; Noonan et al., 2018; Schneider et al., 2021; van Moorselaar and Slagter, 2020). In particular, researchers (a) have tried to test whether distractor suppression is an independent process or tied to the processes that produce enhancement of relevant information, (b) have classified distinct types of suppression and have organized them into taxonomies of distractor suppression, and (c) have started to test models of distractor suppression that might scale up beyond laboratory experiments to realistic and more ecologically

* Correspondence to: Department of Psychology, University of Lübeck, Maria-Goeppert Str. 9a, Lübeck 23562, Germany.

** Correspondence to: Department of Psychological and Brain Sciences, Dartmouth College, Moore Hall, 3 Maynard St Hanover, NH 03755, USA.

E-mail addresses: malte.woestmann@uni-luebeck.de (M. Wöstmann), viola.s.stoermer@dartmouth.edu (V.S. Störmer).

¹ Authors contributed equally and are listed in alphabetical order

valid scenarios.

A rapidly growing body of empirical findings on distractor suppression has recently contributed significantly to our understanding of how the brain processes distracting information across different tasks and environments. However, the lack of a comprehensive framework for studying distractor suppression – and of a definition that explains what distractor suppression even is – has limited the potential to integrate these findings into a coherent model. For instance, research on distractor suppression uses partly conflicting or even undefined concepts of what suppression entails and makes use of different operationalizations of distraction. Furthermore, empirical studies on distractor suppression hypothesize different types of suppression without making it clear how these could be falsified and differentiated in theory. Also, studies intermingle the behavioral observation of suppressing distracting stimuli and suppressing brain activity, and often lack causal evidence for the precise roles of neuro-cognitive mechanisms for distractor suppression.

We believe that it is time to provide a set of guidelines for the study of distractor suppression. Here, we convene a diverse group of authors who represent different sub-disciplines of research on distractor suppression. By agreeing on a set of ten simple rules, we hope to demonstrate their importance and general validity, and we recommend that future investigations on distractor suppression use them to guide their research. The set of rules has important implications for study design, execution, and interpretation of results. Some of the rules propose and discuss concrete approaches that can immediately be implemented in future investigations. Other rules identify important considerations that are necessary to integrate results with existing theories of attention. The authors themselves all have experienced some of the drawbacks associated with violating the present rules. In this sense, the rules should not be perceived as top-down prescriptions formulated by an authority but rather as a means to provide researchers a ‘kick start’ to the study of distractor suppression and to avoid common problems.

Furthermore, the present article provides a comprehensive and concise review of recent approaches to study distractor suppression. In our view, considering these rules will significantly advance our understanding of distractor suppression, improve communication of research on distractor suppression across sub-disciplines, and thus foster the establishment of models of selective attention more generally.

Of note, although in our view the proposed rules are widely applicable to the study of distractor suppression in any sub-discipline, we will cover some aspects in more depth than others. First, we focus on behavioral and neural correlates of distractor suppression in humans but all rules should prove useful to studying distractor suppression in other species (rodents and primates; see Halassa and Kastner, 2017). Second, unimodal visual or auditory distraction is our main concern here, but the same principles apply across modalities; e.g., in somatosensation; see e.g., Frings et al. (2008). Last, we do not focus on distractors of particular emotional valence, and readers interested in suppression of emotional distractors are referred to pertinent reviews in the literature; e.g., Carretié (2014); Schindler and Bublatzky (2020).

The present article is organized in three sections. In the first section, Rules 1–3 will cover necessary aspects of experiment design. We review feasible and established approaches to avoid what we consider typical shortcomings. In the second section, Rules 4–6 will review and demarcate different types of distractor suppression. We will emphasize that adequate study design and interpretation of results critically depend on the particular type of distractor suppression under investigation. In the third section, Rules 7–10 will lay out necessary components for models of distractor suppression, elucidate statistical considerations when studying distractor suppression, and discuss how in-lab studies translate to distractor suppression in ecologically valid, real life scenarios.

2. Part 1: Essential components for experiment design (Rules 1–3)

2.1. Rule 1: Make sure the distractor has the potency to distract

Experiments designed to study distractor suppression need to implement some type of distraction in the first place. Manipulating a stimulus that the experimenter presumes to be distracting (i.e., an intended distractor) and observing changes in behavior is a potent approach to understand the underlying processing of distraction in the brain. Critically, not every object has the potency to distract. There is a longstanding and ongoing debate about whether physically salient stimuli capture attention automatically or whether capture depends on the observer’s task set and/or selection history; this ‘attentional capture debate’ has been reviewed elsewhere (e.g., Luck et al., 2021) and will not be reiterated here.

In many circumstances, it is possible to direct attention to a target without any interference from the intended distractor. This is sometimes considered evidence that the object was successfully suppressed, but it is possible that the so-called ‘distracting’ object never had the potential to cause any distraction to begin with. At the same time, if the intended distractor needs to interfere with target processing to demonstrate that it is truly distracting, successful and complete suppression – that is, nullification of interference with the target – cannot be shown empirically.

A popular but restricted approach often used in visual search paradigms (in particular those using singleton distractors) is to omit the intended distractor on some trials and to compare task performance for distractor-present versus distractor-absent trials. The performance decrement for distractor-present trials is thought to quantify the distraction effect (Fig. 1A). If using this approach, it is important to ensure that the number of trials in both conditions is approximately the same. Otherwise, performance may be biased, for instance when performance is superior for frequent distractor-present trials compared with rare distractor-absent trials (for systematic variation of the proportion of distractor-present trials, see e.g., Geyer et al., 2008). A related approach is to vary the strength of the distractor across trials (e.g., Bonnefond and Jensen, 2012), which is another valid approach if differences in performance are observed.

However, one concern arises in cases when performance does not differ for distractor-present versus distractor-absent trials (or highly-distracting vs. less-distracting trials): This does not necessarily imply that the intended distractor is not distracting; instead, it could also be the case that it was successfully suppressed, or that participants learned through experience that the distractor is irrelevant to the task at hand (e.g., through statistical learning; see Rule 4). Thus, it is essential to test whether the intended distractor has the potency to distract independently of the main task.

In the visual modality, task-irrelevant objects with basic features that stand out from the background are likely candidates of distractors (for a list of such features, see Wolfe, 2018). To precisely quantify an object’s saliency, computational models of saliency can also be useful (Itti and Koch, 2001; Koch and Ullman, 1987; Stilwell et al., 2022), though of course this source of evidence depends on the veridicality of the model (Kotseruba et al., 2021). Alternatively, and applicable across sensory modalities, one could assess the potency of distraction using a psychophysical approach. For example, in an independent task, the saliency of a stimulus could be tested by asking participants to report the presence or location of the intended distractor (Fig. 1B). If reports are of high accuracy even when the stimulus presentation is brief (and followed by a mask), then the stimulus is likely sufficiently potent to be distracting in a main task.

Using this approach, it is important that the distractor stimulus is presented in the same surrounding as in the main task because the same stimulus might be distracting in one situation but not another. For example, a red stimulus will likely be more distracting when surrounded by green versus orange stimuli. Similarly, overall display size and local

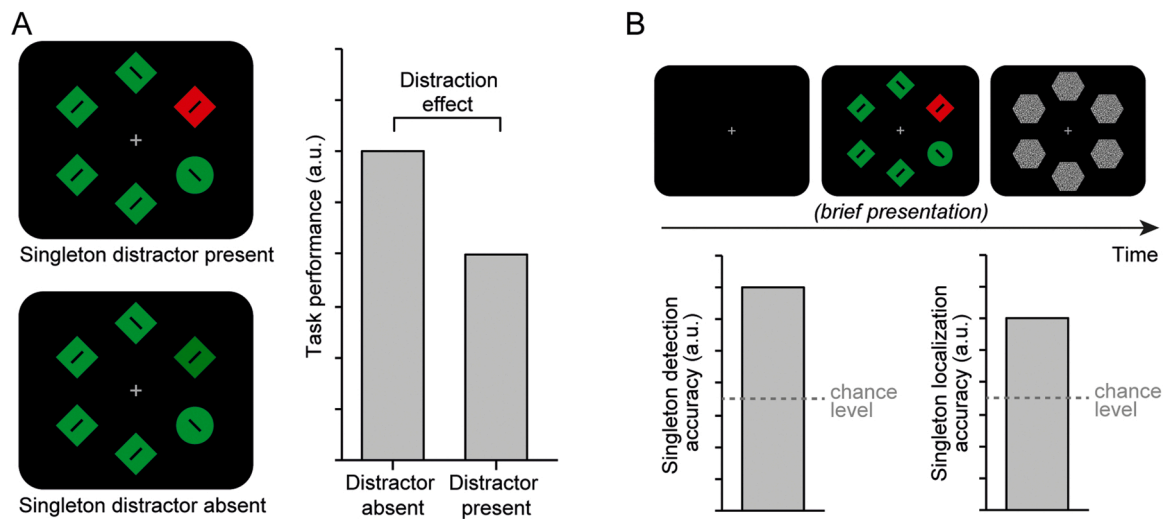


Fig. 1. (A) Schematic example of an approach to test whether an intended distractor has the potency to distract. Left: Participants have the task to report whether the line in the circle shape is left- or right-tilted. The upper search display contains an intended color singleton distractor, which is absent in the lower search display. Right: Enhanced task performance (operationalized by e.g., higher accuracy or faster responses) in distractor-absent trials would give empirical evidence that the intended distractor has the potency to distract. (B) Alternatively, evidence to support the potency of an object to be distracting can be obtained in an independent task, for instance if singleton detection and localization accuracies are high and clearly above chance (see text for details, Figure adapted from Gaspelin and Luck, 2018b, 2019).

feature contrast of the items might matter, as suggested by a recent study of (Wang and Theeuwes, 2020): The authors used visual displays consisting of different numbers of objects and found that the degree of distraction (and potential distractor suppression) depended on the search arrays set size (but see Stilwell and Gaspelin, 2021; Stilwell et al., 2022).

Other approaches can as well provide support that an intended distractor is distracting. One line of work uses distractor cueing paradigms in which participants are cued which visual feature (or location) to ignore on a subsequent visual search trial that contains targets and non-salient distractors (Cunningham and Egeth, 2016; Moher and Egeth, 2012; Reeder et al., 2018; Woodman and Luck, 2007). In these studies, differences in performance for a validly cued (i.e., distractor appears in the cued feature) vs. invalidly cued distractor (i.e., distractor appears in a feature other than the cue) can suggest that the intended distractor is indeed distracting (Addelman and Störmer, 2022).

Beyond the scope of this Rule, it is worth mentioning that any stimulus presented alongside a target might be a distractor, even if it is not cued or not considered salient. Also, non-sensory events, such as internal task-irrelevant thoughts or associations, might be considered sources of distraction.

Any study on distractor suppression, especially those using new and unvalidated paradigms, should provide a test of whether intended distractors are distracting. If we do not comply with Rule 1, we might end up trying to study the phenomenon of distractor suppression without a proper implementation of distraction, which is at best inconclusive and at worst misleading. Whenever we do comply with Rule 1, however, we gain justification to refer to the critical stimulus as a ‘distractor’.

2.2. Rule 2: Manipulate the distractor independently of the target

In addition to the intended distractor being distracting (Rule 1), it is important to ensure that any experimental manipulation of the distractor is independent, i.e., statistically orthogonal, to the manipulation of the target. Indeed, it is a highly debated question whether distractor suppression exists independently of target enhancement or only arises as a consequence of target selection (for example, through center-surround selection, e.g., Andersen and Müller, 2010; Hopf et al., 2006; Störmer and Alvarez, 2014; see also Rule 7). This is a tricky issue, as target and distractor processing are almost always conflated to some extent: every

object that is not a target could potentially be a distractor, and any object that is not a distractor could be a target. How, then, can we isolate distractor processing from processing of the target? The key is to design experiments that manipulate targets and distractors independently of each other.

Many attention studies, however, have not done this. In general, spatial attention studies often only use two locations and present targets and distractors on opposite sides (left and right; Kelly et al., 2006; Wöstmann et al., 2016). For example, in the classic dichotic listening paradigm, two auditory streams are presented simultaneously, one to the left and one to the right ear, and participants are asked to attend to one and ignore the other (Cherry, 1953). These types of designs make it challenging, if not impossible, to link behavioral or neural effects unambiguously to either the enhancement of targets or the suppression of distractors, because these measures can reflect the result of both. For example, lateralized changes in brain activity (e.g., increased alpha activity over the brain region representing a distractor location relative to the target location) may reflect the suppression of distraction at one location, or the enhancement of targets at the other location, or both (Foster and Awh, 2019).

We here review some studies that have successfully separated target and distractor processing. For example, Serences et al. (2004) varied the probability of distractors that co-occurred with targets and found increased neural activity in visual cortex regions (measured with functional Magnetic Resonance Imaging, fMRI) prior to stimulus onset when the probability of distractors was higher, suggesting that this preparatory activity was related to distractor suppression. Another study (Noonan et al., 2016) manipulated the presence of distractors and cued participants to target or distractor locations in advance, which resulted in speeded response times in both cases, though in case of the distractor cues only when spatial locations of distractors were fixed across blocks (possibly suggesting that learned distractor suppression played a role here; see Rule 4). These behavioral effects were accompanied by oscillatory changes in the alpha band (i.e., 10 Hz activity) in the electroencephalogram (EEG) that were related to the cued target location but not the distractor location, suggesting that alpha activity was associated with target enhancement only. In contrast to this, another recent EEG study (Wöstmann et al., 2019) dissociated target and distractor processing in an auditory spatial attention task by fixing either target or distractor position in the front while varying the respective other

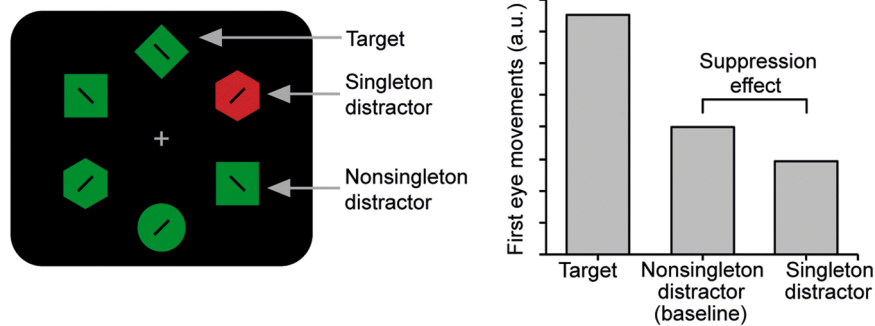


Fig. 2. The additional singleton paradigm and hypothetical results demonstrating oculomotor suppression of a color singleton distractor. (Adapted from Gaspelin and Luck, 2019).

stimulus location to be presented peripherally (left, right). Results showed that lateralized alpha activity tracks both target and distractor locations independently.

A similar logic of fixing the target location centrally while presenting distractors on the left or right side has been used in several other event-related potential (ERP) studies that have examined the P_D component – a lateralized ERP component that has been associated with the suppression of distractors (Gaspar and McDonald, 2014; Hickey et al., 2009; see also Rule 3). Together, these studies exemplify how behavioral and neural correlates of distractor and target processing can be separated by independently and systematically varying features of the distractor (e.g., probabilities or locations), while keeping target processing constant.

Furthermore, it is important to ensure that the distractor is not considered as task-relevant by the observer. This may seem obvious, but a distractor may become indirectly (and unintentionally) task-relevant when it shares certain properties with the target or when it occupies large regions of a visual search display. This task-relevance may prevent the distractor from being suppressed. For example, visual search studies have suggested that when a shape-defined target appears amongst nontarget items that are homogeneous in shape (as in Fig. 1A; Theeuwes, 1992), a color singleton distractor may slow target processing because this type of search encourages participants to search for unique items more generally (singleton detection mode; Bacon and Egeth, 1994; Pashler, 1988). Thus, it is important to consider the attentional set of the observer given a particular task design. Of note, task-relevance of distractors needs to be separated from general psychological relevance: For instance, classic distraction studies operationalized psychological relevance of distractors by presentation of the participant's own name (e.g., Bundesen et al., 1997; Harris and Pashler, 2004; Moray, 1959).

Future studies aiming to investigate the mechanisms of distractor suppression independently of target enhancement should similarly find ways to manipulate targets and distractors separately, and minimize the possibility that distractors are taken as task-relevant by participants. This will allow a clear attribution of behavioral and neural effects to either enhancement or suppression, providing an opportunity to advance theories of attentional selection, in particular whether, when, and how the mind suppresses distracting information.

2.3. Rule 3: Test whether distractors are suppressed below a pre-defined baseline

As discussed in Rule 1, it is essential to first demonstrate that the intended distractor has the potency to distract. Once this is established, one can ask whether the processing of the distractor can be suppressed. In our view, to demonstrate true suppression it is necessary to find a well-reasoned baseline in neural activity or behavior, and processing of the distractor needs to fall below this baseline. In principle, this baseline can be a neutral stimulus or condition (more below), or a neutral time period (i.e., processing before an attention-directing cue). Although this advice may seem relatively straightforward, many studies of

'suppression' have failed to demonstrate that cognitive processing of the distractor is actually suppressed relative to a neutral reference level.

One example of such a baseline may be a search display without a singleton distractor (similar as in Fig. 1A). Evidence for distractor suppression would be obtained if neural populations tuned to features of the distractor show reduced activity in the distractor-present compared to the baseline condition. Other study designs could be used to obtain the baseline and distractor-related responses within the same trial. One such approach has come from studies of saccadic eye movements in the *additional singleton paradigm* (Gaspelin et al., 2017; Gaspelin and Luck, 2018a). Participants searched for a target shape (e.g., a diamond) amongst distractors and reported the orientation of a line inside. One of the distractors was a color singleton that participants were told to ignore. The key result was that eye movements were less likely to be directed to the singleton distractor than the average non-singleton distractor. Thus, a baseline probability of attending a nonsalient object was established and the likelihood of attending a salient distractor was compared to this baseline (Fig. 2). Because the salient distractor attracted overt attention at below-baseline levels, it was assumed to be suppressed. Analogous approaches have been developed to assess whether covert attention to salient stimuli is also suppressed below baseline levels (Gaspelin et al., 2015; see also Stilwell and Gaspelin, 2021).

We acknowledge that some of our own research using event-related potentials has failed to meet this standard. For example, we have found that salient distractors elicit a *distractor positivity* (P_D) component (Sawaki and Luck, 2010), which is a putative index of suppression (Hickey et al., 2009). However, this approach does not yield a separate measure of processing for each item in the array, so it does not provide direct evidence that the salient distractor is suppressed below a neutral baseline. Instead, this approach relies on prior evidence that the P_D co-occurs with other indices that directly demonstrate suppression below some baseline level (e.g., Feldmann-Wüstefeld et al., 2020; Gaspelin and Luck, 2018c). Other evidence has shown that the P_D component can be reduced as a function of statistical learning, which is consistent with the distractor being neutralized and the need for suppression being reduced (van Moorselaar and Slagter, 2019).

It is worth mentioning that establishing a *truly neutral* baseline can be challenging. If the baseline items match some of the features of the target, they might attract attention due to top-down guidance toward the target feature (Folk et al., 1992). This, in turn, may cause the ignored item to appear to be suppressed, even though in actuality the baseline items were boosted (i.e., upweighted). Upweighting and suppression can be challenging to distinguish, in particular when one considers that (a) participants may upweight a large swath of feature or location space, including features that do not exactly match the target feature, for example to increase the distinctiveness between target and distractor (Becker et al., 2010; Geng and Witkowski, 2019; Navalpakkam and Itti, 2007), and (b) the two processes may not be mutually exclusive (Chang and Egeth, 2019). Separating suppression from upweighting, or

enhancement, is an important issue for both studies of featural and spatial accounts of suppression (see also [Rule 2](#)).

In conclusion, establishing a neutral baseline and demonstrating that processing of the distractor falls below this baseline is crucial. In some cases, establishing a baseline may be related to establishing that a distractor is truly distracting ([Rule 1](#)); in other cases, there are approaches that only address [Rule 1](#) (e.g., when using a model to test how salient a potential distractor is; [Stilwell and Gaspelin, 2021](#)) and others that only address [Rule 3](#) (e.g., when measuring neural processing during a pre-trial baseline period; [Andersen and Müller, 2010](#)). Critically, without a baseline, it is unclear whether the processing of the distractor was truly suppressed.

3. Part 2: Distinct types of distractor suppression (Rules 4–6)

3.1. Rule 4: Consider intentions versus experiences as sources of suppression

It is important to consider which mechanism drives a specific case of attentional suppression. On the one hand, suppression could be driven solely by the current intention of the observer (e.g., arising from explicit instructions). On the other hand, the brain's extraction of statistical regularities over time might enable learning to produce suppression of certain features, such as the locations that are likely to contain distractors ([Wang and Theeuwes, 2018a, 2018b, 2018c](#)). These two types of suppression are at the ends of a continuum, and researchers aiming to identify mechanisms of distractor suppression should be mindful of this when designing their experiments.

Experience-driven distractor suppression arises when certain distractors that are repeatedly encountered tend to be ignored more effectively over time. This type of learned suppression is not a unitary phenomenon, instead arising from a variety of sources along two dimensions: temporal extent and feature specificity ([Fig. 3](#)). Experience can transiently bias selective attention for a few trials (as in inter-trial distractor priming; [Kristjánsson and Driver, 2008](#)) and over longer time periods via statistical probability learning ([Ferrante et al., 2018; Goschy et al., 2014; Vatterott and Vecera, 2012](#)). Furthermore, experience can induce feature-general inhibition of singleton distractors within a dimension (e.g., any odd-color singleton among uniformly colored items; [Won et al., 2019](#)), relatively broad suppression of a range of previously encountered distractor features ([Chetverikov et al., 2016](#)), and feature-specific 'distractor templates' (e.g., avoiding distraction by red due to recent rejection of red distractors; [Gaspelin et al., 2015; Geng et al., 2017](#)). A specific instance of experience-driven distractor suppression can exist anywhere in this two-dimensional space, with evidence suggesting that different sources of experience-driven attention recruit dissociable mechanisms ([Addleman and Jiang, 2019](#)).

In most cases, investigations ought to isolate a single form of distractor suppression at a time. Here, we discuss three cases that demonstrate important considerations when the goal is to isolate suppression based on current intentions (case 1) or to separate different forms of experience-driven suppression (cases 2 & 3):

1. It is well known that attention can be directed at will (i.e., top-down) to a location in space ([Posner, 1980](#)). In these types of studies, prior to each search array, a central symbolic cue indicates the likely location of the upcoming target. The typical finding is that observers are faster and more accurate when the target appears at this cued location than at non-cued locations. It is critical that on each trial, observers are required to actively direct attention to a new location or feature to avoid inter-trial priming (lingering biases). With this consideration in mind, one can ask under which circumstances the intentional suppression of distraction is possible. For example, several studies used explicit cues that indicated to participants which features to ignore, varying that feature from trial to trial. These studies found that these cues resulted in behavioral benefits during visual search, suggesting that intention-based suppression can be effective in terms of speeding responses ([Addleman](#)

and [Störmer, 2022; Arita et al., 2012; Carlisle and Nitka, 2019; Wen et al., 2018](#)). In contrast to these studies, another recent study used the *additional singleton paradigm* and cued participants on each trial to suppress one of the locations that was most likely to contain a distractor ([Wang and Theeuwes, 2018b](#)). The results showed no evidence for suppression: capture by the salient distractor was just as strong as when no cue was provided. However, using identical displays, capture was strongly attenuated when distractor locations repeated and participants had instead learned that this location was most likely to contain a distractor. These and similar results ([Cunningham and Egeth, 2016](#)) provide direct evidence that learned distractor ignoring can arise from repeated cues even when intentional ignoring is ineffective when cues vary from trial to trial, further demonstrating the importance of considering trial history in paradigms whose goal is to investigate intention-based suppressive effects only.

2. When inducing learned long-term distractor suppression by making distractors more likely to appear, say, at one location or in one color than others, doing so also makes immediate repetitions of that feature more likely (inter-trial priming). This can yield larger effects than long-term or short-term learning alone ([Sha et al., 2017](#)). Thus, it is important to isolate long-term learning effects from inter-trial priming. This can often be done using a post-learning test phase with equal distractor feature probabilities, allowing researchers to evaluate the potentially durable long-term effects of manipulating distractor probability with inter-trial repetitions equally likely for all features.

3. When studying learned suppression of singletons of a specific feature value (e.g., red), one question is whether participants are learning to suppress singletons more generally (e.g., of any unique color). This can be accomplished by comparing performance for the feature-specific learned distractor both against a neutral baseline (i.e., without a singleton distractor; see [Rule 3](#)) and against a condition in which singletons are in both learned and unlearned colors. This allows for comparisons between performance on singleton-absent trials and trials with either trained or untrained singleton colors.

A benefit of the breadth of recent research on distractor suppression has been the identification of different sources of suppression arising from explicit instructions or incidentally from task experience. As researchers attempt to understand the similarities, differences, and relationships between these mechanisms, it is critical to design experiments to isolate specific sources of distractor suppression. Not doing so risks misattributing contributions of cognitive and neural mechanisms to different kinds of suppression, which impedes the design of comprehensive models for distractor suppression (see [Rule 7](#)).

3.2. Rule 5: Distinguish between proactive and reactive suppression

Distractor suppression is often assumed to be a single neuro-cognitive process that reduces behavioral costs associated with distractor presence. However, accumulating evidence suggests that the reduction of distractor interference may rely on multiple mechanisms ([Chelazzi et al., 2019; Geng et al., 2019; Noonan et al., 2018](#), see [Rules 4 & 7](#)). At a coarse level, these mechanisms can be defined by those that operate *proactively* to attenuate processing before the distractor stimulus appears (e.g., in a search display), and those that operate *reactively* to terminate distractor processing after the stimulus has already initially captured attention ([Geng, 2014](#)). Furthermore, proactive mechanisms that result in suppression prior to distractor onset can be distinguished by whether there is any initial enhancement of the to-be-suppressed distractor location or feature before suppression is implemented. If one wishes to distinguish between these, several considerations should be kept in mind.

First, be clear about how proactive and reactive are defined. Several approaches have been used and it should be noted that there is no current consensus on what mechanism, or mechanisms, defines proactive and reactive suppression. In the cognitive visual search literature, behavioral measures of proactive suppression are frequently defined by

the likelihood of saccadic eye-movements being directed away from the distractor compared to a neutral non-target or the same distractor in a different context (similar to Fig. 2); reactive rejection is indexed by very short fixation durations or saccade paths (Born et al., 2011; Gaspelin et al., 2017; Geng and Diquattro, 2010). In the ERP literature, proactive vs. reactive suppression are indexed by the timing of the P_D component and its relationship to the onset or absence of an N2pc (Hickey et al., 2009). EEG and fMRI studies measure preparatory activity, suppressed distractor-evoked sensory activity, or frontoparietal attentional consequences of reactive mechanisms (Adam and Serences, 2021; Marini et al., 2016; Noonan et al., 2016; van Moorselaar and Slagter, 2020; Won et al., 2020). The vast majority of these methods estimate proactive or reactive processes by behavioral or brain responses to distractors within visual search displays, but some studies examine proactive suppression before the search display using ‘probe trials’ or direct measurements of preparatory activity (Addleman and Störmer, 2022; Huang et al., 2021; van Moorselaar and Slagter, 2020). Given the variety of methods for measuring proactive vs. reactive suppression, it is critical to carefully distinguish between effects we observe and latent mechanisms (Anderson, 2021; Lamy, 2021; see also Rules 4,6,7).

Second, interpret proactive and reactive suppression within the experimental context. It has become clear that certain paradigms elicit stronger effects of proactive suppression than others. Distractors that are predictable are better suppressed than those that are more unexpected or even surprising (Awh et al., 2003; Moher et al., 2011; Noonan et al., 2016; Vatterott et al., 2018; Won et al., 2019). Moreover, proactive suppression occurs more easily when the expectation is derived from experienced statistics than a trial-specific explicit rule in working memory (see Rule 4); rapid reactive mechanisms may be more effective in the latter. This pattern of distractor suppression in which proactive mechanisms operate most effectively on implicit statistical learning and reactive mechanisms are faster when applying explicit rules stands in contrast to the original dual mechanisms of control (DMC) framework

(Braver et al., 2007). DMC posited that proactive control depends on sustained working memory maintenance of goal states and is more resource intensive than reactive processes. Thus, mechanisms of proactive and reactive distractor suppression should potentially be considered as distinct from those involved in cognitive control and target selection.

Third, consider intra- and inter-individual variability. While an average reduction in distractor interference between conditions may be interpreted as proactive suppression, this may not be true if reactive processes are rapid and lead to correct outcomes. Individuals may vary in use of proactive versus reactive mechanisms to reject a distractor on a particular trial (Born et al., 2011; Geng and Diquattro, 2010), or individuals may differ in their proportional use of each (Forster and Lavie, 2016; Fukuda and Vogel, 2011). It is also possible that reactive mechanisms may be called upon when intended proactive ones fail within a trial or over time. These differences are masked within averaged responses but expose a fluidity between strategies.

In sum, the proactive–reactive distinction offers a broad categorical division between mechanisms of suppression, and it is important that these are clearly defined in studies of distractor suppression. At worst, failure to consider the full range of mechanisms could produce errors of reverse inference in interpreting underlying mechanisms (e.g., proactive suppression via top-down control) based on behavioral outcomes or effects (e.g., reduced interference). At best, ignoring these distinctions is a missed opportunity to better understand how optimal distractor suppression may involve synergy between a host of mechanisms, some of which operate proactively and others reactively to minimize distraction.

3.3. Rule 6: Do not confuse psychological with neurophysiological suppression

Describing or dissecting any phenomenon at the psychological level does not equal describing it in neurophysiological terms. This should be

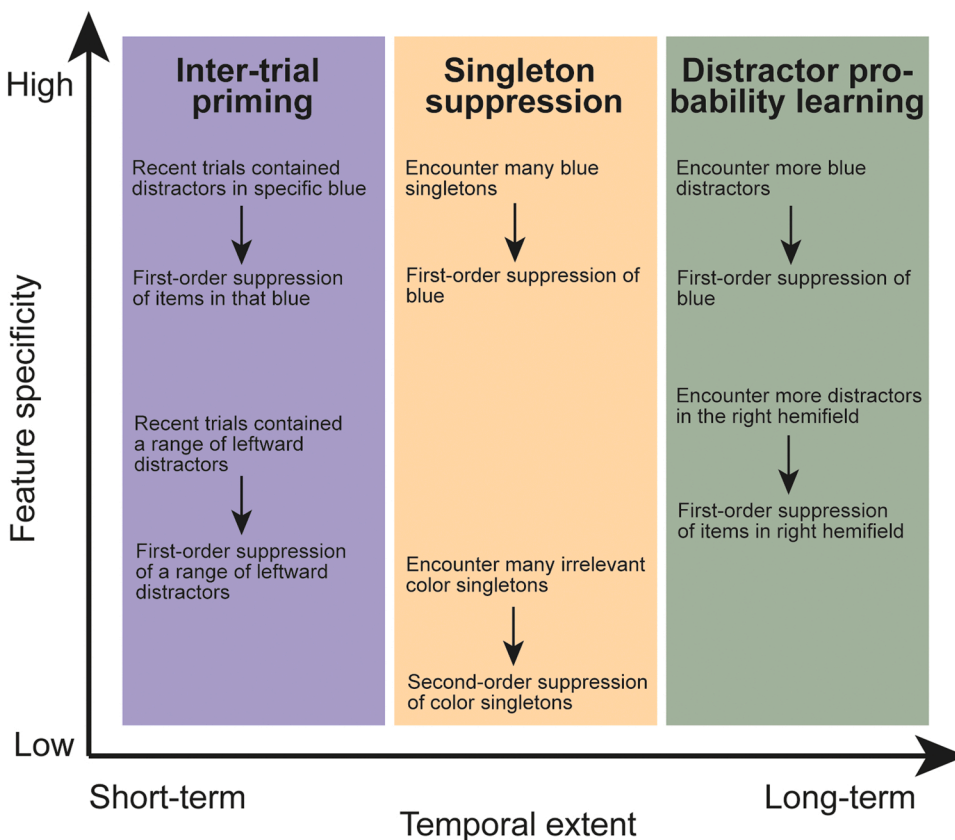


Fig. 3. Experience-driven distractor suppression along the two dimensions feature specificity and temporal extent. Most effects can be categorized as relatively short-term inter-trial priming (Kristjánsson and Driver, 2008), longer-term singleton suppression (Gaspelin et al., 2015), or highly durable effects of distractor probability learning (Vatterott and Vecera, 2012). In all three categories, learning tends to be feature-specific when encountering distractors with a single feature value (e.g., a specific blue), but becomes more general when encountering a wider range of feature values (e.g., Chetverikov et al., 2016). In the case of singleton suppression, learning can generalize beyond encountered feature values when rejecting a wide range of color singletons (Won et al., 2019), while such second-order suppression of distractors is not known to occur for inter-trial priming and probability learning. Here we have treated various feature dimensions (color, orientation, location) together, but it is a matter of future research to test how the temporal extent and feature specificity of effects varies across feature dimensions.

a painfully obvious statement to make, and it has been made many times before (Anderson, 2011; Krakauer et al., 2017; Miller, 2010). Why, then, is it so easy to fall for this epistemological trap time and again? Probably because it is such an alluringly reductionist fallacy: As if we had explained anything yet simply by replacing a behavioral phenomenon with a neurobiological one.

With respect to distractor suppression, the term has been interchangeably intertwined when describing the psychological phenomenon of successfully ignoring distracting information and when referring to the neurobiological processes of suppression. However, there is not a unitary ‘neural level’ to start with; rather, the behavioral observation of ignoring distractors can be accompanied by (not to say, caused by) changes at many neural levels, and there might or might not be a single ‘processing stage’ or neuroanatomical locus to which such behavioral change can be tied to (for a review see e.g., Waschke et al., 2021; see also Rule 8). Within the domain of neuroscience, suppressive effects can sometimes be defined more straightforwardly because they may be directly observed, for example as lateral or recurrent inhibition, or as inhibitory neurons; however, this is not always the case, as EEG or fMRI activity, for example, can reflect excitatory or inhibitory processes (or both). Thus, for both behavioral and many neuroscientific measures even the term ‘suppression’ – and especially the different mechanisms underlying suppression (see Rules 4,5,7) – is not specified clearly enough.

It is critical that attention and suppression researchers ask themselves to what degree their metaphorical language will help or hinder a true understanding of the mechanism under consideration. It is worth reminding ourselves that, for example, the filter metaphor so ubiquitous in attention research is per se neither psychological nor biological in origin. Instead, it was ‘imported’ from engineering in the heyday of the mid-20th century cybernetics and information age. Colin Cherry in his 1953 paper truly broke new ground for attention research when he mused: “how do we recognize what one person is saying when others are speaking at the same time (the ‘cocktail party problem’)? On what logical basis could one design a machine (‘filter’) for carrying out such an operation?” (Cherry, 1953, p. 975 f.). Note this sudden switch from a psychological to an engineering perspective that brought us a whole swath of useful new terminology. Not only terminology, but this new perspective certainly pushed much research to investigate the mechanisms of selection under a new light. But with this switch came also a muddying of the conceptual waters, and it is now easily overlooked that neither the mind per se entertains ‘sets of filters’ nor changes in firing rate or local field potential per se pose such filters.

One more radical conclusion might be to replace the term ‘distractor suppression’ altogether and to use instead descriptions that are more clearly operationalized (as has been suggested for the maybe even broader term ‘attention’, Hommel et al., 2019). No need to put the baby out with the bathwater: We believe that ‘distractor suppression’ can still be useful as an explanatory concept (just like the filtering metaphor is). It is critical though that its meaning is clearly specified by the researcher, both in terms of behavioral measures and neural processes. Then, the goal of research on distractor suppression should be to link brain processes to behavioral observations (Rule 8) in order to understand the underlying mechanism. A study by Gaspar and McDonald (2014) provides such an example: the magnitude of the P_D component of the event-related potential – a presumed neurophysiological marker of suppression (see also Rules 2 & 3) – was associated with faster response times, suggesting that successful distractor suppression measured behaviorally was indexed by a preceding neural correlate of suppression. Another study showed that behavioral costs of attending to two perceptually similar target features could be explained by center-surround suppression in feature space assessed in an independent task using steady-state visual-evoked potentials in the EEG (Störmer and Alvarez, 2014). In other cases, neurophysiological effects have been observed to relate to behavior in one but not another condition, highlighting the importance of not equalizing neural and behavioral indexes

of suppression: For example, Noonan et al. (2016) found behavioral benefits for target and distractor cueing during a spatial attention task, but oscillatory alpha activity – a presumed neural correlate of suppression – was only observed following target cues, questioning a direct association of these neural responses with behaviorally observed distractor ignoring. Studies like these are important to understand neural signatures of distractor suppression. In the future, it is necessary to explore how generally or specifically different neural measures relate to the manifestations of psychological phenomena in behavioral observations.

4. Part 3: Testing models of distractor suppression (Rules 7–10)

4.1. Rule 7: Define your model of distractor suppression

Only if a model is defined can it be tested, falsified and contrasted with other approaches to understand behavior. It is therefore essential to define the operational model in the context of the respective study. Having established that researchers should be aware of the conceptual differences of psychological suppression versus neural suppression (Rule 6), we here review three possible models that deliberately posit neural excitatory and inhibitory processes as the basis for proactive distractor suppression (Rule 5; Fig. 4): direct suppression, secondary suppression, and expectation suppression (Noonan et al., 2018; van Moorselaar et al., 2020).

Direct suppression, as the potential flipside to target facilitation, is when specific representations in lower sensory areas are selectively inhibited by distractor-specific working memory representations in top-down attentional control areas (Arita et al., 2012; see also Rule 4). Secondary suppression, by contrast, may be a secondary effect of target facilitation in which top-down control mechanisms up-regulate activity in target-related neural populations and trigger a spread of inhibition to unattended input via local inhibitory circuits (Andersen and Müller, 2010; Hampshire and Sharp, 2015; Störmer and Alvarez, 2014). Finally, expectation suppression is the selective suppression of predictable or repeated inputs governed by principles of predictive coding or repetition suppression (van Moorselaar and Slagter, 2020). Top-down attention would then be necessary to release task-relevant input from expectation suppression (Summerfield et al., 2008), leaving the inhibitory influence focused on distractors, or potentially intersecting with mechanisms that further downweight prediction errors associated with predictable distractor stimuli.

To complicate our job however, distractor suppression may not be a unitary mechanism (Rule 4), and teasing these processes apart is not necessarily simple (Noonan et al., 2018). Although fundamentally different, the mechanisms sometimes make similar predictions at gross brain and behavioral measurements. For example, in some contexts, when distractors are cued, all models predict reaction time improvements. Similarly, all mechanisms may involve some degree of top-down control – driven either explicitly by intentions or by implicit biases, for example through statistical learning. Further, the anatomical location of the ‘top’ may vary between models; potentially long-ranging prefrontal cortex (PFC) signals in the case of secondary suppression but more local signals from the next level in the processing hierarchy for expectation suppression. And finally, while the temporal accumulation of effects may distinguish fast, trial-wise direct suppression mechanisms, behavioral effects mediated by secondary and expectation mechanisms would be expected to both increase with time, particularly in stable and predictable environments.

It is therefore essential that paradigms are created that clearly delineate the possible paths of distractor suppression. Fortunately, there are some key distinctions between the models, specifically in terms of the flexibility, specificity and selectivity of suppression, and the impact of on-going task demands. Direct, intentional suppression, for example, should be flexible, and specific for the exact cued distracting input. By contrast, secondary suppression may be operationally as flexible as




Fig. 4. Three putative neural mechanisms for preparatory suppression illustrated with reference to a spatially defined target/distractor input. (A) The task requires participants to report the orientation of the grating in the orange box (target). Participants have been forewarned that the top right stimulus will be a distractor. (B) Direct suppression: neurons representing the distractor in the left visual cortex are specifically inhibited via top-down connections from fronto-parietal cortex. (C) Secondary suppression: The left visual field is not inhibited directly but mediated via top-down excitation of right visual cortex with biased competition mechanisms at lower levels. (D) Expectation suppression: Predictable visual inputs are suppressed through inhibitory connections from within the visual processing hierarchy. Predictable representations remain suppressed unless rescued via additional top down facilitation. Green and red lines indicate excitatory and inhibitory connections, respectively.

direct suppression but dependent on the initial target facilitation, and potentially relatively non-specific as it would suppress all uncued locations or features. Furthermore, while expectation suppression may be relatively specific to repeating stimuli, it may require more experience to manifest, in particular for physically highly salient stimuli (Slagter and Moorselaar, 2021). Finally, working memory demands may differentially impact the mechanisms, potentially affecting direct suppression and expectation suppression (St John-Saaltink et al., 2015) but leaving secondary suppression mechanisms unaffected.

Each model has a variable degree of independent support and while potentially mutually exclusive they may actually co-exist; being specialized for different contexts (e.g., proactive vs reactive suppression; Rule 5), processing level, stimuli (feature vs space), as well as factors such as target-distractor similarity and distractor salience (van Moorselaar and Slagter, 2020). In fact, the mechanisms may interact and complement each other. While it may be challenging to separate these explanations, it is important to do so in order to understand how they support typical and atypical behavior. It is critical to get this step right as neglecting this element can limit the conclusions we can draw and hinder our ability to design and test new models.

4.2. Rule 8: Unravel causal implications for distractor suppression

Identifying causes and consequences of distractor suppression is essential for understanding their roles in selective attention. Apparently straightforward ways to test causality are brain stimulation (for reviews, see e.g., Herrmann et al., 2016; Jazayeri and Afraz, 2017; Schneider et al., 2021) and neurofeedback techniques (for reviews, see Enriquez-Geppert et al., 2017; Sulzer et al., 2013). Although brain stimulation does allow for experimental manipulations, it is often debatable to which extent these manipulations are valid operationalizations of the underlying constructs. Furthermore, it is not always feasible to use these techniques and much more common are correlational approaches where it is often unclear what kinds of causal inferences can be drawn (if any).

A key challenge when using correlational approaches lies in the fact that many of the variables between which we may seek to establish causal relationships, such as the brain's sensory stimulus processing, behavioral responses, ongoing neural oscillations, etc., typically take the role of dependent variables in experiments. Thus, the method of choice for establishing causal relationships, experimental manipulation of an independent variable while measuring its effect on a dependent variable, may not be viable for many questions. Hence, observed relationships may often be of a correlational nature, for which a multitude of competing causal relationships could potentially account. It can be a daunting task to establish a causal relationship between dependent

variables with any certainty, but by taking some additional steps, it is possible to rule out at least some of the (potentially infinite) number of competing accounts.

Let us consider three possible approaches in a simple thought experiment: imagine that we observe that, compared to a neutral baseline (Rule 3), invalid spatial cues elicit slower reaction times and reduced P1 amplitudes roughly 100 ms after stimulus onset in the distractor-elicited visual ERP. We might suspect that attentional modulation of sensory processing reflected in the P1 amplitude is the cause for slower reaction times, because (a) both are correlated and (b) the P1 precedes the manual response. In fact, neither of these two arguments is necessarily valid. All that we observed is that both dependent variables change with the same manipulation of an independent variable, but we do not know which relationship (if any) exists between these variables.

A meaningful additional analysis might be to examine the covariation between both measures across participants or, even better, trials within the same conditions, to establish whether the relationship between the measures is really as tight as might be expected if there were a direct causal link (e.g. Gaspelin and Luck, 2018b; Steinhauser and Andersen, 2019). What about the timing? The P1 component, with a latency roughly around 100 ms, will precede even the fastest manual responses on any given trial. However, to establish whether attentional modulation of the P1 precedes attentional modulation of reaction time, we should examine the time-course of attention shifts by measuring different cue-stimulus intervals (CSIs). If our putative causal explanation holds, attentional modulation of both measures should be closely linked across different CSIs. Specifically, there should be no CSI at which attentional modulation of reaction time, but not the P1, occurs (for an example of this logic involving the alpha band see Antonov et al., 2020).

Imagine the P1-reaction time relationship in our example had passed the additional tests above. While this would provide some support for a causal link between P1 and behavior, in a potentially multicausal system, we still do not know how important attentional modulation of sensory processing reflected in the P1 is for the observed behavioral effects – in the end, both are still dependent variables. Thus, it still could be one of many causes, or even just an epiphenomenon (e.g. an ‘echo’ of more relevant unobserved effects elsewhere). By testing quantitative models of the relationship between the variables, we can address whether changes in the putative cause are sufficient to account for the observed effects (e.g. Itthipuripat et al., 2014).

In summary, caution is advised when interpreting seemingly tight relationships between variables in experiments on distractor suppression (as well as target enhancement). When the underlying data is of correlational nature, it is of little help to notice that the observations are consistent with a certain causal interpretation: the problem is that there

are potentially too many causal interpretations that are consistent with the data. Explicit efforts should be directed at ruling out at least some of these. As pointed out above, other techniques such as transcranial magnetic stimulation (TMS) may be used to more directly test for causality, though of course these techniques do come with their own limitations (Jazayeri and Afraz, 2017). Most importantly, any claims made should be mindful of those that couldn't be ruled out.

4.3. Rule 9: Beware of what statistical tests do and do not reveal about suppression

Statistical inference is a key element of any empirical study on distractor suppression. Often, it is the (non-) significance of a statistical test that is taken as indication for or against the existence of distractor suppression in patterns of brain activity or in their behavioral manifestation. While common statistical mistakes have been reviewed elsewhere (e.g., Makin and Orban de Xivry, 2019), we will here caution against problematic types of statistical inference observed in the recent literature on distractor suppression.

Studies on selective attention often contrast (a) target enhancement versus distractor suppression or (b) different kinds of distractor suppression (Rules 4–6). In the literature, it is easy to find studies that observed statistically significant target enhancement but not distractor suppression, while other studies found both mechanisms to be statistically significant. The probability to detect an existing effect – that is, statistical power – depends on the (unknown, true) effect size, the sample size (N), as well as on the significance level (α), which is typically fixed at .05. In practice, effects related to target processing are often considerably larger than effects related to distractor suppression. Among others, one likely reason is that participants' task performance most often directly relates to target processing but only indirectly to distractor suppression. Note that it is furthermore likely that different kinds of distractor suppression induce effects of different sizes.

Let us consider an example of how the choice of a too small sample can affect the statistical results, which, in turn, can make researchers draw misleading conclusions. In an auditory spatial attention task (Wöstmann et al., 2019; $N = 33$), the authors found two statistically significant neural responses, one related to attention to lateralized targets (Cohen's d effect size = 1.24), and the other related to ignoring lateralized distractors (Cohen's d effect size = 0.42). Based on these results, the authors concluded that the human brain implements both, target enhancement and distractor suppression. Given these effect size estimates, let us consider possible results of the same study with a smaller yet typical sample size, say $N = 20$. While this study would still detect the target processing effect with high probability (power > 0.99), it would likely miss the much smaller distractor processing effect (power < 0.5). In this case, it would be tempting but erroneous if the researchers concluded that the human brain implements target enhancement but not distractor suppression. These reflections illustrate that a non-significant statistical test does not tell whether the distractor (suppression) effect under consideration is non-existent or simply too small to be detected given the present study design (see also Amrhein et al., 2017). Furthermore, a non-significant statistical test on one particular measure of distractor suppression (e.g., suppression of initial saccades towards the distractor) does not preclude that the distraction effect would manifest at a later point in the stimulus processing cascade (see Rule 8), or in another behavioral or neural measure (e.g., suppression of distractor-evoked neural responses).

To improve interpretability of (non-) significant results of statistical tests in research on distractor suppression, researchers should be aware that effect sizes for target enhancement and distinct types of distractor suppression differ considerably. An experiment that is adequately powered to detect one kind of distractor suppression can still be underpowered when the goal is to detect another type of distractor suppression. A useful statistical approach to differentiate 'evidence for the absence of distractor suppression' and 'absence of evidence for

distractor suppression' is to compute the Bayes Factor (BF; Dienes, 2014; Keyesers et al., 2020; Rouder et al., 2009). As some research has shown, the BF does indeed speak to the absence of some types of distractor suppression effects (e.g., Reeder et al., 2018; Sörqvist et al., 2013).

It is necessary to not only test the existence of target enhancement and distractor suppression effects (i.e., by testing each against zero), but also to directly compare target enhancement with distractor suppression statistically (Nieuwenhuis et al., 2011). Furthermore, consistent reporting of effect sizes is recommended as these allow to adequately power future studies (Funder and Ozer, 2019). Open science practices of accessibility and transparency facilitate the combined analyses of data across studies to enhance power and to improve precision of effect size estimates.

If we do not comply with Rule 9, we risk drawing conclusions that are not supported by the empirical data. If we do comply with Rule 9, we avoid overinterpreting results on distractor suppression and we provide effect size estimates necessary for adequate design of future studies.

4.4. Rule 10: Consider distraction in the lab versus in the real world

Neural and behavioral consequences of distraction have so far been mostly studied in controlled laboratory settings. Yet, there are some key differences between these typical experimental set ups and our natural environments that are important to consider as they can present theoretical challenges. Thus, a long-term goal should be to not only understand how our brain deals with distraction in the laboratory, but also in the real world.

First, in the vast majority of laboratory experiments, participants have to sit still, keep their eyes at fixation, and respond with a button press to stimuli that are presented to them. This is particularly true in EEG and fMRI studies, given their sensitivity to motion artifacts. However, humans are not passive observers, but active agents in their environments. Our brains are not at the whim of external influences, but can exert a powerful influence over their own sensory sampling: To block out unwanted noises we can simply close our office door, and we move our eyes away from a distractor (Abeles and Yuval-Greenberg, 2021) to reduce its sensory impact in the first place.

Indeed, eye tracking studies show that scanning paths can deviate away from known distractors (Van der Stigchel and Theeuwes, 2006). Visual search in the real world also often entails reaching for, or turning our heads towards objects, which automatically means not reaching for, or turning our heads towards another object (i.e., the exclusion of certain sensory inputs). Strikingly, one study found that highly physically salient stimuli were better suppressed during reaching movements in visual search than less salient stimuli (Moher et al., 2015), which is opposite from what is typically found in studies with passive participants and contrary to predictions by influential theories of visual search (Wolfe, 1994). Another study found that while in a typical button-press condition, distractor interference was largely determined by perceptual salience, in search conditions requiring acting upon the target object, distractor interference was determined by the extent to which distractor features were relevant to the to-be-performed action (Pavese and Buxbaum, 2002). Nonetheless, in the lab, action is typically treated as separate from perception or as a confound. The consequence of this is that we know very little about how action selection may affect distractor processing.

Second, distraction is typically studied using paradigms in which participants are presented with relatively simple stimuli. Yet, our natural environments are visually rich, three-dimensional, inherently multisensory, and often dynamic. They also provide predictive structure. Researchers have only just begun to investigate to what extent basic findings from the literature on distraction obtained with simple, artificial stimuli may generalize to more naturalistic stimuli and contexts. This research suggests that while many basic findings may replicate in more ecologically valid conditions (Olk et al., 2018), there may also be qualitative differences compared to more artificial lab settings.

One recent study, for example, reported ultrafast distractor rejection using natural scene stimuli (Hickey et al., 2019), suggesting that global scene information can rapidly constrain naturalistic search. Indeed, our natural environments typically contain predictive structure and regularities that can not only guide attention (Kaiser et al., 2016; Theeuwes, 2021; Vö et al., 2019), but are also informative about what can be safely ignored (Vadillo et al., 2021). Sources of initial distraction (e.g., traffic outside your new office window), may simply disappear into the background once familiar to the brain. In line with this, recent studies show that statistical learning about the likely location or color of a distractor greatly reduces distractor interference, and can even prevent capture by a physically salient distractor (Gaspelin and Luck, 2018c; see also Rule 4). Yet, how natural structure and statistical learning, in particular over longer time periods as is common in daily life, facilitate the ability to ignore distracting information is still largely unclear (van Moorselaar and Slagter, 2020).

Third and finally, in daily life, the most potent sources of distraction are not meaningless red circles flashed on a computer screen, but stimuli (e.g., your smart phone) that in the current task context (e.g., paying attention in class) should be ignored, but are generally valuable sources of information (e.g., messages from your friends). In the lab, value-based attention is typically induced using monetary rewards, and studies show enhanced capture by distractors previously associated with monetary reward even when physically inconspicuous (B. A. Anderson, 2016; Failing and Theeuwes, 2018). Yet, in daily life, reward is much more multifaceted and attention and distraction may be better explained by the main imperative of the brain to sample the most informative signals to achieve goal-directed action over the long-term (Friston, 2009; Hommel et al., 2019; Slagter and Moorselaar, 2021). There is indeed some empirical evidence showing that objects that one previously acted on can drive subsequent visual search (Buttaccio and Hahn, 2011). Naturalistic driving studies furthermore reveal the powerful influence of irrelevant objects that invite information sampling on attention (Dingus et al., 2016). The goal of future work should be to better understand how natural attentional habits are established and interfere with goal-directed performance.

Thus, consider and determine the ecological validity of your

paradigms and findings, and how lab-based theories may extent to distraction in the real world. After all, in the long run, we wish to capture and explain natural human cognition and performance.

5. Conclusion

Science benefits from both specialization and interdependence (Utzerath and Fernández, 2017). Here, we present a consensus on ten simple rules to study distractor suppression, agreed upon by experienced researchers from diverse sub-disciplines of attention research. Just like for most sets of rules, in practice it might turn out necessary to prioritize some of the rules at the expense of others. In particular, we recommend that Rules 1–3 should generally be complied with. However, breaking these rules might be an option if the goal is to study distraction instead of suppression. Applicability of Rules 4–6 depends on the type of study design and research question under consideration. These rules are furthermore important when the goal is to integrate results across studies. Rules 7–10 discuss theoretical considerations that have important implications for the type of inference that can (or cannot) be derived from a certain kind of study. Thus, not complying with some of these rules is fine if the researchers make sure that this does not interfere with the primary goals of the investigation, and is considered when theorizing about the study's implications.

Compliance with all rules is no guarantee for an excellent study, just as noncompliance with some of the rules might still result in an important and insightful contribution. Table 1 provides a checklist that summarizes the benefits and costs of following – or not following – each rule. We hope that this checklist serves as a useful guide for researchers who are setting up new studies to investigate distractor suppression and those who are reviewing and critically examining the literature on distractor suppression.

CRedit authorship contribution statement

Malte Wöstmann: Conceptualization, Writing – original draft, Writing – review & editing. **Viola S. Störmer:** Conceptualization, Writing – original draft, Writing – review & editing. **Jonas Obleser:**

Table 1
Checklist of ten simple rules to study distractor suppression.

Rule	Complied?	If complied ...	If not complied ...
1. Make sure the distractor has the potency to distract	o	We gain justification to refer to the task-irrelevant stimulus as 'distractor'.	We might end up trying to study the phenomenon of distractor suppression without a proper implementation of distraction.
2. Manipulate the distractor independently of the target	o	We can unambiguously assign effects to either target enhancement or distractor suppression.	We hazard (partial) confounds of target enhancement and distractor suppression, which might limit interpretability or results.
3. Test whether distractors are suppressed below a pre-defined baseline	o	We employ a neutral baseline to test whether distractors are truly 'suppressed' below baseline levels.	It remains unclear if distractors were truly suppressed or if they were processed similarly to other nontarget items.
4. Consider intentions versus experiences as sources of suppression	o	We isolate specific sources of distractor suppression, particularly those that might change over time.	We risk confounding sources of suppression and misattributing observed effects.
5. Distinguish between proactive and reactive suppression	o	We obtain precise temporal profiles of distractor suppression effects, which are important for the construction of testable models.	We accept underspecification of distractor processing effects in time, which complicates integration of studies.
6. Do not confuse psychological with neurophysiological suppression	o	We discern suppression of distractors in behavior and suppression of brain activity, which also enables us to test whether and how these two relate.	We risk the use of vague terminology and confuse neural processes and behavioral observations, which might give rise to unjustified conclusions about suppression.
7. Define your model of distractor suppression	o	We use a precisely defined model of suppression that allows us to derive testable hypotheses.	We risk the assignment of empirical results to misspecified mechanisms that lack a model.
8. Unravel causal implications for distractor suppression	o	We can address whether changes in the putative cause are sufficient for observed suppression effects.	We might fall into the trap of drawing unjustified conclusions on causality (implicitly or explicitly).
9. Beware of what statistical tests do and do not reveal about suppression	o	We avoid overinterpreting results on distractor suppression and provide effect size estimates necessary for good study design.	We risk drawing conclusions not supported by the empirical data, which impedes progress beyond the level of the individual study.
10. Consider distraction in the lab versus in the real world	o	We are aware that generalization of laboratory studies is limited and we aim for experiments in more realistic environments.	We might mistake distractor suppression in a standardized laboratory task for suppression in the real world.

Conceptualization, Writing – original draft, Writing – review & editing. **Douglas A. Addleman:** Writing – original draft, Writing – review & editing. **Søren K. Andersen:** Writing – original draft, Writing – review & editing. **Nicholas Gaspelin:** Writing – original draft, Writing – review & editing. **Joy J. Geng:** Writing – original draft, Writing – review & editing. **Steven J. Luck:** Writing – original draft, Writing – review & editing. **MaryAnn P. Noonan:** Writing – original draft, Writing – review & editing. **Heleen A. Slagter:** Writing – original draft, Writing – review & editing. **Jan Theeuwes:** Writing – original draft, Writing – review & editing.

Conflict of interest

The authors declare no conflict of interest.

References

- Abeles, D., Yuval-Greenberg, S., 2021. Active sensing and overt avoidance: Gaze shifts as a mechanism of predictive avoidance in vision. *Cognition* 211, 104648. <https://doi.org/10.1016/j.cognition.2021.104648>.
- Adam, K.C.S., Serences, J.T., 2021. History modulates early sensory processing of salient distractors. *J. Neurosci.* 41 (38), 8007–8022. <https://doi.org/10.1523/JNEUROSCI.3099-20.2021>.
- Addleman, D.A., Jiang, Y.V., 2019. Experience-driven auditory attention. *Trends Cogn. Sci.* 23 (11), 927–937. <https://doi.org/10.1016/j.tics.2019.08.002>.
- Addleman, D.A., Störmer, V.S., 2022. No evidence for proactive suppression: Enhancement of explicitly cued distractor features. *Psychon. Bull. Rev.* <https://doi.org/10.3758/s13423-022-02071-7>.
- Amrhein, V., Korner-Nievergelt, F., Roth, T., 2017. The earth is flat ($p > 0.05$): Significance thresholds and the crisis of unreplicable research. *PeerJ* 5, e3544. <https://doi.org/10.7717/peerj.3544>.
- Andersen, S.K., Müller, M.M., 2010. Behavioral performance follows the time course of neural facilitation and suppression during cued shifts of feature-selective attention. *Proc. Natl. Acad. Sci.* 107 (31), 13878–13882. <https://doi.org/10.1073/pnas.1002436107>.
- Anderson, B., 2011. There is no such thing as attention. *Front. Psychol.* 2, 246. <https://doi.org/10.3389/fpsyg.2011.00246>.
- Anderson, B.A., 2016. The attention habit: How reward learning shapes attentional selection. *Ann. N. Y. Acad. Sci.* 1369 (1), 24–39. <https://doi.org/10.1111/nyas.12957>.
- Anderson, B.A., 2021. Time to stop calling it attentional “capture” and embrace a mechanistic understanding of attentional priority. *Vis. Cogn.* 0 (0), 1–4. <https://doi.org/10.1080/13506285.2021.1892894>.
- Antonov, P.A., Chakravarthi, R., Andersen, S.K., 2020. Too little, too late, and in the wrong place: Alpha band activity does not reflect an active mechanism of selective attention. *NeuroImage* 219, 117006. <https://doi.org/10.1016/j.neuroimage.2020.117006>.
- Arita, J.T., Carlisle, N.B., Woodman, G.F., 2012. Templates for rejection: Configuring attention to ignore task-irrelevant features. *J. Exp. Psychol. Hum. Percept. Perform.* 38 (3), 580–584. <https://doi.org/10.1037/a0027885>.
- Awh, E., Matsukura, M., Serences, J.T., 2003. Top-down control over biased competition during covert spatial orienting. *J. Exp. Psychol.: Hum. Percept. Perform.* 29 (1), 52–63. <https://doi.org/10.1037/0096-1523.29.1.52>.
- Bacon, W.F., Egeth, H.E., 1994. Overriding stimulus-driven attentional capture. *Percept. Psychophys.* 55 (5), 485–496. <https://doi.org/10.3758/bf03205306>.
- Becker, S.I., Folk, C.L., Remington, R.W., 2010. The role of relational information in contingent capture. *J. Exp. Psychol.: Hum. Percept. Perform.* 36 (6), 1460–1476. <https://doi.org/10.1037/a0020370>.
- Bonnefond, M., Jensen, O., 2012. Alpha oscillations serve to protect working memory maintenance against anticipated distractors. *Curr. Biol.* 22 (20), 1969–1974. <https://doi.org/10.1016/j.cub.2012.08.029>.
- Born, S., Kerzel, D., Theeuwes, J., 2011. Evidence for a dissociation between the control of oculomotor capture and disengagement. *Exp. Brain Res.* 208 (4), 621–631. <https://doi.org/10.1007/s00221-010-2510-1>.
- Braver, T.S., Gray, J.R., Burgess, G.C., 2007. Explaining the many varieties of working memory variation: Dual mechanisms of cognitive control. *Var. Work. Mem.* 75, 76–106.
- Bundesen, C., Kyllingsbaek, S., Houmann, K.J., Jensen, R.M., 1997. Is visual attention automatically attracted to one’s own name? *Percept. Psychophys.* 59 (5), 714–720. <https://doi.org/10.3758/bf03206017>.
- Buttaccio, D.R., Hahn, S., 2011. The influence of action on visual search: Behavioral response toward stimuli modifies the selection process. *Atten., Percept. Psychophys.* 73 (5), 1453–1466. <https://doi.org/10.3758/s13414-011-0121-y>.
- Carlisle, N.B., Nitka, A.W., 2019. Location-based explanations do not account for active attentional suppression. *Vis. Cogn.* 27 (3–4), 305–316. <https://doi.org/10.1080/13506285.2018.1553222>.
- Carretié, L., 2014. Exogenous (automatic) attention to emotional stimuli: A review. *Cogn., Affect., Behav. Neurosci.* 14 (4), 1228–1258. <https://doi.org/10.3758/s13415-014-0270-2>.
- Chang, S., Egeth, H.E., 2019. Enhancement and suppression flexibly guide attention. *Psychol. Sci.* 30 (12), 1724–1732. <https://doi.org/10.1177/0956797619878813>.
- Chelazzi, L., Marini, F., Pascucci, D., Turatto, M., 2019. Getting rid of visual distractors: The why, when, how, and where. *Curr. Opin. Psychol.* 29, 135–147. <https://doi.org/10.1016/j.copsyc.2019.02.004>.
- Cherry, E.C., 1953. Some experiments on the recognition of speech, with one and with two ears. *J. Acoust. Soc. Am.* <https://doi.org/10.1121/1.1907229>.
- Chetverikov, A., Campana, G., Kristjánsson, Á., 2016. Building ensemble representations: How the shape of preceding distractor distributions affects visual search. *Cognition* 153, 196–210. <https://doi.org/10.1016/j.cognition.2016.04.018>.
- Cunningham, C.A., Egeth, H.E., 2016. Taming the white bear: Initial costs and eventual benefits of distractor inhibition. *Psychol. Sci.* 27 (4), 476–485. <https://doi.org/10.1177/0956797615626564>.
- Dienes, Z., 2014. Using Bayes to get the most out of non-significant results. *Front. Psychol.* 5, 781. <https://doi.org/10.3389/fpsyg.2014.00781>.
- Dingus, T.A., Guo, F., Lee, S., Antin, J.F., Perez, M., Buchanan-King, M., Hankey, J., 2016. Driver crash risk factors and prevalence evaluation using naturalistic driving data. *Proc. Natl. Acad. Sci.* 113 (10), 2636–2641. <https://doi.org/10.1073/pnas.1513271113>.
- Enriquez-Geppert, S., Huster, R.J., Herrmann, C.S., 2017. EEG-Neurofeedback as a tool to modulate cognition and behavior: A review tutorial. *Front. Hum. Neurosci.* 11 <https://www.frontiersin.org/article/10.3389/fnhum.2017.00051>.
- Failing, M., Theeuwes, J., 2018. Selection history: How reward modulates selectivity of visual attention. *Psychon. Bull. Rev.* 25 (2), 514–538. <https://doi.org/10.3758/s13423-017-1380-y>.
- Feldmann-Wüstefeld, T., Busch, N.A., Schubö, A., 2020. Failed suppression of salient stimuli precedes behavioral errors. *J. Cogn. Neurosci.* 32 (2), 367–377. <https://doi.org/10.1162/jocn.a.01502>.
- Ferrante, O., Patacca, A., Di Caro, V., Della Libera, C., Santandrea, E., Chelazzi, L., 2018. Altering spatial priority maps via statistical learning of target selection and distractor filtering. *Cortex: A J. Devoted Study Nerv. Syst. Behav.* 102, 67–95. <https://doi.org/10.1016/j.cortex.2017.09.027>.
- Folk, C.L., Remington, R.W., Johnston, J.C., 1992. Involuntary covert orienting is contingent on attentional control settings. *J. Exp. Psychol.: Hum. Percept. Perform.* 18 (4), 1030–1044. <https://doi.org/10.1037/0096-1523.18.4.1030>.
- Forster, S., Lavie, N., 2016. Establishing the Attention-Distractibility Trait. *Psychol. Sci.* 27 (2), 203–212. <https://doi.org/10.1177/0956797615617761>.
- Foster, J.J., Awh, E., 2019. The role of alpha oscillations in spatial attention: Limited evidence for a suppression account. *Curr. Opin. Psychol.* 29, 34–40. <https://doi.org/10.1016/j.copsyc.2018.11.001>.
- Frings, C., Bader, R., Spence, C., 2008. Selection in touch: Negative priming with tactile stimuli. *Percept. Psychophys.* 70 (3), 516–523. <https://doi.org/10.3758/PP.70.3.516>.
- Friston, K., 2009. The free-energy principle: A rough guide to the brain? *Trends Cogn. Sci.* 13 (7), 293–301. <https://doi.org/10.1016/j.tics.2009.04.005>.
- Fukuda, K., Vogel, E.K., 2011. Individual differences in recovery time from attentional capture. *Psychol. Sci.* 22 (3), 361–368. <https://doi.org/10.1177/0956797611398493>.
- Funder, D.C., Ozer, D.J., 2019. Evaluating effect size in psychological research: Sense and nonsense. *Adv. Methods Pract. Psychol. Sci.* 2 (2), 156–168. <https://doi.org/10.1177/2515245919847202>.
- Gaspar, J.M., McDonald, J.J., 2014. Suppression of salient objects prevents distraction in visual search. *J. Neurosci.* 34 (16), 5658–5666. <https://doi.org/10.1523/JNEUROSCI.4161-13.2014>.
- Gaspelin, N., Luck, S.J., 2018a. Distinguishing among potential mechanisms of singleton suppression. *J. Exp. Psychol.: Hum. Percept. Perform.* 44 (4), 626–644. <https://doi.org/10.1037/xhp0000484>.
- Gaspelin, N., Luck, S.J., 2018b. The role of inhibition in avoiding distraction by salient stimuli. *Trends Cogn. Sci.* 22 (1), 79–92. <https://doi.org/10.1016/j.tics.2017.11.001>.
- Gaspelin, N., Luck, S.J., 2018c. Combined electrophysiological and behavioral evidence for the suppression of salient distractors. *J. Cogn. Neurosci.* 30 (9), 1265–1280. <https://doi.org/10.1162/jocn.a.01279>.
- Gaspelin, N., Luck, S.J., 2019. Inhibition as a potential resolution to the attentional capture debate. *Curr. Opin. Psychol.* 29, 12–18. <https://doi.org/10.1016/j.copsyc.2018.10.013>.
- Gaspelin, N., Leonard, C.J., Luck, S.J., 2015. Direct evidence for active suppression of salient-but-irrelevant sensory inputs. *Psychol. Sci.* 26 (11), 1740–1750. <https://doi.org/10.1177/0956797615597913>.
- Gaspelin, N., Leonard, C.J., Luck, S.J., 2017. Suppression of overt attentional capture by salient-but-irrelevant color singletons. *Atten., Percept., Psychophys.* 79 (1), 45–62. <https://doi.org/10.3758/s13414-016-1209-1>.
- Geng, J.J., 2014. Attentional mechanisms of distractor suppression. *Curr. Dir. Psychol. Sci.* 23 (2), 147–153. <https://doi.org/10.1177/0963721414525780>.
- Geng, J.J., DiQuattro, N.E., 2010. Attentional capture by a perceptually salient non-target facilitates target processing through inhibition and rapid rejection. *J. Vis.* 10 (6), 5. <https://doi.org/10.1167/10.6.5>.
- Geng, J.J., Witkowski, P., 2019. Template-to-distractor distinctiveness regulates visual search efficiency. *Curr. Opin. Psychol.* 29, 119–125. <https://doi.org/10.1016/j.copsyc.2019.01.003>.
- Geng, J.J., DiQuattro, N.E., Helm, J., 2017. Distractor probability changes the shape of the attentional template. *J. Exp. Psychol.: Hum. Percept. Perform.* 43 (12), 1993–2007. <https://doi.org/10.1037/xhp0000430>.
- Geng, J.J., Won, B.-Y., Carlisle, N.B., 2019. Distractor ignoring: Strategies, learning, and passive filtering. *Curr. Dir. Psychol. Sci.* 28 (6), 600–606. <https://doi.org/10.1177/0963721419867099>.

- Utzerath, C., Fernández, G., 2017. Shaping science for increasing interdependence and specialization. *Trends Neurosci.* 40 (3), 121–124. <https://doi.org/10.1016/j.tics.2016.12.005>.
- Vadillo, M.A., Giménez-Fernández, T., Beesley, T., Shanks, D.R., Luque, D., 2021. There is more to contextual cuing than meets the eye: Improving visual search without attentional guidance toward predictable target locations. *J. Exp. Psychol.: Hum. Percept. Perform.* 47 (1), 116–120. <https://doi.org/10.1037/xhp0000780>.
- Van der Stigchel, S., Theeuwes, J., 2006. Our eyes deviate away from a location where a distractor is expected to appear. *Exp. Brain Res.* 169 (3), 338–349. <https://doi.org/10.1007/s00221-005-0147-2>.
- van Moorselaar, D., Slagter, H.A., 2019. Learning what is irrelevant or relevant: Expectations facilitate distractor inhibition and target facilitation through distinct neural mechanisms. *J. Neurosci.: Off. J. Soc. Neurosci.* 39 (35), 6953–6967. <https://doi.org/10.1523/JNEUROSCI.0593-19.2019>.
- van Moorselaar, D., Slagter, H.A., 2020. Inhibition in selective attention. *Ann. N. Y. Acad. Sci.* 1464 (1), 204–221. <https://doi.org/10.1111/nyas.14304>.
- van Moorselaar, D., Lampers, E., Cordesius, E., Slagter, H.A., 2020. Neural mechanisms underlying expectation-dependent inhibition of distracting information. *eLife* 9, e61048. <https://doi.org/10.7554/eLife.61048>.
- Vatterott, D.B., Vecera, S.P., 2012. Experience-dependent attentional tuning of distractor rejection. *Psychon. Bull. Rev.* 19 (5), 871–878. <https://doi.org/10.3758/s13423-012-0280-4>.
- Vatterott, D.B., Mozer, M.C., Vecera, S.P., 2018. Rejecting salient distractors: Generalization from experience. *Atten., Percept. Psychophys.* 80 (2), 485–499. <https://doi.org/10.3758/s13414-017-1465-8>.
- Võ, M.L.-H., Boettcher, S.E., Draschkow, D., 2019. Reading scenes: How scene grammar guides attention and aids perception in real-world environments. *Curr. Opin. Psychol.* 29, 205–210. <https://doi.org/10.1016/j.copsyc.2019.03.009>.
- de Vries, I.E.J., Slagter, H.A., Olivers, C.N.L., 2020. Oscillatory control over representational states in working memory. *Trends Cogn. Sci.* 24 (2), 150–162. <https://doi.org/10.1016/j.tics.2019.11.006>.
- Wang, B., Theeuwes, J., 2018a. Statistical regularities modulate attentional capture. *J. Exp. Psychol.: Hum. Percept. Perform.* 44 (1), 13–17. <https://doi.org/10.1037/xhp0000472>.
- Wang, B., Theeuwes, J., 2018b. How to inhibit a distractor location? Statistical learning versus active, top-down suppression. *Atten. Percept. Psychophys.* 80 (4), 860–870. <https://doi.org/10.3758/s13414-018-1493-z>.
- Wang, B., Theeuwes, J., 2018c. Statistical regularities modulate attentional capture independent of search strategy. *Atten., Percept. Psychophys.* 80 (7), 1763–1774. <https://doi.org/10.3758/s13414-018-1562-3>.
- Wang, B., Theeuwes, J., 2020. Saliency determines attentional orienting in visual selection. *J. Exp. Psychol.: Hum. Percept. Perform.* 46 (10), 1051–1057. <https://doi.org/10.1037/xhp0000796>.
- Waschke, L., Kloosterman, N.A., Obleser, J., Garrett, D.D., 2021. Behavior needs neural variability. *Neuron* 109 (5), 751–766. <https://doi.org/10.1016/j.neuron.2021.01.023>.
- Wen, W., Hou, Y., Li, S., 2018. Memory guidance in distractor suppression is governed by the availability of cognitive control. *Atten., Percept., Psychophys.* 80 (5), 1157–1168. <https://doi.org/10.3758/s13414-018-1511-1>.
- Wolfe, J.M., 1994. Guided Search 2.0 A revised model of visual search. *Psychon. Bull. Rev.* 1 (2), 202–238. <https://doi.org/10.3758/BF03200774>.
- Wolfe, J.M. (2018). Visual Search. In Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience, J.T. Wixted (Ed.). (Bd. 2). Wiley.
- Won, B.-Y., Kosoyan, M., Geng, J.J., 2019. Evidence for second-order singleton suppression based on probabilistic expectations. *J. Exp. Psychol.: Hum. Percept. Perform.* 45 (1), 125–138. <https://doi.org/10.1037/xhp0000594>.
- Won, B.-Y., Forloines, M., Zhou, Z., Geng, J.J., 2020. Changes in visual cortical processing attenuate singleton distraction during visual search. *Cortex J. Devoted Study Nerv. Syst. Behav.* 132, 309–321. <https://doi.org/10.1016/j.cortex.2020.08.025>.
- Woodman, G.F., Luck, S.J., 2007. Do the contents of visual working memory automatically influence attentional selection during visual search? *J. Exp. Psychol. Hum. Percept. Perform.* 33 (2), 363–377. <https://doi.org/10.1037/0096-1523.33.2.363>.
- Wöstmann, M., Herrmann, B., Maess, B., Obleser, J., 2016. Spatiotemporal dynamics of auditory attention synchronize with speech. *Proc. Natl. Acad. Sci.* 113 (14), 3873–3878. <https://doi.org/10.1073/pnas.1523357113>.
- Wöstmann, M., Alavash, M., Obleser, J., 2019. Alpha oscillations in the human brain implement distractor suppression independent of target selection. *J. Neurosci.* 39 (49), 9797–9805. <https://doi.org/10.1523/JNEUROSCI.1954-19.2019>.