Visual Singleton Detection (‘Pop-out’) is Mediated by Dimension-based Attention

Joseph Krummenacher[1], Hermann J. Müller[2], Brit Reimann[2] & Dieter Heller[1]
[1]RWTH Aachen, [2]LMU München, Germany

Abstract

The work reported in this chapter is concerned with a seemingly simple problem, namely: how an odd-one-out object, a single-feature target, is segmented from a background of homogeneous nontarget objects when the target-defining dimension is not known in advance, that is, when the critical dimension varies from trial to trial, resulting in cross-dimensional target uncertainty. Our work argues that, under such conditions, the target does not simply ‘pop out’ of the field on the basis of some early, pre-attentive, segmentation mechanism operating in a purely bottom-up fashion. Rather, target segmentation involves an attentional mechanism that modifies the processing system by allocating a limited ‘selection weight’ to the various dimensions that potentially define the target.

Cross-dimension Costs and Inter-trial Transition Effects in Singleton Feature Search

<u>Visual search for odd-one-out feature targets.</u> It is well established that targets which differ from distractors in certain single salient attributes, or *features*, can be rapidly discerned irrespective of the number of items in the display (the set size). Phenomenally, the target appears to ‘pop out’ of the display (pop-out effect). Visual features that support set size-independent search are generally assumed to be registered *in parallel* across the visual field. Such features are regarded as primitive image descriptors organised along a set of *feature dimensions* (e.g., color and orientation). A number of feature dimensions have been shown to support parallel search, including: orientation, size, colour, stereo depth, and motion.

There are various accounts of how salient feature differences in the field may be detected. One influential account is Guided Search (GS) (Cave & Wolfe, 1990; Wolfe, 1994). GS assumes that the visual field is initially represented, in parallel, as a set of basic stimulus attributes in different dimension-specific ‘modules’ (such as colour, orientation etc.). Each module computes saliency signals for all stimulus locations, indicating the feature contrast between one particular item relative to the various other items represented within the same module: The more dissimilar an item is compared to the others, the greater its saliency. Maps of saliency signals are computed in parallel in all modules, and then these signals are summed onto a *master map of activations*. The activity on the master map guides *focal attention,* the most active location being sampled with priority. Focal attention gates the passage of information to higher stages of processing (visual object recognition and response systems). Thus, any odd-one-out feature target will generate a strong contrast signal within its own dimension. Even given some variability due to noise, the target’s saliency signal on the master map should always be larger than those of distractor items, and attention should always be deployed first to its location.

However, our recent work demonstrates that bottom-up models such as GS are, in a crucial respect, incomplete as an account of singleton feature search – in particular, when the dimension defining the target is uncertain (i.e., variable) on a trial. Dimensional uncertainty produces a cost in discerning the presence of a target (see also Treisman, 1988), which is inconsistent with the assumption that saliency signals from relevant dimensions are integrated by the master map units in a parallel and equally weighted fashion.

Visual search for singleton feature targets across dimensions. We (Müller, Heller & Ziegler, 1995) have recently investigated search for singleton feature targets within and across stimulus dimensions. In an initial experiment, search for three possible targets all defined *within* the orientation dimension (left-tilted, horizontal, and right-tilted small grey bars) was compared with search for three possible targets defined *across* three different - orientation, colour and size - dimensions (a *right-tilted* grey small bar, a vertical *black* small bar or a grey vertical *large* bar). The distractors in both uncertainty conditions, intra- and cross-dimension, were the same: small grey vertical bars. There was also a no-uncertainty control condition in which the target was always known to be a small grey *right-tilted* bar among small grey vertical bars. Observers were instructed to simply respond to the detection of any heterogeneity in the display, without processing its source any further. According to bottom-up accounts, search performance ought to be unaffected by whether or not observers can predict the dimension and the feature value defining the target on a particular trial.

However, although search was parallel in all conditions, the detection of the common right-tilted target was 60 ms slower in the cross-dimension condition relative to both the intra-dimension condition and the control condition - a considerable RT cost in view of the fast base RTs. That there was a RT cost only in the cross-dimension condition, but not the intra-dimension condition, suggests that, to detect the presence of a target, observers had to determine in which dimension a feature difference was present: orientation, color, or size.

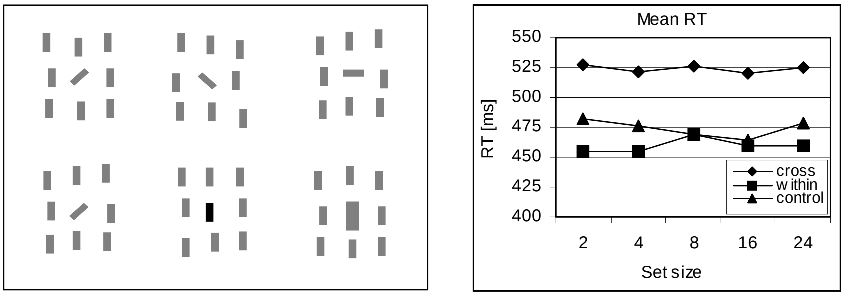

a) b)

Figure 1. Müller et al. (1995), Experiment 1. Figure 1a. Illustration of possible target displays in the *intra*-dimension (upper displays) and *cross*-dimension search conditions (lower displays). Figure 1b. Reaction times (RT) to displays with a right-tilted orientation target (target present) as a function of set size in the *control*, *intra*-dimension, and *cross*-dimension search conditions.

One further aspect of cross-dimension search performance is noteworthy: There was a RT advantage for a target on a given trial if the previous trial contained a target defined on the same, as compared to a different, dimension (dimension-specific inter-trial transition effect).

Dimension-specific inter-trial transition effect. A further experiment (Found & Müller, 1996, Experiment 1) demonstrated that the inter-trial facilitation is indeed *dimension*-specific rather than *feature*-specific in nature. Displays in these experiments contained, on positive trials, one of four possible targets: either a *left*- or a *right-tilted* white bar (orientation target) or a *red* or a *blue* vertical bar (colour target). If the inter-trial effect is dimension-specific, it should always be evident when the target dimension (e.g., colour) is repeated on consecutive trials irrespective of whether or not the target feature value (e.g., red) is repeated.

The results showed clear inter-trial facilitation of 30 to 40 ms when consecutive trials contained targets defined in the same dimension, relative to targets defined in different dimensions. This was the case irrespective of whether a target (on trial N) was preceded by a featurally identical target (on trial N-1) or by a dimensionally identical, but featurally non-identical target. For example, there was a RT advantage for a red target preceded by either a red or a blue target, relative to a preceding orientation target; but there was little (extra) advantage for a red target preceded by a red target, relative to a preceding blue target.

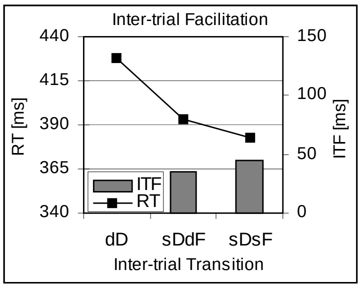

Figure 2. Found & Müller (1996), Experiment 1: Reaction times (RT) to a target on trial n dependent on the dimensional and featural identity of the target on trial n-1 (intertrial transition: dD = different dimension; sDdF = same dimension, different feature; sDsF = same dimension, same feature). Also presented is the intertrial facilitation (ITF) for same-dimension (sDdF, sDsF) targets relative to different-dimension (dD) targets.

Dimensional Weighting

We (Müller et al., 1995; Found & Müller, 1996) took the cross-dimension cost and inter-trial facilitation observed in our experiments to argue for a *dimension weighting account* of visual search for feature targets. Similar to GS, for a feature target to generate fast parallel search requires that it rapidly attracts focal attention. Focal attention operates on a master map of integrated (summed) dimension-specific saliency signals. However, unlike GS, dimension-specific saliency information is attentionally 'weighted' as it is transmitted to the master map of activations. Sufficient weight must be assigned to the target dimension for the target's saliency signal at the master map level to exceed the response threshold. In the intra-dimension conditions described above, the target dimension was always known, and so weighted in advance, permitting rapid search. That is, weights may be assigned according to the known likelihood of a target appearing in a particular dimension. However, in the cross-dimension conditions (without knowledge of the likely target dimension), the search involved a time-consuming *weight-shifting* process to determine the target's dimension and amplify its activity at the master map level. The weight pattern established in this process persists into the next trial, producing a dimension-specific RT advantage for a target defined within the same dimension as the preceding target.

Top-down Weighting of Dimensions

One further important question concerns the extent to which the weighting of dimensions is, or can be, top-down controlled. There is psychophysical evidence that, in simple singleton feature search tasks of the type described above, the target-defining dimension is determined and weighted relatively automatically, without involving deliberate (top-down) control operations. For example, Müller, Krummenacher, and Heller (2001) found that, in cross-dimensional search, observers did not explicitly encode and retain the target-defining dimension (or the target feature) on a given trial, and, even if they were required to do so by the task, this did not alter the pattern of dimension-specific inter-trial effects. On the other hand, there is evidence that observers can modulate the dimensional weight setting in a top-down fashion in response to symbolic pre-cues indicating the dimension within which the target is likely to be defined on a given trial. For example, Müller et al. (1995; Experiment 3) made one particular dimension the likely target-defining dimension throughout a block of trials (p=.80), symbolically indicating this dimension to the observers at the start of the block (blockwise dimension cueing). The result was that targets defined in the likely dimension were detected faster than targets in the unlikely dimensions. However, one problem with interpreting this cueing effect in terms of top-down control of dimensional weighting is that, by making one dimension more likely to define a target, targets actually appeared more often in the indicated dimension; that is, there was also a greater likelihood for consecutive targets to be defined in the cued (i.e., the same) dimension compared with targets in the unlikely dimensions. Given that a target on trial n defined within the same dimension as the target on trial n-1 benefits from the persistence of the dimensional weight setting from trial n-1 to trial n, the cueing effects revealed in this experiment might simply reflect passive, stimulus-driven, priming (see also Maljkovic & Nakayama, 1994), rather than active top-down control. To rule out such passive priming effects, it is necessary to demonstrate dimension cueing effects using trial-by-trial (rather than blockwise) cueing of the target-defining dimension.

We (Reimann, Müller, & Krummenacher, 2001) have recently conducted such a trial-by-trial cueing study. Precueing of the likely target-defining dimension on a given trial (by the cue words "color" or "orientation") produced RT benefits for valid-cue trials, on which the target was defined in the cued dimension, and costs for invalid-cue trials, on which the target was defined in an uncued dimension, relative to a neutral-cue condition (with the cue word "neutral"). Furthermore, the dimension-specific inter-trial effects were reduced for valid trials relative to neutral trials. Note that, even when a specific target feature (e.g., red) was pre-cued to be likely, while other features in the same dimension (i.e., color: e.g., blue, yellow) were unlikely, the cueing effects were dimension-specific in nature; that is, the were benefits of the cueing even for unlikely features within the same dimension as the cued features (while there were costs only for features in a different dimension). This was the case even when an (uncued) feature within the dimension of the cued feature was extremely rare. This pattern of results is consistent with the idea that observers can use the advance cue to set themselves for (i.e., allocate attentional weight to) the likely target dimension. However, the fact that there remained a residual inter-trial transition effect even with valid precues suggests that top-down control processes cannot completely overcome automatic priming processes.

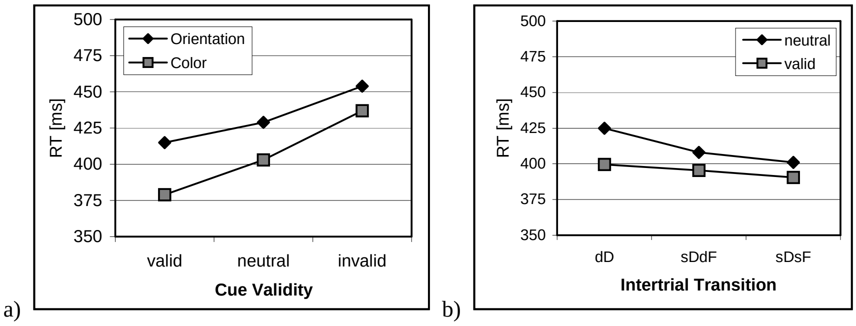

Figure 3. Reimann, Müller, & Krummenacher (2001), Experiment 1. Figure 3a. Reaction time (RT) as a function of cue validity, separately for color and orientation-defined targets. Figure 3b. RT as a function of intertrial transition, separately for neutral and valid-cue trials (intertrial transition: dD = different dimension; dDdF = same dimension, different feature; sDsF = same dimension, same feature).

Taken together, these results suggest that dimension switching can operate relatively automatically, in a largely stimulus-driven manner, once the basic operating parameters are set (e.g., between which dimensions switches must be carried out). However, dimension switching may also be top-down controlled when there is an advantage, or a need, to do so.

Parallel or Serial Weighting of Dimensions?

Although there is convergent evidence as to the existence of dimensional weighting, how does the weighting process actually work: Does it operate in a parallel, continuous, fashion across dimensions or in a serial, all-or-none, fashion? Although some theorists have advocated serial processing of dimensions (e.g., Treisman, 1988; Grossberg, et al., 1994), other evidence such as the increased variability of observers' RTs in cross-dimension search (Müller et al., 1995) points towards parallel processing. The dimension-weighting account as such makes no prediction as to whether singleton feature search across dimensions is serial in nature or parallel, and, if the latter, whether a parallel-race model is true or a parallel-coactivation model. Thus, the issue of serial versus parallel processing of dimensions is an empirical one, which has recently been investigated by Krummenacher, Müller, and Heller (2001a) by examining visual search for singleton feature targets *redundantly* defined in multiple dimensions; more specifically, by adapting the redundant-target detection paradigm (e.g., Mordkoff, Yantis, & Egeth, 1990; Mordkoff & Yantis, 1993) to cross-dimension search, permitting Miller's (1982) 'race model inequality' (RMI) to be tested.

Normally in redundant-target search, there can be one or two targets in the display (on present trials). Serial search models predict a redundancy gain such that mean RTs should be faster when there are two targets than when there is only one, simply because one of two targets has a higher chance of being encountered early in the search than a single target. However, when the entire distributions of RTs are analyzed, a form of redundancy gain may be revealed that is inconsistent with any strictly serial model. Miller (1982) demonstrated that all models that assume that each target produces an independent, separate activation must satisfy the following RMI: $P(RT<t/T_1\&T_2) \leq P(RT<t/T_1) + P(RT<t/T_2)$, where t is the time since display onset and T_1 and T_2 are targets 1 and 2. Importantly, this inequality entails that the fastest RTs to displays with redundant targets be no faster than RTs to displays with single targets; however, fast RTs may

occur more often with redundant targets. Violations of this inequality constitute evidence against serial processing, and in favor of parallel-coactive processing.

Applied to cross-dimension search, Krummenacher et al. (2001a, Experiment 1) varied the number of dimensions in which a single target was defined (instead of varying the number of targets in a display), for example: color only or orientation only (singly defined targets, e.g.: a red target or a 45°-tilted target), or color and orientation simultaneously (redundantly defined target: e.g.: a red 45°-tilted target). They could then examine, by testing for violations of the RMI, whether only one dimension (dimension-specific saliency signal) at a time can activate a response-relevant representation, or whether there is coactivation from multiple dimensions.

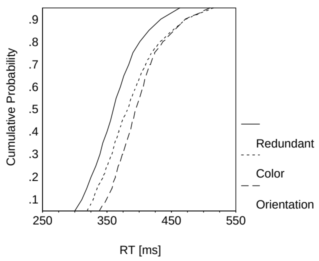

Figure 4. Krummenacher, Müller, & Heller (2001), Experiment 1: Cumulative RT distribution functions (CDFs) for singly-defined and redundantly-defined color and orientation targets (singly: dotted and dashed CDFs; redundantly: solid CDF).

Krummenacher et al. found that not only were RTs to redundantly defined targets on average faster than RTs to singly defined targets, but also that the fastest RTs to redundantly defined targets were faster than the fastest RTs to singly defined targets, violating the RMI. The second finding constitutes strong evidence in favor of dimension-specific saliency signals coactivating, or being integrated by, a common response-relevant (output) unit. The implication is that cross-dimension search for singleton feature targets does indeed proceed in parallel in multiple dimensions (e.g., Mordkoff & Yantis, 1993; Müller et al., 1995), rather than serially, in only one dimension at a time (e.g., Grossberg et al., 1994; Treisman, 1988).

Krummenacher, Müller, and Heller (2001b) went on to show that there is no evidence of coactivation when there are dual (redundant) saliency signals, at separate locations, defined within the same dimension (e.g., a red and a blue color target), consistent with Mordkoff and Yantis (1993). Furthermore, when there are dual (redundant) signals defined in different dimensions (e.g., a red color target and a 45°-tilted orientation target), evidence for co-activation is found only when the two signals are spatially adjacent, and even in this case the evidence tends to be weaker compared to when there is a single target redundantly defined on two dimension (e.g., a red 45°-tilted orientation target), that is, with two saliency signals at the same location. This pattern of effects suggests that there is signal integration only for saliency signals from separate dimensions, and the integration is spatially specific. This is consistent with the dimension-weighting account, according to which saliency signals from multiple dimensions can combine to raise the activation of the master map unit signaling the presence and location of the target above the value achieved by a single dimensional saliency signal.

Locus of Dimension Weighting: Perceptual or Response-related?

Although the notion of dimension weighting as such is 'agnostic' with respect to the locus of dimensional-uncertainty and redundancy effects, Müller and his colleagues (e.g., Müller et al., 1995) interpreted these effects as arising at a perceptual stage of processing. This interpretation has recently been challenged by Cohen and Magen (1999) who argued that these effects reflect response stage processes. According to their 'response-based' account, the various dimensional processing modules (e.g., color, motion, orientation etc.) possess separate response selection mechanisms. Effects of dimensional uncertainty in singleton feature search can then be explained by assuming that: "an intradimensional [search] task [target-defining dimension fixed] requires the use of a single response selection mechanism. By contrast, cross-dimensional tasks [target-defining dimension variable] require the use of multiple response selection mechanisms" (Cohen & Magen, 1999, p. 292). Similarly, dimension-specific intertrial facilitation in cross-dimension search can be explained by assuming that the relevant response selection mechanism is primed by repeated targets within the same dimension (irrespective of whether the target feature is repeated or not). A similar account could be derived for detection RTs to dimensionally redundant targets violating the RMI: Redundant targets activate separate dimension-specific response selection mechanisms which, in turn, drive a common response output stage in a parallel-coactive manner.

There is psychophysical evidence for the perceptual account, in particular, the demonstration by Krummenacher et al. (2001b) that saliency signal integration is spatially specific (see above). This finding would require response-based accounts to assume dimension-specific response selection mechanisms for each location in the field, which would make them indistinguishable from the perceptual account. Further evidence is provided by Reimann, Schröger, Müller, and Krummenacher's (2001) analysis of event-related brain potentials (ERPs) in visual search for dimensionally redundant singleton feature targets. Reimann et al. found that, initially, redundant targets were processed separately, similar to the 'strongest activated' dimension at frontal and, respectively, occipital electrode sites. ERP evidence for integration (i.e., violations of the additivity of the ERP effects expected on separate-activation models) emerged 210 ms after display onset, but 150-160 before overt responding, at both frontal and occipital sites. Reimann et al. took this to suggest that the integration occurs at a relatively early stage of (perceptual) processing: overall-saliency computation.

However, these findings do not tell what is actually 'weighted' attentionally in dimension weighting: the computation of dimension-specific saliency signals within the respective visual input modules, or the transfer of dimension-specific saliency signals to the integration stage. That is, does dimensional weighting influence (e.g., enhance) early saliency signal computation or (e.g., amplify) the subsequent signal transfer? Evidence in favor of the former is provided by a recent functional magnetic resonance imaging (fMRI) study of singleton feature search (Pollmann, Weidner, Müller, & von Cramon, 2000), which examined event-related activation changes accompanying changes in the target-defining dimension across trials, specifically, changes from color to motion and vice versa. Changes in the target-defining dimension (but not changes in the target feature within a constant dimension) led to increased activation in a fronto-posterior network consisting of left frontopolar cortex (BA 10) and inferior frontal gyri, high-level visual processing areas in parietal cortex and temporal cortex, and dorsal occipital visual areas. When attention was shifted to a new target-defining dimension, activation increased in the visual areas involved in the processing features of this dimension. Pollmann et al. hypothesized that frontopolar cortex is involved in controlling attentional weight shifting and that inferior frontal gyri and high-level parietal and temporal areas mediate attentional weighting via feedback to extrastriate visual areas that process the features of the new target dimension.

In summary, the functional-imaging evidence suggests that dimension weighting is mediated by frontal-lobe mechanisms and involves the modulation of neuronal activity in extrastriate visual areas specialized in the processing of features of the respective dimensions. The latter is consistent with the view that dimension weighting is perceptual in nature, influencing (enhancing) the computation of dimension-specific saliency signals.

Acknowledgements. The work presented in this article was supported by DFG grants HE 1192/5-1,2 to D. Heller and H.J. Müller, SCHR 375/8-1,2 to H.J. Müller, S. Pollmann & D.Y. von Cramon (Leipzig working memory research group), and MU 773/3-1 to H.J. Müller and J. Krummenacher. Correspondence concerning this article should be addressed to: J. Krummenacher, RWTH Aachen, Department of Psychology, Jägerstrasse 17/19, 52056 Aachen; e-mail: joseph.krummenacher@post.rwth-aachen.de.

References

Cave, K.R. & Wolfe, J.M. (1990). Modelling the role of parallel processing in visual search. *Cognitive Psychology*, **22**, 225-271.

Cohen, A., & Magen, H. (1999). Intra- and cross-dimensional visual search for single feature targets. *Perception & Psychophysics*, **61**, 291-307.

Found, A. P., & Müller, H. M. (1996). Searching for feature targets on more than one dimension: Investigating a dimension weighting account. *Perception & Psychophysics*, **1**, 88-101.

Grossberg, S., Mingolla, E., & Ross, W.D. (1994). A neural theory of attentive visual search: Interaction of boundary, surface, spatial, and object representations. *Psychological Review*, **101**, 470-489.

Krummenacher, J., Müller, H.J. & Heller, D. (2001a). Visual search for dimensionally redundant pop-out targets: Evidence for parallel-coactive processing of dimensions. *Perception & Psychophysics*, (in press)

Krummenacher, J., Müller, H.J., Heller, D. (2001b). Visual search for dimensionally redundant pop-out targets: Parallel-coactive processing of dimensions is location-specific (submitted for publication).

Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, **22**, 657-672.

Mordkoff, J.T., Yantis, S. (1993). Dividing attention between color and shape: evidence of coactivation. *Perception & Psychophysics* ,**53**, 357-66.

Müller, H.J., Heller, D., & Ziegler, J. (1995). Visual search for singleton feature targets within and across feature discriminations. *Perception & Psychophysics*, **57**, (1), 1-17.

Müller, H.J., Krummenacher, J. & Heller, D. (2001). Dimension-specific inter-trial facilitation in visual search for pop-out targets: evidence for an implicit visual short-term memory effect, (submitted for publication).

Pollmann, S., Weidner, R., Müller, H.J., von Cramon, D.Y. (2000). A fronto-posterior network involved in visual dimension changes. *Journal of Cognitiv Neuroscience*, **12**, 480-94.

Reimann, B. Müller, H.J. & Krummenacher, J. (2001). Stimulus-driven and expectancy-driven Effects in Dimensional Weighting, (submitted for publication).

Reimann, B., Schröger, E., Müller, H.J., Krummenacher, J. (2001). Search for redundant targets is based on co-activation of dimensional information. A combined behavioral and electrophysiological study, (submitted for publication).

Treisman, A. (1988). Features and objects: the fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology*, **40A**, 201-236.

Wolfe, J. M. (1994). Guided search 2.0 - a revised model of visual search. *Psychonomic Bulletin and Review*, **1**, 202-238.