

The redundant-signals paradigm and preattentive visual processing

Michael Zehetleitner¹, Hermann Josef Muller¹, Joseph Krummenacher²

¹Ludwig-Maximilians-Universitat, Germany, ²Universite de Fribourg, Switzerland

TABLE OF CONTENTS

1. Abstract
2. What redundancy gains reveal about preattentive visual processing
 - 2.1. The redundant-target paradigm
 - 2.2. The question of architecture: parallel, co-active, or serial?
 - 2.3. Is integration spatially selective?
 - 2.4. Is integration feature-based or dimension-based?
 - 2.5. Weighting or priming?
 - 2.6. Implementation of saliency maps and dimensional weighting in the brain
 - 2.7. Conclusion
3. Acknowledgments
4. References

1. ABSTRACT

Physiological and cognitive models of vision agree that the preattentive processing of visual stimuli is organized in a parallel and segregated fashion. However, several incompatible models have been proposed for the subsequent processing stages. They differ in their assumptions about architecture (serial, parallel, or coactive/integrative), stopping-rules (self-terminating, or exhaustive), spatial specificity of saliency signal coding (signal pooling across locations, or spatially distinct processing), and dependency of target detection on the prior allocation of attention (preattentive, or post-selective). We review how studies employing the redundant-signals paradigm in visual pop-out search contribute to discerning between the different assumptions. We find strong support for the notion of a saliency map, into which feature contrast signals are pooled, and especially the dimension weighting account (1) receives further support: Instead of a priming mechanism that could increase weights for several dimensions independently, evidence favors a weighting mechanism that effectively limits the total weight available for allocation to the various dimensions through competitive interactions, whereby increasing the weight for one dimension goes along with decreasing the weights for other dimensions.

2. WHAT REDUNDANCY GAINS REVEAL ABOUT PREATTENTIVE VISUAL PROCESSING

In their seminal work, Livingstone and Hubel showed that visual processing operates in separate and parallel 'channels' from the retinal level onwards (2-6). The separation of retinal cells specialized for high temporal- and, respectively, high spatial-frequency information is maintained in the laminar network of the lateral geniculate nucleus (LGN) and further in cortical areas. One pathway specialized for extracting motion information runs through distinct layers of the LGN, V1, and V2 on to the medial temporal area (MT), whereas the other pathway coding color and form information runs through distinct layers and sections (blobs, inter-blobs of V1, and thin-, inter-stripes of V2) on to V4 and higher-level areas in infero-temporal cortex. Although the exact nature of these pathways is under debate (see e.g. 7 for a review), the basic finding that neuronal visual processing operates in functionally separate, parallel pathways is undisputed.

Many cognitive and neuro-cognitive models are based on this initial parallel processing of different aspects of the visual scene – for example, Feature Integration Theory (FIT: 8, 9), Guided Search (GS: 10), Dimension Weighting (DW: 1, 11), and the neuro-computational models of Koch and Ullman (12) and Itti and Koch (13).

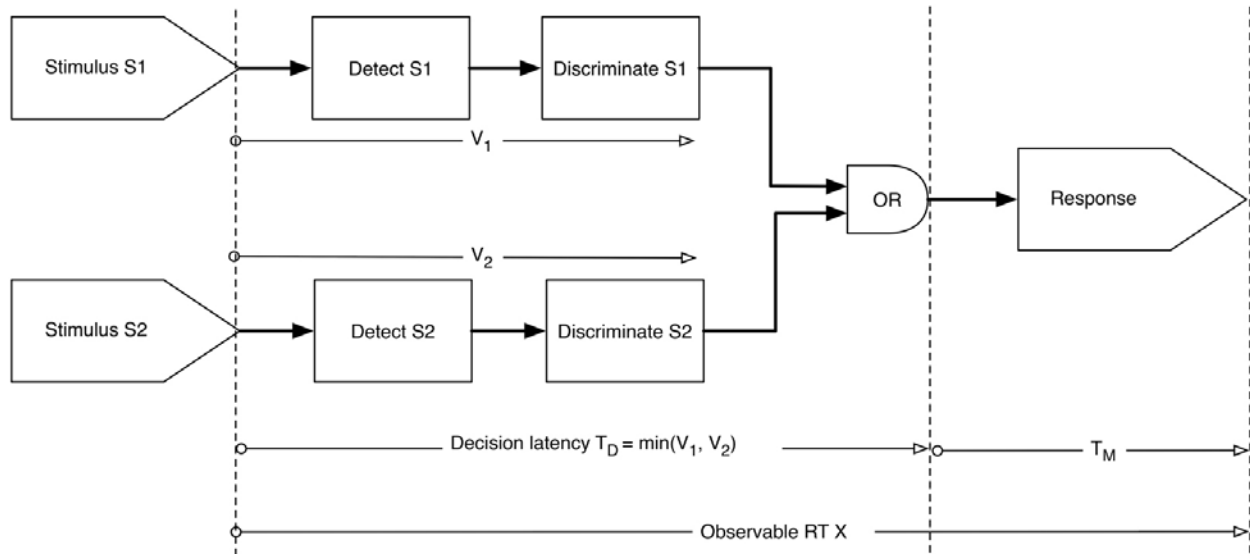


Figure 1. The observed reaction time in a race model is the sum of T_M , the time necessary for response and motor processing, and the minimum of the two detection times T_D for signals S_1 and S_2 . Each single signal has its own decision unit that can trigger the response. Figure adapted with permission from ref. 28.

These models largely agree on the initial processing stages: Local feature contrast is computed in parallel for separate visual dimensions (e.g., color, motion, orientation; see 14, 15, for reviews). The models differ, however, in the assumptions about later stages that follow the initial parallel coding stage. In this paper, we summarize recent evidence from studies of redundancy gains in visual pop-out search, which can help to determine the nature of these further processing steps. In pop-out search, the observers' task is to detect a singleton target which differs from distracting (non-target) items in one or more dimensions, such as color or orientation. Search time for such targets is usually independent of the number of display items, that is, all items in the display are searched efficiently (14).

While there is general agreement that features are initially registered in parallel in dimensionally segregated pathways, there are various models of the subsequent processing stages required for successful target detection. These models differ in a number of respects with regard to the further stages of processing necessary for arriving at a response decision: (i) the basic processing architecture (serial, parallel, vs. integrated processing of dimensions), (ii) stopping rules for the search process (self-terminating vs. exhaustive), (iii) spatial specificity of saliency signal coding (signal pooling across locations vs. spatially distinct processing), and (iv) dependency of target detection on the prior allocation of focal attention (preattentive vs. post-selective).

First, we will describe the basic redundant-target paradigm and its applicability to visual pop-out search. Second, we will elaborate each of the above questions and review relevant studies that show how analysis of redundancy gains can contribute to answering them. Third, we will review evidence from single-cell recordings in

monkeys and fMRI studies in humans concerning the neural implementation of preattentive search processes.

2.1. The redundant-target paradigm

In a redundant-target paradigm, participants have to respond as soon as a critical stimulus – that is, an element of a predefined set of target stimuli – appears. Each target of this set is mapped to the same response. Performance for trials on which only a single target is presented (single-signal trial, SST) is then compared to performance for trials on which two targets are presented (redundant-signals trial, RST). The presentation of redundant signals can be achieved by presenting the same target element simultaneously at two locations, or two different elements of the target-defining set at the same location or at different locations. Frequently, a benefit of RSTs over SSTs has been found, which is referred to as redundancy gain or redundant-signals effect (RSE) (e.g., 16-26).

Regardless of the specific nature of the two signals, the redundant-signals effect can be accounted for in two mutually exclusive ways: by 'independent parallel race' models or by 'coactivation' models. Raab (27) explained the redundant signal effect in terms of statistical facilitation. He proposed that the processing of redundant targets is analogous to a horse race. A RST is composed of two simultaneous single signals that are processed in parallel. The signal that triggers the response first wins the race and determines the observed reaction time (RT) (Figure 1). Thus, on each RST, the two single signals are processed in parallel by independent processors which accumulate activation in such a way that, once a threshold is exceeded for any of the two signals, the response is triggered. If the processing times of both single signals on a RST are drawn from the two SST reaction time distributions, it is highly probable that the processing time

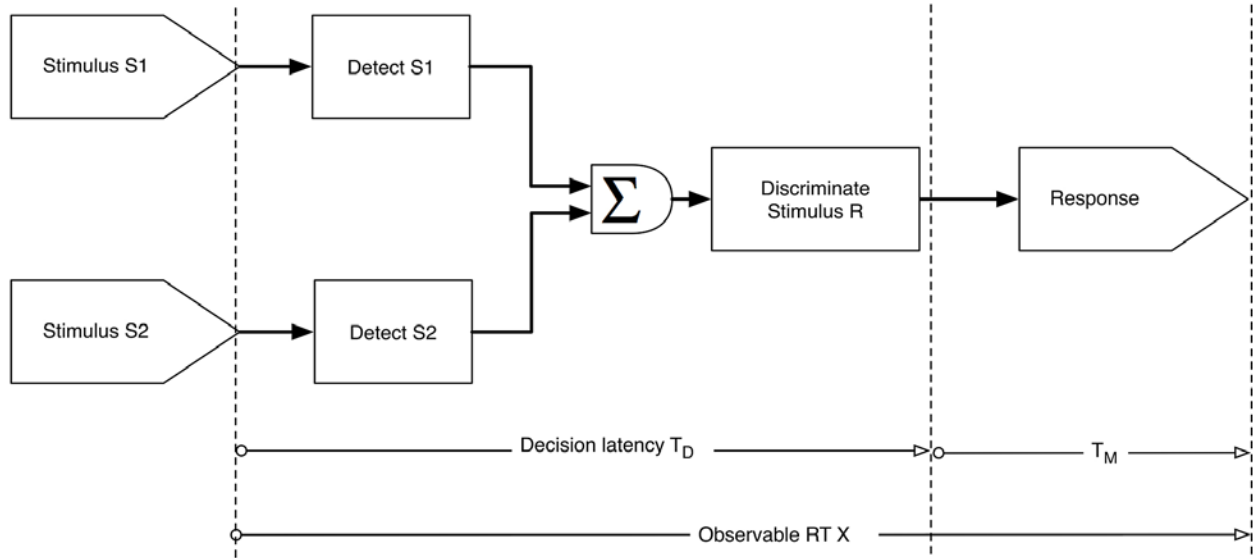


Figure 2. The observed reaction time in a co-activation model is the sum of T_M the time necessary for response and motor processing, and the decision latency T_D . Both single signals are fed into a common decision unit. Figure adapted with permission from ref. 28.

in one channel is faster/slower than that in the other, which leads to a faster mean processing time, because the faster of both channels determines the decision time. In more formalized terms, the minimum of two distributions (the processing times of both single signals on a RST) is less than each of the single distributions. In summary, Raab explained the redundant-signals effect by assuming an independent parallel race between the two signals on a RST, which leads to a reaction time distribution that is shifted to the left of the distributions of both SSTs, resulting in a faster mean reaction time for RSTs than for SSTs (i.e., statistical facilitation).

Miller (17) showed that there is an upper bound for the size of the redundant-signals effect under the assumption of a race model. If the redundancy gains exceed a certain boundary, given by the ‘race model inequality’ (RMI), then statistical facilitation can no longer account for this gain. The upper bound of the benefit deriving from redundant targets under the assumption of a race model is formulated in terms of reaction time distributions, rather than mean reaction times. The RMI makes use of the fact that, in the race model, the reaction time distribution for redundant targets is the minimum of the distributions of both single targets. That is, on a RST, the processing times for the single targets are drawn from the corresponding single-signal distributions. Statistically, one of the targets is almost always processed faster than the other, giving rise to the redundancy gain. This gain can be maximized if the distribution of processing times for single signals, rather than being stochastically independent, are maximally negatively dependent (29) – such that, if the processing time of one single signal is fast, then the processing time of the other single signal is slow. This upper bound is given by the sum of both single-signal distributions, and, under the assumption of a race model, the distribution of reaction times for redundant signals is always below (or does not

differ from) the sum of the reaction time distributions for both SSTs:

$$P(RT < t | S_{12}) \leq P(RT < t | S_1) + P(RT < t | S_2)$$

where S_i denotes channel i for a SST, and S_{12} denotes a RST with a target presented in both channel 1 and channel 2.

Given that a race model cannot account for observed redundancy gains, if the RMI is violated, the alternative proposed by Miller (17, 18) are co-activation models (Figure 2). In a co-activation model, the signals in the two channels of a RST do not race against each other; rather, they are integrated/summed before triggering a response. Hence activity accumulates faster and to a higher level on RSTs compared to SSTs, resulting in redundancy gains that cannot be accounted for by parallel-race models.

The redundant target paradigm has been applied in a variety of areas, such as bi-modal (e.g., visual and auditory) processing (17, 30, 31), divided attention (between two locations: 20, 21, 32), and neuro-psychological research (26, 33, 34).

The studies relevant to the present review used a redundant-signals paradigm in visual pop-out search. In pop-out search, targets possess a unique feature relative to the non-targets/distracters – for example, a horizontal (target) bar is presented among vertical (distracter) bars or a red bar appears among green bars. As already stated, several cognitive and neuronal models assume that feature contrast signals are computed in parallel in different dimensions (e.g., orientation and color). In terms of the notation introduced above in the abstract description of the redundant-signals paradigm, feature contrast in one




Figure 3. Example displays with targets defined in a single dimension or redundantly in two dimensions. (a) presents a color target, (b) an orientation target, and (c) a target defined redundantly by orientation and color contrast. (d) presents a dual-target display, with two pop-out targets defined in separate dimensions, which are in close spatial proximity.

dimension (e.g., orientation) can be denoted as a single stimulus S_1 , and feature contrast in another dimension (e.g., color) as single stimulus S_2 . Redundant targets are then the combination of both S_1 and S_2 , such as a red horizontal among green vertical bars. If such a redundant pop-out target is presented, feature contrast is processed initially in both dimensions in parallel. The RMI provides then a tool to decide between different architectures of how responses may be elicited: If feature contrast in both dimensions can independently elicit a detection response, then redundancy gains are expected without violations of the RMI. However, if the RMI is violated, a parallel architecture of feature contrast signals independently racing to trigger a detection response can be excluded, and an integrative/co-active model can account for the redundancy gains.

2.2. The question of architecture: parallel, co-active, or serial?

Thus, the initial question concerns the fundamental processing architecture: how are the separate dimensional feature contrast signals processed prior to eliciting a response? Does the initial parallel processing continue until response selection? That is, can both dimensional signals elicit a response in parallel? Or are dimensional feature contrast signals integrated in some kind of overall-saliency (or ‘master’) map, in which the signals from separate dimensions are summed for each location (10, 13)? Or are dimensional feature contrast signals processed in serial for response decision, either in a self-terminating search (i.e., the search stops as soon as a target is found in one dimension; e.g., 9, 35), or in an exhaustive manner (i.e., all dimensions are serially checked even if a target is found in one; e.g., 36)? For divided-attention paradigms, in which a target may be presented at two possible locations, Mordkoff and Yantis (20, 21) have proposed an interactive-race model in which there may be cross-talk between the separate channels prior to response selection.

In summary, the processing architecture of the initially separate dimensional contrast signals could be parallel (with independent channels or with cross-talk between the channels), it could be serial (either self-terminating or exhaustive), or it could be integrative/co-

active (as with models that assume an overall-saliency map).

Krummenacher, Müller, and Heller (22) used the redundant-signals paradigm in visual pop-out search to address these questions. In their Experiment 1, the displays consisted of an array of distracters (green vertical bars) with a singleton target (presented on 50% of the trials) that could be either color-defined (red vertical bar), orientation-defined (green 45°-tilted bar), or redundantly defined (i.e., differ from the distracters in both color and orientation, red 45°-tilted bar, see Figure 3 for an illustration of the stimuli).

When comparing reaction times to targets defined in one dimension only to redundantly defined targets, they found a significant redundancy gain of about 20 ms. Furthermore, when comparing the cumulative density functions (CDFs) of reaction times for redundant targets to the sum of the CDFs for both single-dimension targets, violations of the RMI were observed. These violations exclude serial self-terminating and parallel self-terminating or parallel exhaustive models of visual processing in pop-out search.

The more theoretical alternatives of serial exhaustive and parallel interactive models have been examined by Zehetleitner, Krummenacher, and Müller (37), combining a redundant-signals paradigm with a double-factorial design (36). The redundant signal-paradigm permits the RMI to be tested, as a means to exclude parallel-race models. However, even if violated, the RMI cannot exclude serial exhaustive (38) or interactive-race models (20, 21). Townsend and Nozawa demonstrated that, by using a factorial design together with a redundant-signals paradigm, it is possible to test the model architectures (serial, parallel, or co-active) and stopping rules (self-terminating or exhaustive). The double factorial design is derived from Sternberg’s (39) additive-factors method. Applied to a visual-search paradigm, it combines the presentation of a pop-out target defined in two possible dimensions with the factorial manipulation of a second variable, such as feature contrast. So, observers are presented with an ‘absent’ display, with single dimension displays (in which the target differs from distracters in one dimension), or with redundant-dimension

displays (in which the target differs from distracters simultaneously in two dimensions). The feature contrast can be manipulated by varying the difference between targets and distracters. Townsend and Nozawa proved that analyzing the interaction between feature contrasts in both dimensions of redundant targets can differentiate between different architectures and stopping rules. There are four possible types of redundant target in the double factorial paradigm (2 dimensions x 2 levels of feature contrast). For instance, with orientation and luminance as the critical dimensions, orientation targets may differ from distracters by an angle of 6° (low feature contrast) or 45° (high feature contrast), while luminance targets may be either dim (low feature contrast) or bright (high feature contrast). Thus, the four different types of redundant targets are: (i) tilted 45° and bright, (ii) tilted 45° and dim, (iii) tilted 6° and bright, and (iv) tilted 6° and dim. If the two factors are independent, they should have additive (non-interacting) effects on the processing speed of redundant targets. Sub-additivity occurs if lowering the feature contrast in one dimension has a smaller slowing effect on RTs when the feature in the other dimension is already of low feature contrast. If lowering feature contrast in one dimension has a larger effect when the feature in the second dimension is of low contrast, then super-additivity is said to occur. Townsend and Nozawa (36) proved that (under general conditions) parallel-race models predict super-additivity in the mean interaction contrast, parallel exhaustive models predict sub-additivity, and both exhaustive and self-terminating models predict simple additivity when looking at the interaction of feature contrast for both dimensional components of redundant targets.¹

Using the paradigm and stimuli described above, Zehetleitner *et al.* (37) found a super-additive interaction of the intensity levels in redundant targets, hence excluding serial models of both stopping rules (self-terminating and exhaustive).

Interactive parallel-race models that can lead to violations of the RMI assume cross-talk between the parallel channels, that is, exchange of information between the two channels before the decision unit is reached. This information is helpful only if there are correlations/contingencies between the signals in the two channels. For instance, if the color channel identifies the presence of a pop-out target, this information could be made available to the orientation channel via cross-talk. If the presence of feature contrast defined in the color dimension is uncorrelated with the presence of feature contrast in the orientation dimension, this information is not beneficial for processing in the other channel. For example, if the probability of color feature contrast being present is greater given the presence, rather than the absence, of orientation contrast (i.e., if the probability of a redundant target is greater than the probability of a simple color target), then exchange of information about the presence of feature contrast favors redundant targets – because information about the presence of feature contrast in one dimension (e.g., color) makes the presence of feature contrast in the other dimension (e.g., orientation) more probable than its absence. Under these circumstances, detection of redundant targets can be expedited as

compared to single targets, which can lead to violations of the RMI that are not due to co-activation (20, 21).

It is possible to design an experiment in which the correlations between channels do not favor redundant targets, if more than one type of ‘absent’ stimulus is introduced – which is, however, not possible logically in visual pop-out search. Thus, in order to test whether an interactive race-model could account for observed violations of the RMI (22, 23), Zehetleitner *et al.* (37) manipulated the amount of information that can aid detection of redundant targets via channel cross-talk, assuming that this manipulation would modulate redundancy gains and the degree to which the RMI is violated. Zehetleitner *et al.* varied the ratio of ‘present’ compared to ‘absent’ displays (50:50 vs. 75:25), crossed with variation of the proportion of single targets compared to redundant targets (50:50 vs. 66:33). Each of these combinations of ratios leads to contingencies that differ in the strength of benefit for redundant over single targets. The interactive-race model predicts redundancy gains and violations of the RMI to be the greater the stronger these contingencies are. However, at variance with this, Zehetleitner *et al.* failed to find any variation of redundancy gains or of the magnitude of the RMI violations when manipulating inter-channel contingencies.

In summary, the available evidence most strongly supports co-active/integrative models of visual processing of feature contrast signals. Independent parallel models can be excluded because redundant pop-out targets lead to violations of the RMI. Serial exhaustive models, which in theory, can also lead to violations of the RMI, can be excluded because of an over-additive interaction for redundant targets. Interactive-parallel models are unlikely, because manipulating the amount of information that would facilitate processing of redundant targets via channel-cross-talk did not alter redundancy gain or the magnitude of RMI violations.

2.3. Is integration spatially selective?

Both neuronal and cognitive models of visual processing agree that the initial parallel computation of feature contrast is topographically organized. The evidence of Krummenacher *et al.* (22) as well as of Zehetleitner *et al.* (37) strongly support co-activation models, consistent with the idea of an overall-saliency (master) map into which feature contrast signals are summed. The models of both Wolfe (10, see also 1) and Itti and Koch (13) assume that the integration stage is topographically organized, that is, the integration is spatially specific (Figure 4). In models of this type, redundant signals can only be integrated if they are in close spatial proximity.

An alternative model would assume that dimensional signals are spatially pooled before the integration stage. An example model of this type is the original feature integration theory (FIT; 8, 40), which assumed that there are dimensional pooling units that signal presence of feature contrast in one dimension, irrespective of its precise location, as presented in Figure 5. That is, crucial with

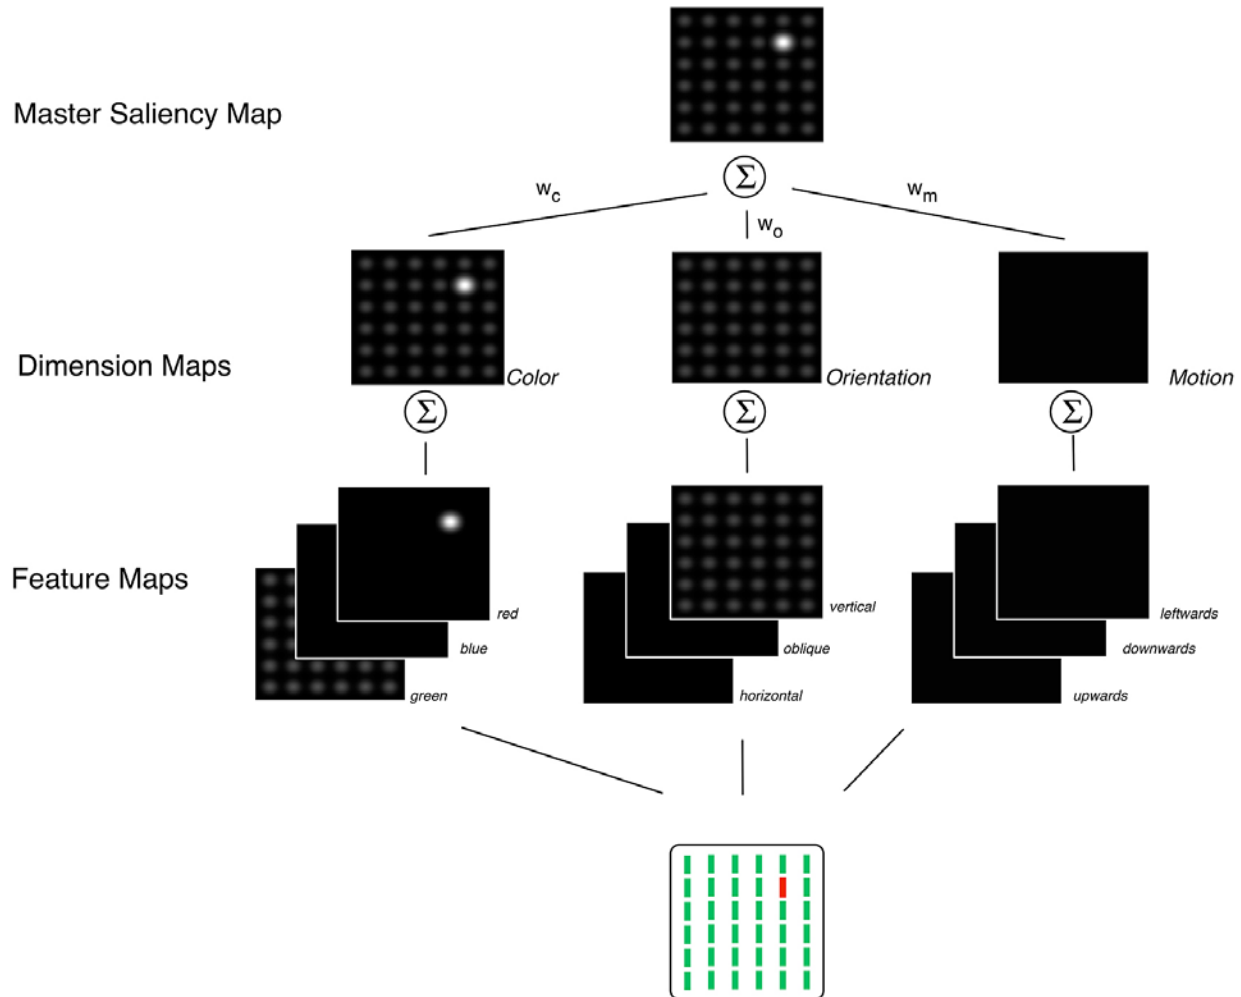


Figure 4. Example saliency summation model (e.g., 1, 10, 13). The display is first analyzed by spatio-topically organized feature analyzers for different colors, orientations, motion directions, etc. Each map is a topographical representation of the display, with black representing no activity and white strongest activity. Feature maps are summed into dimension maps, which are then summed into the master saliency map. The contribution of each dimension map to the activity of the master saliency map can be modulated by dimension weights w_c , w_o , and w_m (for the color, orientation, and motion dimensions, respectively).

regard to the question of spatial specificity, these dimensional units convey information only about signal presence in a dimension, but not spatial information about where the signal originates from. This notion of dimensional signals, which are spatially unspecific has been revived recently by Chan and Hayward (41) and Mortier *et al.* (42). This notion predicts, that processing of redundant pop-out target signals, which are defined by feature contrast in two dimensions, is independent of the spatial arrangement of the individual signals (for example, a single target redundantly defined in two dimensions compared to dual targets defined in different dimensions).

Another model makes exactly the opposite prediction, namely, that integration of dimensionally redundant target signals happens only if both dimensional signals originate from the same location: the dimension action (DA) model of Cohen and Feintuch (43), which is based on the cross-dimensional response selection model of

Cohen and Magen (44, see also 45, 46). The DA model assumes that there are dimension-specific feature analyzer units as well as multiple response selection units, one per visual dimension (43, p. 589). While the dimensional response selection units compute responses in parallel, the response decision of only one such unit can be transferred to an executive (working memory) stage which mediates overt reactions. Thus, if targets defined in multiple dimensions are present in the display, their critical features will be analyzed in parallel in a dimension-specific manner. Likewise, the response for each target feature will be selected separately in parallel. However, which specific response will be transferred to the reaction stage depends on focal attention: responses will be processed further only for signal locations to which spatial attention is allocated. If the two dimensional signals originate from the same position (i.e., with a single dimensionally redundant target),




Figure 5. Example model in which presence of feature contrast in one dimension is represented in a non-spatial fashion (e.g., 8, 40). The display is analyzed in terms of color, orientation, and motion contrast. The possible activity is represented by the light grey boxes. Large activity is represented by a tall dark grey bar, low activity by a small dark grey bar.

response units from both dimensions will be activated and their signals transferred on to working memory in a co-active manner, provided that focal attention is directed to the target location (43, pp. 591-592). Regarding the spatial range of cross-dimensional signal integration, this model predicts that co-activation can take place only if the two component signals of a dimensionally redundant target are presented at one location.

In contrast, guided-search-type models (10, 13) assume that integration of redundant dimensional signals happens on the overall-saliency (master) map in a spatially specific manner. Thus, these models predict – in contrast to models with a FIT-type architecture – that feature contrast signals from two dimensions that are too far apart spatially will not be integrated, because summation is spatially specific. At the same time, they predict – in contrast to the DA model – that dimensionally redundant signals that are spatially not too far apart may still be integrated, albeit to a lesser extent.

To examine these alternatives, Krummenacher *et al.* (23) examined the role of spatial information in a variant of the redundant-signals paradigm for pop-out targets. Targets differed from distracters in either color or orientation. Redundancy could be of two types, either in one target (which differed from the distracters in both dimensions), or in two targets, one defined in orientation and the other in color. By introducing dual (redundant) targets, spatial distance between the redundant signals could be manipulated. Krummenacher *et al.* found that redundancy gains decreased with increasing spatial distance between the dual targets, as predicted by models that assume spatially specific integration. Violations of the RMI occurred only for redundant target signals separated by one to two units of distance (1.5°–3.6° of visual), but not for redundant targets separated by three units of distance (more than 3.5° of visual angle). In addition, the range of reaction

times for which the RMI was violated was smaller for the medium than for the short distance.

Two conclusions can be drawn from these findings: (i) integration of redundant signals is spatially specific, and (ii) the strength of integration decreases with increasing distance between the two redundant signals. If the two targets are too far apart spatially, no integration occurs at all. This is at odds both with theories that assume spatial (intra-dimensional) pooling of the separate signals prior to the integration stage (e.g., FIT) and with models that assume integration of redundant signals to be possible only at one position, the locus of focal attention (DA model). In contrast, it is consistent with models that assume a topographically organized overall-saliency (master) map: integration is spatially scaled, so that redundant signals that originate from the same location benefit maximally in terms of integration strength, and strength of integration decreases with increasing spatial distance.

If integration of redundant signals requires spatial proximity, does integration also require spatial (focal) attention? The answer provided by the DA model is ‘yes’, because dimensional response units require deployment of focal attention for co-active processing (i.e., the integration of redundant dimensional signals is assumed to be a post-selective process). To address this question, Krummenacher *et al.* (23) combined the redundant-signals paradigm with a spatial-cueing procedure. At the start of a trial, participants were informed of the display quadrant, a circumscribed region (with the maximum center-to-center distances varying between 2.05° horizontally and 2.9° vertically), highly likely ($p=.79$) to contain a target by a central-arrow indicator. Observers did make use of the cues, as evidenced by overall faster reaction times for (valid) targets that appeared at the cued location, compared to (invalid) targets at uncued locations. Importantly, Krummenacher *et al.* found violations of the RMI to be independent of the locus

of attention. The RMI was violated both for redundant targets that were placed within the cued quadrant and for targets placed within an uncued quadrant. This finding is at variance with the assumption of post-selective signal integration made by the DA model, but consistent with models that assume an overall saliency map that guides spatial-attentional selection. In these models, integration of redundant signals is assumed to be a preattentive process, that is, independent of the locus of focal attention – a view advocated by Krummenacher *et al.* (23).

In summary, co-activation of redundant pop-out target signals is spatially specific (or spatially scaled): the two components of redundant targets must be in spatial proximity in order to be integrated. This is consistent with models that assume a center-surround computation of feature contrast (13). Although spatially specific, spatial attention is not necessary for integration: integration of redundant pop-out targets is a preattentive process in the sense that it affects build-up of activation on the salience master map, before this activation is used to guide deployment of spatial attention (10, 13).

2.4. Is integration feature-based or dimension-based?

The redundant-signals paradigm applied to pop-out search provides support in favor of a co-active/integrative architecture of processing of dimensionally different feature contrast signals and against serial and, respectively, parallel (independent or interactive) models. It also provides evidence that the integration of feature contrast signals is spatially specific and preattentive. Cognitive and neuronal models that are based on an overall-saliency map (1, 10, 13) additionally assume that feature contrast signals are dimensionally pooled before being integrated into the overall-saliency map: as illustrated in Figure 4, according to these models, feature contrast is computed in parallel for each feature (e.g., red, green, vertical, right-tilted, bright, dim, etc.). Before being pooled into a master saliency map, feature contrast signals that stem from the same dimension (see 14) are first summed into a dimension-specific map. Dimension-specific signals may be weighted prior to being transferred to the master saliency map (1). If the weight is set higher for a particular dimension, activity from this dimension has an earlier and/or greater impact on the activity on the master saliency map.

Explicit computational models (13) would also predict integration to be possible only between different dimensions, not between different features within the same dimension – because the signals on the dimensional maps are assumed to be normalized before further processing. If a redundant target consists of two pop-out signals defined in the same dimension (e.g., a red and a blue target among green distracters), these would initially produce larger signals on the color-specific (dimension) map, due to the summation of signals from the feature maps (red and blue), compared to single targets. That is, the activity on the color map for two targets defined by separate feature contrasts (e.g., red and blue vs. green) would be higher than the activity generated by a target defined by only one feature, if the two targets are in close spatial proximity. In more

detail, the activity produced by a target on any map, rather than being confined to a single point, is spread out (e.g., in the way of a two-dimensional Gaussian distribution). When there are two targets in close proximity, they produce spatially overlapping activations on different feature maps. When such signals are pooled (i.e., when their overlapping activations are summed), there would be co-activation of dimension-specific units by separate features within a given dimension, analogous to the co-activation of master saliency units by dual pop-out targets defined in separate dimensions (23). But in order to assure that all dimensions contribute equally to overall-saliency, dimensional signals are normalized (to values between 0 and 1; 13, p. 1493) before being summed into the master map. Due to this normalization, enhanced activity on dimension-specific maps generated by dual pop-out targets in the same dimension is not propagated to the overall-saliency map. In summary, the critical difference between dimension-specific maps and the master saliency map is that activity is normalized on the former (a process by which redundancy gains are lost), but not the latter (which permits for redundancy gains to have an effect).

To examine this assumption, Krummenacher *et al.* (23) analyzed the processing of redundant target signals that were spatially separate, but in close proximity, for two conditions: within-dimension and cross-dimension. In the within-dimension condition, the two signals on RSTs were defined both in either the color (red and blue) or the orientation dimension (tilted to the left and to the right). In the cross-dimension condition, one of the targets was defined in the orientation, the other in the color dimension. Violations of the RMI were observed only when the two targets were defined in separate dimensions, but not when they were defined within the same dimension. This finding strongly supports models that assume dimensional pooling with some kind of normalization before contrast signals are fed into a master saliency map (1, 13).

2.5. Weighting or priming?

Maljkovic and Nakayama (47) observed that target detection on trial n was affected by the target on trial $n-1$. In their experiments, a pop-out target was present on each trial. For example, in the color condition, the target could be either red or green, and the distracters were green or red, respectively. If the target definition stayed the same on successive trials, search performance was faster than when the target definition changed from one trial to the next. To rule out that this intertrial effect is due to top-down processes, Maljkovic and Nakayama varied the predictability of the target/distracter feature swap. They found that, even when the sequence was made perfectly predictable (i.e., the target definition changed regularly every two trials in AABBA...manner) and observers were informed about this rule, change of the target/distracter features still produced substantial reaction time costs. Hence, Maljkovic and Nakayama interpreted this effect in terms of the passive (top-down impenetrable) ‘priming of pop-out’.

A similar effect of the previous target definition on search performance on a given trial was described by

Müller *et al.* (1). Their observers had to discern the presence (vs. the absence) of an orientation pop-out target in either a within-dimension or a cross-dimension condition. In within-dimension search, targets, if present, were always defined in the orientation dimension. In cross-dimension search, targets could be defined in the orientation, the size, or the color dimension. Comparison of the reaction times to the right-tilted orientation target presented in both conditions revealed costs of about 60 ms for cross-dimension as compared to within-dimension search. Similar to Maljkovic and Nakayama (47), Müller *et al.* (1) took this result as evidence for a bottom-up modulation of search performance. But, in contrast to Maljkovic and Nakayama who swapped target and distractor features, Müller *et al.* used dimensionally variable targets (with a constant distractor background), emphasizing the importance of dimensional changes for producing search reaction time costs.

Subsequently, Found and Müller (11) examined the relative contributions of the dimensional and featural effects on change costs across consecutive trials. Observers had to discern the presence of a target among green vertical distracter bars. The target, if present, could be either color-defined (red or blue) or orientation-defined (tilted to the left or right of vertical). The target type was varied randomly from trial to trial. Hence, on two consecutive trials, the target could either be repeated (i.e., be defined in same dimension by the same feature, e.g., a red target following a red target), it could change feature while remaining defined in the same dimension (e.g., a red target following a blue target), or it could be defined in a different dimension (e.g., a red target following a right-tilted target). Found and Müller found substantial reaction time costs when the target changed dimension, and only slight costs (if any) when the target changed feature within a repeated dimension – relative to the condition in which the target was unchanged. They concluded that the slowing of reaction times was (mainly) due to dimension changes rather than feature changes. In their ‘dimension weighting’ account, Müller and colleagues (1, 11) assumed that dimensional signals can be modulated by dimension-specific weights, prior to integration into the overall-saliency (master) map (see weights w_c , w_o , and w_m in Figure 4). These weights can be affected in two ways: (i) in bottom-up fashion by the trial history, and (ii) in top-down fashion by the observer’s intentions (48, 49).

It is possible to examine changes of dimensional weights dependent on trial history in more detail. The dimension weighting account assumes that the target’s feature contrast signal(s) on trial $n-1$ gives rise to weight adaptations, which then modulate saliency computations on the following trial n . On the assumption that preattentive feature contrast signals, rather than attentionally analyzed conjunctions of the target features, are determining changes of dimensional weights, there are in principle two fundamentally different ways in which such a weight-based signal modulation may be implemented: ‘priming’ or ‘weighting’. On the priming account, presenting a target in one dimension on trial $n-1$ increases the weight for that dimension. The increased weight leads to a faster build-up

and a larger final signal on the overall-saliency (master) map when the target on trial n is defined in the same dimension – which becomes evident in a reaction time benefit for repeated-dimension targets. Nearly the same is true on a weighting account, with one important difference: Similar to the priming account, presenting a target in a specific dimension on trial $n-1$ increases the weight for signals defined in this dimension (expediting reaction times to targets defined in the same dimension on the subsequent trial n). But, in contrast to the priming account, the weighing hypothesis states that increasing the weight for one dimension entails decreasing the weights for other dimensions. If targets occur in only two possible dimensions in an experiment, other, irrelevant dimensions would never receive a weight increase (as they are never reinforced by a target event). However, they would receive an inhibitory signal (by the target events in the relevant dimensions), driving the dimensional weights towards the minimum. That is, while the priming account assumes dimensional weights to be an unlimited resource, the weighting account assumes that the total weight is limited such that the weight of one dimension cannot be increased without decreasing the weight of other dimensions.

Again, the redundant-signals paradigm in pop-out search can help to decide between the two alternatives: weighting versus priming, since the two accounts differ in their predictions about the effect of a dimensionally redundant – such as a color- plus orientation-defined target – signal on trial $n-1$. On the assumption of priming, a redundant target increases the weights for both dimensions: the color component of the redundant target increases the weight for the color dimension (just like a singly defined color target), and the orientation component increases the weight for the orientation dimension (just like a singly defined orientation target). Hence, the priming account predicts performance for a singly defined color target on trial n to be the same, regardless of whether the preceding – trial $n-1$ – target was color-defined or redundantly defined, because the weight for the color dimension is changed by both types of target in the same way. By comparison, responses to a color target on trial n are predicted to be slower if the target $n-1$ was defined by orientation. The same would hold for a singly defined orientation target on trial n , responses to which would be independent of whether target $n-1$ was orientation-defined or redundantly defined, while responses would be slower when the orientation target is preceded by a color target.

The weighting account makes a very different prediction, based on the assumption that the weights for a given dimension are dynamically adjusted by competitive interactions that strengthen the weight for a given target-defining dimension by withdrawing weight from other dimensions. That is, the weight for a given dimension can only be increased by decreasing the weight for one or more of the other dimensions, implementing a limited-capacity weight resource. Thus, on a redundant-target trial, the color component strengthens the weight for color, while simultaneously reducing that for orientation. Concomitantly, the orientation component of the redundant target would increase the weight for orientation and

decrease that for color. That is, on redundant-signals trials, the weight of each (relevant) dimension would be simultaneously strengthened and weakened. Consequently, responses to a color target on trial n are fastest if it follows a color-defined target, intermediate if it follows a redundantly defined target, and slowest if it follows an orientation-defined target. The reason is that, if target $n-1$ is color-defined, the weight for the color dimension receives only a facilitatory input; if it is redundantly defined, the color weight receives both a facilitatory (from the color component of the redundant target) and an inhibitory input (from the orientation component); and if target $n-1$ is orientation-defined, the color weight receives only an inhibitory input. (Analogous predictions hold for an orientation target on trial n following an orientation-defined, a redundantly defined, or a color-defined target on trial $n-1$.)

Thus, the two different models of how dimensional weights are changed – priming versus weighting – lead to differential predictions regarding the effect of a redundantly defined target on search performance for the next target. The data of Krummenacher *et al.* (22) permit these predictions to be tested, by analyzing the effects of orientation, color, and redundant targets on trial $n-1$ on reaction times to targets on trial n . They found that targets defined in a given dimension were detected fastest when the preceding target was defined in the same dimension, and slowest when it was defined in a different dimension. This is consistent with both the priming and the weighting account. However, search performance for a singly defined target (whether by color or by orientation) preceded by a redundant target was in-between performance for same-dimension and different-dimension targets. This pattern is inconsistent with the priming account (which predicts a redundant target on trial $n-1$ to lead to the same reaction time performance as a same-dimension target), but expected on the weighting account.

Thus, dimensional weights do not behave in terms of an unlimited resource that can be increased for each dimension without any constraints. Rather, dimensional weights may be conceived of as a limited resource, in the sense that it is impossible to increase the weight of one dimension without decreasing the weight of other dimensions, as originally proposed by Müller and colleagues (1, 11).

2.6. Implementation of saliency maps and dimensional weighting in the brain

Although an integral part of many cognitive and computational theories, the neural implementation of an overall-saliency map is not yet fully clear. The properties of such a neural map have to include: (i) topographical organization, (ii) featureless representation of stimuli (locations), and (iii) strength of activity related to strength of local center-surround contrast. Several structures are currently hypothesized to provide an implementation of an overall-saliency map: the pulvinar (e.g. 50), the lateral intraparietal area (e.g. 51), and the frontal eye fields (FEF; e.g. 52, 53) – with the FEF being a particularly promising

structure. Quite likely, though, there is not only a single implementation of a saliency map in the primate brain, but rather a network of multiple, interacting areas: the oculomotor network (54). In the following section, we will focus on evidence relating to the FEF, which has been gained using visual-search paradigms.

The FEF fulfills all three of the above criteria. It is topographically organized such that neighboring neurons represent neighboring points in retinotopic coordinates (55, 56 for monkeys; 57 for humans). The featureless and feature-contrast dependent response characteristics of FEF neurons have been demonstrated by Sato *et al.* (58). They manipulated two aspects of a visual search task in a monkey single-cell study that both lead to increased reaction times: (i) they varied the saliency of the target, that is, the similarity between the target and distracters, and (ii) they introduced response interference by infrequently changing the location of the target. The task was to saccade to an odd-ball target defined by either color or motion contrast. Search difficulty was manipulated by varying the color similarity of the target to the distractors and, respectively, the proportion of dot stimuli moved coherently in one direction within a pattern of randomly moving dots. In the response interference condition, the target and one distracter changed locations after initial presentation of the search array. The monkey had to cancel the initial saccade and shift gaze to the new target location. Although both manipulations affected the latency and variability of reaction times, only the perceptual manipulation had an influence on the time taken by visually responsive neurons in the FEF to select the target. Thus, activity of visually responsive FEF neurons reflects the strength of feature contrast, and responses are triggered by feature contrast whether defined by motion or color differences.

Given the evidence for FEF neurons signaling overall-saliency, two new questions arise within the present context: can FEF neuronal activity explain redundancy gains and intertrial change effects in visual pop-out search? Concerning the former, cognitive and computational models of visual search (1, 10, 13) assume that redundant targets give rise to both a faster and a higher build-up of activation on the master saliency map compared to single targets. Accordingly, if there is a neural implementation of the master saliency map in the FEF, visually responsive FEF neurons should be able to select targets faster when they are redundantly defined by feature contrast in two dimensions, compared to being defined by feature contrast in one dimension only. This prediction still needs to be tested in single-cell studies.

Relating to the second question, the DWA assumes that repetition of the target-defining dimension (e.g., color) across trials leads to an increase of the weight for this dimension, and a decrease of weight for other dimensions. Thus, feature contrast signals from the repeated dimension (e.g., color) would have an earlier and greater impact on the master saliency map, compared to signals from a non-repeated dimension. Thus, if visually responsive FEF neurons signal saliency, they should show

an earlier, enhanced response to a pop-out target defined in a repeated (rather than changed) dimension. Indeed, Bichot and Schall (53) observed a similar pattern using a pop-out search paradigm adapted from Maljkovic and Nakayama (47), in which the identity of targets and distractors (rather than the target-defining dimension) could change from trial to trial. In Bichot and Schall's task, monkeys were trained to saccade to an odd-one-out target defined by either color or shape, with the target and distractor features varying randomly across trials. Bichot and Schall found that activity of visually responsive FEF neurons to targets and, respectively, distractors separated faster if the target and distractor features were repeated, rather than changed, from the previous trial. This suggests that a similar pattern of FEF neuronal responses would also be observed for dimension repetitions versus changes across trials; however, this prediction still requires explicit testing.

Concerning the brain mechanisms responsible for controlling the assignment of dimensional weights, data from fMRI studies suggest that these comprise a fronto-posterior network. Pollmann *et al.* (59) found that changes (vs. repetitions) in the dimension defining a pop-out target lead to increased activation in the left frontopolar cortex and inferior-frontal gyri, as well as high-level visual processing areas in parietal and temporal cortex, and dorsal occipital visual areas. Follow-up studies (60-62) support the view that the mechanisms responsible for controlling the change of dimensional weights involve fronto-polar cortex and that the effect of changes in dimensional weights is mediated via feedback connections to the extrastriate visual areas that process the features of the new target dimension. For example, in the fMRI study of Pollmann *et al.* (60), the target-defining dimensions were either color or motion direction. BOLD activity for trials with targets successively defined in the same dimension was tonically increased in posterior fusiform gyrus (which contains human area V4) for repeated color targets and in lateral occipital cortex (which contains the hMT + complex) for repeated motion targets. This supports the view that dimension-specific feature contrast signals can be weighted before being summed onto a master saliency map, where signals from a weighted dimension lead to a faster build-up of activation (for further details see 62, this issue).

While the framework discussed thus far assumes that the master saliency map is a relatively high-level representation, an alternative – low-level – representation was recently proposed by Li (64). According to Li, V1 computes a saliency map that is not based on the summation of feature-contrast signals (summation models, as proposed by others) (1, 8, 10, 13). Instead, the saliency of a location is determined by the firing rate of the most active V1 cells responding to the feature singleton (maximum model). The firing rates of V1 do not only depend on input strength, but also on the 'context' (65). Thus, for example, in a display with a horizontal bar surrounded by vertical bars, the cells responding to vertical orientation would be subject to iso-orientation suppression, whereas the cells responding to horizontal orientation would not be suppressed. As a result, the feature singleton target would lead to more active V1 cells tuned to

horizontal orientation and to less active V1 cells tuned to vertical orientation. The most salient location is then simply signaled by the most active V1 cells (maximum selection rule). – This notion of a V1 saliency map resembles that of an overall-saliency map as conceived in summation accounts: saliency is signaled in a topographical and featureless manner, with saliency strength being related to strength of local center-surround contrast.

This alternative model is also relevant to the present question at issue, namely, how redundantly defined pop-out targets are processed. In summation models, redundant targets are processed faster because activity on the master map is generated by signals originating from two dimensions simultaneously. In contrast, in Li's maximum model, all dimensions contribute independently of each other to overall-saliency. If there were only dimension-specific cells in V1 (e.g., cells tuned to either color or orientation), processing would resemble a parallel horse race model – which is, however, excluded by established violations of the RMI (22). But V1 contains also cells that respond to features of more than one dimension (e.g., tuned for red vertical bars, or for bright bars that are moving upwards). Such cells have been demonstrated by, for example, Leventhal *et al.* (66) who analyzed responses of single V1 neurons in V1 of anaesthetized, paralyzed monkeys: they reported most V1 cells to be responsive simultaneously to color, orientation, and motion. Most importantly, there was no negative correlation between color- and orientation-sensitive cells, as would be the case if each cell were tuned to a feature of one dimension exclusively. Following Leventhal *et al.*, there have been more frequent reports of cells conjunctively tuned to features of two dimensions in V1 and V2 (67-69).

The existence of such conjunction cells is relevant to the detection of dimensionally redundant pop-out targets. For instance, for a red vertical target among green horizontal distractors, there are three types of cell in V1 which are most active at the target location: color cells tuned to red, orientation cells tuned to vertical, and conjunction cells tuned to red and horizontal. Based on the V1 maximum model of saliency, Koene and Zhaoping (70) contended that saliency is larger for dimensionally redundant, relative to singly defined, pop-out targets if there exist conjunction cells in V1 for the respective combination of dimensions. Based on neuronal evidence (71-73), they argued that there are no conjunction cells in V1 for the combination of color and motion (CM), whereas there are such cells for the combinations of color and orientation (CO) and motion and orientation (MO). They therefore hypothesized that the RMI would be violated only for the combinations of CO and MO, but not the combination CM. This dissociation was supported by their experiment, in which participants had to respond to the location (i.e., left or right half of the display) of a pop-out target that was defined in either one dimension (of color, motion, or orientation) or redundantly in two dimensions (CO, MO, or CM): there were reliable violations of the RMI only for CO and MO targets, but not for CM targets. Koene and Zhaoping took this pattern as support for their

V1 maximum model of saliency, based on the non-existence of cells in V1 cells conjunctively tuned to color and motion.

However, at variance with Koene and Zhaoping's (70) null-result, Krummenacher and Müller (74) found pop-out targets redundantly defined by a combination of color and motion to significantly violate the RMI (see also 24, who reported redundancy gains and violations of the RMI for targets defined by color and motion). Thus, given the non-existence of color-motion cells in V1 (for which the evidence is actually mixed: while 66, reported finding such cells, others, such as 71, failed to do so), the findings of Krummenacher and Müller as well as of Katzner *et al.* would argue against the V1 maximum model of saliency advocated by Li (64) and Koene and Zhaoping (70). In summary, it is not yet possible to unequivocally decide between the summation saliency (10, 13) and the V1 maximum models (64), but if color-motion cells are indeed non-existent in V1, as assumed by Koene and Zhaoping, the results of Krummenacher and Müller provide further good grounds to argue in favor of the summation saliency model.

2.7. Conclusion

In the present review, we have summarized research on several critical questions concerning the nature of early visual processing and we have shown how the redundant-signals paradigm in visual pop-out search provides a powerful tool for answering these questions. Applied to pop-out search, this paradigm yields a number of dependent measures, including mean redundancy gains, violations of the RMI, and effects of redundant targets on cross-dimensional intertrial transitions. The findings strongly support co-activation models that assume summation of feature contrast signals in a master saliency map (e.g., Guided Search model, the dimension-weighting account, and the model of 13). Exclusion of parallel- and interactive-race models, as well as of serial models supports summation models in general. Several other findings, such as spatial specificity of integration, the preattentive nature of integration, and the dimensional organization of feature contrast signals are supported by various studies that have employed the redundant-signals paradigm as a tool. Especially the dimension weighting account receives further support regarding the limited-resource nature of dimensional weights. Instead of a priming mechanism that could increase weights for several dimensions independently, a weighting mechanism (as proposed by 1) seems to determine stimulus-driven changes in the dimensional weight set: increasing the weight for one dimension goes along with decreasing the weights for one or several other dimensions – by a competitive interaction that implies a limit to the total weight available to be allocated to the various dimensions.

Issues for further research include the effect of redundant targets, as well as that of cross-trial changes in the target-defining dimension, on FEF neuronal activity: the DWA predicts that redundantly defined targets as well as targets defined in a repeated dimension would lead to expedited discrimination between targets and distracters in

FEF neurons, compared to targets defined in a single dimension, and targets defined in a changed dimension relative to the target on the previous trial. Also, the alternative to summation models of saliency, namely: Li's V1 maximum model of saliency (64) requires further behavioral and neuro-physiological research to permit an unequivocal decision to be made between the alternative models.

3. ACKNOWLEDGMENTS

This research was supported by grants from the DFG (excellence cluster EXC142 'Cognition for Technical Systems – CoTeSys', research group FOR480) and the Swiss National Science Foundation (PP001 – 110543/1). Correspondence concerning this article should be addressed to Michael Zehetleitner.

4. REFERENCES

1. H.J. Müller, D. Heller, J. Ziegler: Visual search for singleton feature targets within and across feature dimensions. *Perception & Psychophysics* 57, 1–17 (1995)
2. M.S. Livingstone, D.H. Hubel: Anatomy and physiology of a color system in the primate visual cortex. *J Neurosci* 4, 309–356 (1984)
3. M.S. Livingstone, D.H. Hubel: Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *J Neurosci* 7, 3416–3468 (1987, Nov)
4. M. S. Livingstone, D.H. Hubel: Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science* 240, 740–749 (1988, May)
5. D.H. Hubel, M.S. Livingstone: Complex-unoriented cells in a subregion of primate area 18. *Nature* 315, 325–327 (1985)
6. D.H. Hubel, M.S. Livingstone: Segregation of form, color, and stereopsis in primate area 18. *J Neurosci* 7, 3378–3415 (1987)
7. L.C. Sincich, J.C. Horton: The circuitry of V1 and V2: integration of color, form, and motion. *Annu Rev Neurosci* 28, 303–326 (2005)
8. A. Treisman, G. Gelade: A feature-integration theory of vision. *Cognitive Psychology* 12, 97–136 (1980)
9. A. Treisman: Features and objects: the fourteenth bartlett memorial lecture. *Quarterly Journal of Experimental Psychology* 40, 201–237 (1988, May)
10. J.M. Wolfe: Guided search 2.0 -a revised model of visual search. *Psychonomic Bull Rev* 1, 202–238 (1994)
11. A. Found, H.J. Müller: Searching for unknown feature targets on more than one dimension: investigating a "dimension-weighting" account. *Perception &*

Psychophysics 58, 88–101 (1996)

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

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

14. J.M. Wolfe: Visual Search. In: Attention. Ed: H. Pashler, (13-73). *Psychology Press*, Hove, UK (1998)

15. J.M. Wolfe, T.S. Howoritz: What attributes guide the deployment of visual attention and how do they do it? *Nat Rev Neurosci* 5, 495-501 (2004)

16. J.W. Todd: Reaction to multiple stimuli. *Science Press*, New York (1912)

17. J. Miller: Divided attention: evidence for coactivation with redundant signals. *Cognitive Psychology* 14, 247–279 (1982)

18. J. Miller: Timecourse of coactivation in bimodal divided attention. *Perception & Psychophysics* 40, 331-343 (1986)

19. M. Giray, R. Ulrich: Motor coactivation revealed by response force in divided and focused attention. *J Exp Psychol: Hum Perc Perf* 19, 1278–1291 (1993)

20. J.T. Mordkoff, S. Yantis: An interactive race model of divided attention. *J Exp Psychol: Hum Perc Perf* 7, 520-538 (1991)

21. J.T. Mordkoff, S. Yantis: Dividing attention between color and shape: Evidence of coactivation. *Perception & Psychophysics* 53, 357-366 (1993)

22. J. Krummenacher, H.J. Müller, D. Heller: Visual search for dimensionally redundant pop-out targets: Evidence for parallel-coactive processing of dimensions. *Perception & Psychophysics* 63, 901-917 (2001)

23. J. Krummenacher, H.J. Müller, D. Heller: Visual search for dimensionally redundant pop-out targets: Parallel-coactive processing of dimensions is location specific. *J Exp Psychol: Hum Perc Perf* 28, 1303-1322 (2002)

24. S. Katzner, L. Busse, S. Treue: Feature-based attentional integration of color and visual motion. *Journal of Vision* 6, 269–284 (2006)

25. C.A. Marzi, N. Smania, M.C. Martini, G. Gambina, G. Tomelleri, A. Palamara, F. Alessandrini, M. Prior: Implicit redundant-targets effect in visual extinction. *Neuropsychologia* 34, 9–22 (1996, Jan)

26. M.C. Corballis: Hemispheric interactions in simple reaction time. *Neuropsychologia* 40, 423–434 (2002)

27. D.H. Raab: Statistical facilitation of simple reaction times. *Transactions of the New York Academy of Sciences* 24, 574-590 (1962)

28. B. Schönwälder: Mechanisms of target selection and feature binding in visual object recognition: evidence from the asynchronous presentation of target features. Logos, Berlin, Germany (2006)

29. H. Colonius, A. Diederich: The race model inequality: Interpreting a geometric measure of the amount of violation. *Psychological Rev* 113, 148-154 (2006)

30. J. Miller: Channel interaction and the redundant-targets effect in bimodal divided attention. *J Exp Psychol: Hum Perc Perf* 17, 160-169 (1991)

31. A. Diederich, H. Colonius: Bimodal and trimodal multisensory enhancement: effects of stimulus onset and intensity on reaction time. *Perception & Psychophysics* 66, 1388-1404 (2004)

32. J.T. Mordkoff, J. Miller: Redundancy gains and coactivation with two different targets: the problem of target preferences and the effects of display frequency. *Perception & Psychophysics* 53, 537-535 (1993)

33. S. Pollmann, E. Zaidel: Redundancy gains for visual search after complete commissurotomy. *Neuropsychology* 13, 246-258 (1999)

34. S. Savazzi, C.A. Marzi: The superior colliculus subserves interhemispheric neural summation in both normals and patients with a total section or agenesis of the corpus callosum, *Neuropsychologia* 42, 1608-1618 (2004)

35. S. Grossberg, E. Mingolla, W.D. Ross: A neural theory of attentive visual search: interactions of boundary, surface, spatial, and object representations. *Psychological Rev* 101, 470–489 (1994)

36. J.T. Townsend, G. Nozawa: Spatio-temporal properties of elementary perception: An investigation of parallel, serial, and coactive theories. *Mathematical Psychology* 39, 321-359 (1995)

37. M. Zehetleitner, J. Krummenacher, H.J. Müller: Coactivation – neither serial exhaustive nor interactive race – models can explain redundancy gains in visual pop-out search. Unpublished manuscript (2007)

38. J.T. Townsend, G. Nozawa: Serial exhaustive models can violate the race model inequality: Implications for architecture and capacity. *Psychological Rev* 104, 595-602 (1997)

39. S. Sternberg: Memory-scanning: mental processes revealed by reaction-time experiments. *American Scientist* 57, 421-457 (1969)

40. A. Treisman, S. Gormican: Feature analysis in early vision - evidence from search asymmetries. *Psychological Rev* 95, 15-48 (1988)

41. L.K.H. Chan, W.G. Hayward: Feature Integration Theory revisited: Dissociating feature detection and attentional guidance in visual search. Unpublished manuscript (2007)
42. K. Mortier, W. van Zoest, M. Meeter, J. Theeuwes: No top-down control of selection in singleton search: Evidence from manual and eye movement responses in singleton detection and localization tasks. Unpublished manuscript (2007)
43. A. Cohen, U. Feintuch: The dimensional-action system: A distinct visual system. In: Attention and performance: XIX. Common mechanisms in perception and action. Eds: Prinz W., Hommel B.N., (587-608). *Oxford University Press*, Oxford, UK (2002)
44. A. Cohen, H. Magen: Intra- and cross-dimensional visual search for single feature targets. *Perception & Psychophysics* 61, 291-307 (1999)
45. A. Cohen, R. Shoup: Perceptual dimensional constraints in response selection processes. *Cognitive Psychology* 32, 128-181 (1997)
46. A. Cohen, R. Shoup: Response selection processes for conjunctive targets. *J Exp Psychol: Hum Perc Perf* 26, 391-411 (2000)
47. V. Maljkovic, K. Nakayama: Priming of pop-out: I. role of features. *Memory & Cognition* 22, 657-672 (1994, Nov)
48. H.J. Müller, B. Reimann, J. Krummenacher: Visual search for singleton feature targets across dimensions: Stimulus- and expectancy-driven effects in dimensional weighting. *J Exp Psychol: Hum Perc Perf* 29, 1021-1035 (2003)
49. H. J. Müller, J. Krummenacher: Locus of dimension weighting: Preattentive or postselective? *Visual Cognition* 14, 490-513 (2006)
50. C. Bundesen, T. Habekost, Søren Kyllingsbæk: A neural theory of visual attention: Bridging cognition and neurophysiology. *Psychological Rev* 112, 291-328 (2005)
51. J. Gottlieb: Parietal mechanisms of target representation. *Curr Opin Neurobiol* 12, 134-140 (2002)
52. K.G. Thompson, N.P. Bichot: A visual salience map in the primate frontal eye field. *Prog Brain Res* 147, 251 (2005)
53. N.P. Bichot, J.D. Schall: Effects of similarity and history on neural mechanisms of visual selection. *Nat Neurosci* 2, 549-554 (1999)
54. J.H. Fecteau, D.P. Munoz: Salience, relevance, and firing: a priority map for target selection. *Trends Cogn Sci* 10, 382-390 (2006)
55. C.J. Bruce, M.E. Goldberg, M.C. Bushnell, G.B. Stanton: Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *J Neurophysiol* 54, 714-734 (1985)
56. D.A. Robinson, A.F. Fuchs: Eye movements evoked by stimulation of frontal eye fields. *J Neurophysiol* 32, 637-648 (1969)
57. S. Kastner, K. DeSimone, C.S. Konen, S.M. Szczepanski, K.S. Weiner, K.A. Schneider: Topographic maps in human frontal cortex revealed in memory-guided saccade and spatial working-memory tasks. *J Neurophysiol* 97, 3494-3507 (2007)
58. T.R. Sato, A. Murthy, K.G. Thompson, J.D. Schall: Search efficiency but not response interference affects visual selection in frontal eye field. *Neuron* 30, 583-591 (2001)
59. S. Pollmann, R. Weidner, H.J. Müller, D.Y. von Cramon: A fronto-posterior network involved in visual dimension changes. *J Cog Neurosci* 12, 480-494 (2000)
60. S. Pollmann, R. Weidner, H.J. Müller, M. Maertens, D.Y. von Cramon: Selective and interactive neural correlates of visual dimension changes and response changes. *NeuroImage* 30, 254-265 (2006)
61. S. Pollmann, R. Weidner, H.J. Müller, D.Y. von Cramon: Neural correlates of visual dimension weighting. *Visual Cognition* 14, 877-897 (2006)
62. R. Weidner, S. Pollmann, H.J. Müller, D.Y. von Cramon: Top-down controlled visual dimension weighting: An event-related fMRI study. *Cerebral Cortex* 12, 318-328 (2002)
63. A. Kristjansson: "I know what you did last trial" – A selective review of research on priming in visual search. *Frontiers in Bioscience* 13, 1171-1181 (2008)
64. Z. Li: A saliency map in primary visual cortex. *Trends Cogn Sci* 6, 9-16 (2002)
65. J.J. Knierim, D.C. van Essen: Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J Neurophysiol* 67, 961-980 (1992)
66. A.G. Leventhal, K.G. Thompson, D. Liu, Y. Zhou, S.J. Ault: Concomitant sensitivity to orientation, direction, and color of cells in layers 2, 3, and 4 of monkey striate cortex. *J Neurosci* 15, 1808-1818 (1995)
67. E.N. Johnson, M.J. Hawken, R. Shapley: The spatial transformations of color in the primary visual cortex of the macaque monkey. *Nature Neuroscience* 4, 409-416 (2001)
68. H.S. Friedman, H. Zhou, R. von der Heydt: The coding of uniform colour figures in monkey visual cortex. *J Physiol* 548, 593-613 (2003)
69. K.R. Gegenfurtner, D.C. Kiper, S.B. Fenstermaker: Processing of color, form, and motion in macaque area V2.

Visual Neuroscience 13, 161-172 (1996)

70. A.R. Koene, L. Zhaoping: Feature-specific interactions in salience from combined feature contrasts: Evidence for a bottom-up saliency map in V1. *Journal of Vision* 7, 7 1-14 (2007)

71. G.D. Horowitz, T.D. Albright: Paucity of chromatic linear motion detectors in macaque V1. *Journal of Vision* 5, 6 525-533 (2005)

72. D.H. Hubel, T.N. Wiesel: Receptive fields of single neurons in the cat's visual cortex. *J Physiol* 148, 574-591 (1959)

73. D.Y. Ts'o, C.D. Gilbert: The organization of chromatic and spatial interactions in the primate striate cortex. *J Neurosci* 8, 1712-1727 (1988)

74. J. Krummenacher, H.J. Müller: The role of saliency signals in parallel coactive processing of dimensional information. Unpublished manuscript (2007)

Footnotes: ¹. The original proof is of course independent of the realization of the redundant signal or the type of additional factorial manipulation. For better readability the ideas have been formulated in terms of visual pop-out search.

Key Words: Redundant Signals Paradigm, Coactivation, Saliency Map, Dimension Weighting Account, Race Model Inequality, Review

Send correspondence to: Michael Zehetleitner, LMU Department Psychologie, Leopoldstr. 13, 80802 München, Germany, Tel: 0049-089-2180 5209, Fax: 0049-089-2180-5211, E-mail: zehetleitner@psy.uni-muenchen.de

<http://www.bioscience.org/current/vol13.htm>