
**Saliency from the decision
perspective:
Inferring the processing architecture
of pre-attentive vision with mental
chronometry**

Michael Zehetleitner



München 2007

**Saliency from the decision
perspective:
Inferring the processing architecture
of pre-attentive vision with mental
chronometry**
Michael Zehetleitner

Dissertation

an der Fakultät für Psychologie und Pädagogik

der Ludwig–Maximilians–Universität

München

vorgelegt von

Michael Zehetleitner

aus Kempten im Allgäu

München, den 2. November 2007

Erstgutachter: Prof. Dr. Hermann J. Müller

Zweitgutachter: Prof. Dr. Joseph Krummenacher

Tag der mündlichen Prüfung: 20. Dezember 2007

To Norbert Bischof, without whom not.

ACKNOWLEDGMENTS

Norbert Bischof, in his lectures and in personal communication, ignited my curiosity for animal and human psychology and provided me with a psychological and philosophical framework as footholds for investigating such issues. Additionally his cybernetic, supported by Felix Tretter's system theoretical approach influenced my thinking intensely.

In particular, I would like to thank my mentor and supervisor of this thesis, Hermann Müller, for focussing my mind to the art of psychological experimentation and for providing me with the freedom to pursue very interesting questions answerable with mental chronometry. I thank Joseph Krummenacher for introducing me to the redundant-signals paradigm and being my contact person for experimental, writing, or statistical questions from the beginning of my studies of neuro-cognitive psychology. Being part of a large unit of colleagues also was of great help for developing and discussing experiments, as well as for gaining new insights. Specifically, I would like to mention Thomas Geyer, Thomas Töllner, Zhuanghua Shi, and Dragan Rangelov for valuable discussions and input for this thesis. I had the chance to be supported by Henning Bumann, Yvonne Schiller, Michael Hegenloh, and Frieder Wormser, who carried out the experiments and collected the data.

All work related support would be nothing without my family: my wife Ilona, our sons Franz and Peter, as well as our parents. Finally, I thank the Deutsche Forschungsgemeinschaft for funding the CoTeSys Excellence Cluster, which financed my research.

TABLE OF CONTENTS

	Page
ABSTRACT	ix
1 INTRODUCTION	1
1.1 Saliency map models	1
1.2 Modulation of saliency	3
1.2.1 The redundant-signals effect	3
1.2.2 Dimension switch costs and cueing benefits	4
1.3 Alternative processing architectures	6
1.4 Scope of the present study	8
1.5 Summary of findings	9
1.5.1 Chapter 2	9
1.5.2 Chapter 3	9
1.5.3 Chapter 4	11
1.5.4 Chapter 5	13
1.6 Conclusion and outlook	14
2 WHAT THE REDUNDANT-SIGNALS PARADIGM CAN REVEAL	17
2.1 The redundant-signals paradigm	18
2.2 The question of architecture: parallel, co-active, or serial?	22
2.3 Is integration spatially selective?	27
2.4 Is integration feature-based or dimension-based?	32
2.5 Weighting or priming?	34
2.6 Implementation of saliency maps and dimensional weighting in the brain	38
2.7 Conclusion	43
3 CO-ACTIVATION VS. SERIAL AND PARALLEL MODELS	45
3.1 Experiment 1	53
3.1.1 Method	54
3.1.2 Data Analysis	57
3.1.3 Results	60
3.1.4 Discussion	63
3.2 Experiment 2	64
3.2.1 Method	64
3.2.2 Results and Discussion	65
3.3 Experiment 3	66
3.3.1 Method	72
3.3.2 Results	74

	Page
3.3.3 Discussion	77
3.4 General Discussion	78
4 INTENTION AND TRIAL HISTORY IN LOCALIZATION	89
4.1 Experiment 1	95
4.1.1 Method	96
4.1.2 Results	99
4.1.3 Discussion	102
4.2 Experiment 2	105
4.2.1 Method	106
4.2.2 Design and Procedure	106
4.2.3 Results	107
4.2.4 Discussion	107
4.3 General Discussion	110
4.3.1 Summary of findings	110
4.3.2 Relations to post-selective accounts	112
4.3.3 Relation to further studies	114
4.3.4 Summary and conclusion	118
5 DECISION PERSPECTIVE ON SALIENCY	121
5.1 The RSE	133
5.1.1 Manipulation of feature contrast	133
5.1.2 Manipulation of response bias	134
5.1.3 The effect of spatial attention	136
5.2 Benefits from Dimensional Cues	136
5.3 Dimension Switch Costs	137
5.4 Experiment 1	138
5.4.1 Method	139
5.4.2 Results	141
5.4.3 Discussion	148
5.5 Experiment 2	150
5.5.1 Method	150
5.5.2 Results	152
5.5.3 Discussion	153
5.6 General Discussion	153
5.6.1 Relation to previous studies	156
5.6.2 Theoretical Implications	164
5.7 Conclusion	165
REFERENCES	167
VITA	178

ABSTRACT

There are several cognitive and neuro-scientific models of early, pre-attentive visual processing, with saliency map models being the particular dominant ones. Although they are very specific about how feature contrast and salience are being computed (including stimulus- and observer-driven influences), there usually is a theoretical gap between models dealing with visual analysis (such as the dimension weighting account, Müller, Heller, & Ziegler, 1995; Found & Müller, 1996) and models describing decision and response selection processes (e.g. the Ratcliff Diffusion Model, Ratcliff, 1978). Consequently, I propose that investigating saliency from a decision perspective that is by applying mathematical theories of decisions to several tasks that can be performed supposedly based on a salience map (e.g. detection, localization, attentional selection), explanatory power is increased and new hypothesis can be generated.

Further, several issues of visual pre-attentive processing are currently still under debate. Specifically, the exact nature of the pre-attentive architecture is disputed with regard to trial sequence, intention, and redundancy effects. The present study targets at shedding new light on the question of pre-attentive processing architecture (serial, parallel independent or interactive, co-active), top-down penetrability of pre-attentive vision, and pre-attentive vs. post-selective locus of dimensional inter-trial effects.

In summary, the findings presented in this thesis support the predictions of the dimension weighting account that is that dimension-specific feature contrast signals are computed in parallel and are dimensionally weighted before being pooled on the master saliency map. They further support the view that dimensional weighting operates at a pre-attentive processing stage and is modulated by trial history (i.e. bottom-up) as well as by intention (i.e. top-down). Further, the present results suggests that different tasks such as detection, localization, and compound tasks act upon

the same (salience map) representation. Finally, examining salience from a decision perspective provides an explanatory framework of how the observed performance difference induced by the redundant-signals effect, dimension switch costs, or benefits from dimensional cues depend on the duration of the underlying decision process.

1. INTRODUCTION

All animals are limited in their possibilities to interact with the environment. Humans and other primates have two hands to grasp or point with, one head to turn, one pair of eyes with one fovea each to look at with, while the amount of possible locations and objects in the environment to use these effectors on seems unlimited in comparison. Additionally the processing system to control behavior, the brain, is limited in its capacity to deal with the environmental complexity, as well (Tsotsos, 1990; Rensink, O'Regan, & Clark, 1997). Given these limitations of animals with regard to the complexity of the environment it was necessary to adaptively solve a problem of selection. For an adaptive control of selection there had to be a trade-off between responding to the affordances of the environment and choosing actions based on the internal state of the organism. The effect of internal states can operate on the sensory level, by enhancing certain sensory properties that are known to be more relevant over others, or on a semantic level, for which object recognition is necessary. Both response and choice behavior considered in the present study include binary decisions about the absence or presence, or the rough location (e.g. left vs. right) of a target, as well as selection decisions. Selection decisions determine where in space to move an effector or where to commit the limited processing capacity (i.e. covert attention) to next.

1.1 Saliency map models

Both binary decisions (e.g. about presence/absence, left/right location of a target) as well as multiple decisions (where to move an effector or attention next) depend on determining where the most interesting and important location is, at any given

moment - depending both on the properties of the environment (i.e. bottom-up or stimulus driven), as well as on the internal state of the animal (i.e. top-down or driven by intention). The solution to the problem of selection convergingly assumed by cognitive (Treisman, 1986, 1988; Wolfe, 1994; Müller, Heller, & Ziegler, 1995; Theeuwes, Kramer, & Belopolsky, 2004), neuro-cognitive/computational (Koch & Ullman, 1985; Itti & Koch, 2000; Li, 2002) models, as well as neurophysiology (Fecteau & Munoz, 2006; Gottlieb, 2002; Goldberg, Bisley, Powell, & Gottlieb, 2006) is a topographical representation of the environment that signals distinctiveness of locations irrespective from what sensory properties this distinctiveness originates from, mostly called a saliency map. Additionally to representing properties of the environment, a saliency map also represents importance of a location based on the internal state of the organism, in that sense it is also called relevance or priority map. The present study is focussed on the sensory processing aspect of saliency maps and how sensory processing is affected by the internal state of animals.

The most explicit model that describes a master saliency map in the visual domain is the computational model of Itti and Koch (2000) and Itti and Koch (2001). It is based on the neuro-computational work of Koch and Ullman (1985) as well as on cognitive models such as Guided Search (Wolfe, 1994) or extensions thereof (e.g. the dimension weighting account Müller et al., 1995; Found & Müller, 1996). Itti and Koch (2000) assume that initially local feature contrast is extracted in parallel for all locations and a number of features (i.e. different colors, orientations, motion directions, intensities of luminance). At each location not only strength of feature presence is taken into account, but also the difference between presence of a feature at a given location and its immediate surroundings. That way the local distinctiveness can be calculated. The feature contrast signals are then pooled into dimension specific maps, which are again further pooled into a master saliency map. The topographical representation of the visual scene on the saliency map signals strength of feature contrast or distinctiveness at each location in a feature unspecific manner. If there is strong activity at one location of the saliency master map, this activity signals presence of distinctive-

ness but carries no information about what feature this contrast derives from, e.g. whether from a red among green spots, or one single moving among stationary items.

1.2 Modulation of salience

Activity on the salience map can be modulated in various way. Trivially, strength of feature contrast influences both the level as well as the time course of activity on the salience map. Important for the present study are modulations due to targets being defined in two dimensions simultaneously (redundancy), modulations stemming from the trial history (dimension specific intertrial effects), as well as modulations due to intention or expectation of the observers (benefits from dimensional cues).

1.2.1 The redundant-signals effect

In the redundant-signals paradigm (Todd, 1912; Raab, 1962; Miller, 1982) targets can either have one or two response defining features. A usual finding is a performance benefit of targets with two response defining features (redundant targets) over targets with only one response defining feature (redundant signals effect, RSE). Such a benefit can be explained in several mutually exclusive ways: Raab (1962) explained the RSE with statistical facilitation. In that view, both target defining features in redundant targets are processed in parallel and can trigger a response independently. Due to the distribution of processing times it is probable that one of the both parallel channels is processed faster than the other. That is, compared to single signal trials, where the processing time can be modeled as being drawn from one random distribution, in redundant signal trials the processing time can be modeled as the minimum of two processing times, being drawn independently from the respective random distributions. That way, statistical facilitation can explain the RSE. However, Miller (1982) showed that statistical facilitation has an upper boundary of how much redundant targets can be faster than single targets. In his well-known race model inequality (RMI, Miller, 1982) he showed that the minimum of two distribu-

tions (i.e. the maximal gain from statistical facilitation) is bounded by the sum of both single target distributions. That is, statistical facilitation predicts that the fast reactions to redundant signal trials are not faster than the fast reactions to single signal trials. If the RMI is violated that is if the fast reactions to redundant targets are faster than the fast reactions to single targets, a parallel race model can not be the architecture underlying processing of the task. As an alternative, Miller (1982) proposed the signals of both processing channels for single targets to be summed, or integrated before triggering of the response (i.e. a co-activation model).

Krummenacher, Müller, and Heller (2001) applied this reasoning to visual pop-out search. They compared detection performance for targets that differed from distracters in only one dimension (e.g. a red vertical bar among green vertical bars) to performance for targets that differed from distracters in two dimensions, redundantly (e.g. a red and horizontal among green vertical bars). Krummenacher et al. (2001) found a RSE that is a mean reaction time benefit of redundant over single targets, and further reported violations of the RMI, excluding a parallel race model. That is models, which assume parallel processing of feature contrast for different dimensions are restricted regarding how detection responses are triggered: a parallel architecture, in which dimension specific feature contrast signals race independently for triggering a response are excluded due to violations of the RMI. The alternative of salience map based models are well possible to account these data, for they assume that dimension specific feature contrast signals are pooled/summed into a salience map before a response can be triggered. That way redundant targets lead to a faster and/or higher build-up of activation on the salience map compared to single targets.

1.2.2 Dimension switch costs and cueing benefits

A second effect relevant for the present study concerns trial history, which has been reported both when targets and distracters are repeated or change their roles

(e.g. Maljkovic & Nakayama, 1994) and when the target defining dimension is repeated or changes (e.g. Müller et al., 1995; Found & Müller, 1996). For instance, Found and Müller (1996) compared performance for detection of a pop-out target for sequences of trials in which the target defining dimension stayed the same to when the dimension changed. They found response times for targets in trial n to be faster, when in trial $n - 1$ a target of the same dimension (e.g. color in trials $n - 1$ and n) and to be slower, when in trial $n - 1$ a target of different dimension (e.g. color in trial n and orientation in trial $n - 1$) has been shown. That is, although the target in trial n is physically the same in both cases, performance depends on the trial history: if the previous target was defined in the same dimension, reaction times are faster than if the previous target was defined in a different dimension. Müller and colleagues proposed a dimension weighting account in order to explain these findings (Müller et al., 1995; Found & Müller, 1996). They extended a-historic salience map models, such as Guided Search (Wolfe, 1994) with the conception of dimensional weights. These dimensional weights are assumed to modulate dimension specific feature contrast signals before they are pooled into the salience map. That is, the higher the weight of a specific dimension, the greater it's impact on the salience map. Increased weights are assumed to lead to a faster and/or higher build-up of activation on the salience map. Shifting of weights according to the dimension weighting account can occur passively depending on the trial history (for active modulations see the following paragraph). Additionally increasing the weight of one dimension, according to the dimension weighting account, leads to decrease of weights of all other dimensions (i.e. the weights act as a limited resource). If for example the dimensional weights are equal for color and orientation, and a color target is presented (in trial $n - 1$), this leads to an increase of color and to a decrease of orientation weights. In the next trial, n , a second color target has a higher impact on the salience map due to the increased color weights, whereas an orientation target has a lower impact on the salience map. Consequently, selection or detection processes based on activity on the salience map operate in a facilitated fashion for color compared to orientation targets in trial n .

In addition to this passive change of weights, the dimension weighting account proposes that dimensional weights are subject to influence of the observer. That is, if the observer knows the target defining dimension in advance, she can use that prior knowledge to prepare intentionally for a target in that dimension. The dimension weighting account assumes that this preparation is reflected in an increase of weight in the prepared-for and a decrease of weight in the unprepared dimensions. Müller, Reimann, and Krummenacher (2003) investigated this possibility of top-down influences on weights with a trial-by-trial cueing procedure. In each trial the observers received either a neutral or a dimensional cue. Dimensional cues could be either valid or invalid. They found reaction times to be fastest, if the cue was valid, followed by neutral cues and slowest for invalidly cued trials. Müller et al. (2003) concluded that prior knowledge about the upcoming target dimension did change performance. In terms of the dimension weighting account preparation for the next target dimension based on the cue leads to an increase of weight for the cued and a decrease of weight for all other dimensions. Therefore, impact of feature contrast signals from the cued dimension on the salience map is higher, as they are modulated by increased weights, whereas impact of feature contrast from the other dimensions is weaker. This leads to better performance for validly cued compared to invalidly cued trials, because according to the dimension weighting account activity on the salience map is further processed by detection and selection mechanisms, which work faster the higher and more distinct signals on the salience map are.

1.3 Alternative processing architectures

However, although salience summation models of visual search with the extension of the dimension weighting account are a theory driven framework that can explain several phenomena in search behavior, these models are not undisputed. For several aspects of salience based models of visual search, there are alternative accounts of

how the processing architecture looks like. These aspects include the architecture of dimension based feature contrast signals: are they processed in a coactive/integrative fashion, as salience map models propose, or are they rather processed in a parallel or serial fashion? If several dimensions are relevant for a task, how does search terminate? Saliency map models assume a self-terminating search, but there is the alternative stopping rule of exhaustive search. Is the pre-attentive processing stage (i.e. computation of salience until integration in the salience map, guiding deployment of attention) modulable by top-down signals (i.e. internal states of the observer) or is it totally stimulus-driven? An additional question about the processing architecture is about processing routes for different tasks: as a topographical representation a salience map inherently contains spatial signals. How does processing of tasks look like, which do not depend on spatial information, such as detection tasks? Are detection tasks also processed via the salience map, or is there an additional processing route for detection?

These alternatives are substantiated by either theoretical considerations or empirical evidence. Townsend and Nozawa (1997) for instance have shown that violations of the RMI can also occur in serial models with an exhaustive stopping rule, whereas Mordkoff and Yantis (1991, 1993) argued that a parallel architecture, in which the channels can exchange information also can lead to violations of the RMI. Therefore, the RMI loses its decisiveness in determining the processing architecture (i.e. coactive/integrative processing), although it still is fit to exclude the possibility of parallel race models if violated. Regarding top-down penetration of pre-attentive processing, recent studies failed to find performance benefits from dimensional cues for compound (Theeuwes, Reimann, Brit, & Mortier, 2006) or localization tasks (Mortier, Zoest, Meeter, & Theeuwes, 2007). Further, dimension specific intertrial effects have been found reduced or even absent in compound and localization, as compared to detection tasks (Chan & Hayward, 2007; Mortier et al., 2007). These divergent findings lead to a revival of Feature Integration Theory (FIT: Treisman & Gelade, 1980; Treisman, 1988) in assuming different processing routes for tasks that require spatial informa-

tion compared to tasks can be solved without spatial information. FIT proposes dimension specific, spatially pooled signals that can be searched in a serial fashion in order to simply detect a target. If however the task requires spatial information, such as coarse localization or further processing possible only after attentional selection, the salience map is employed to solve these tasks. In summary, the alternatives to salience map based models such as the dimension weighting account (Müller et al., 1995) include the question of serial (Townsend & Nozawa, 1997; Treisman, 1988) or parallel interactive (Mordkoff & Yantis, 1991) architectures without the possibility of top-down penetration (Theeuwes et al., 2006; Mortier et al., 2007), as well as different processing routes for localization and compound vs. detection tasks (Chan & Hayward, 2007; Mortier et al., 2007).

1.4 Scope of the present study

The present study aims at investigating these alternatives with the means of mental chronometry. In addition to analysis of mean reaction times (among other things Sternberg's method of additive factors: Sternberg, 1969a; Townsend & Nozawa, 1995), distributional measures such as the RMI (Miller, 1982), as well as fitting the Ratcliff Diffusion Model (Ratcliff, 1978), a type of sequential sampling models, was applied in the present study. Chapter 2 comprises a theoretical review about salience map models and how the redundant-signals paradigm can contribute to the examination of pre-attentive vision. Chapter 3 tests alternatives to salience summation models, which also can lead to violations of the RMI (specifically serial exhaustive and parallel interactive models). Chapter 4 investigates the question of top-down modulations of pre-attentive processing. Chapter 5 investigates salience from the decision perspective with respect to modulations of pre-attentive modulations of visual processing in different tasks.

1.5 Summary of findings

1.5.1 Chapter 2

Chapter 2 reviews contributions of the redundant-signals paradigm in visual search to answering the questions of architecture, spatial selectivity of integration, and of feature- or dimension-based integration. Further, there the distinction between dimensional weighting (where increasing preparedness for one dimension leads to a decrease for other dimensions) and dimensional priming (where one dimension is prepared for, while leaving the other dimensions unaffected). Finally, evidence for neuronal implementations of saliency maps are reviewed and opposed critically to recent proposals of salience maps in V1 (Li, 2002).

1.5.2 Chapter 3

In Chapter 3, the alternative architectures to co-activation (as in salience map models), namely serial exhaustive and parallel interactive models are tested. In two experiments the redundant-signals paradigm has been implemented, in one to test against serial models and various stopping rules, employing a procedure proposed by Townsend and Nozawa (1995), dubbed the double-factorial design, and in the other to test against interactive parallel models. Both models assume different processing architectures for redundant targets. Both models, as well as parallel independent and co-activation models agree that initially feature contrast of different dimensions is calculated in parallel. In model terms, there is one channel processing feature contrast of one dimension (e.g. luminance) and the other channel computes feature contrast of a different dimension (e.g. orientation). All models differ in their assumptions of how responses are triggered. As described previously, parallel independent models assume that in case of a redundant target, both channels process feature contrast, and a response is triggered as soon as one channel detects presence of a target. Remember that redundant targets are differing from distracters in two dimensions simultane-

ously. Hence, each channel can detect presence of one component of the redundant target, enough for triggering a 'present' response. Co-activation models in contrast assume that signals from both channels are integrated or pooled, and that a detection module monitors this sum of signals. A response here is triggered, as soon as this pooled signal reaches a certain threshold. Serial models in contrast assume that when encountering a redundant target, both channels are examined for presence of a target in a serial fashion, first one - then the other. A self-terminating serial model would stop processing as soon as in one channel a high level of feature contrast would be detected. Townsend and Nozawa (1997) however showed that exhaustive serial models can lead to violations of the RMI, if detection of presence of a target is faster than detection of absence. In case of a serial exhaustive search of both potentially relevant dimensions, in case of a target defined in only one dimension, the decision time is composed of detecting presence in one channel and absence in another channel. In case of a redundant target, this serial exhaustive process detects presence in both channels. Under the assumption that detection of absence takes longer than detection of presence, this leads to a comparative slowing of single as opposed to redundant targets, making violations of the RMI possible. Finally, in interactive parallel models, channels are assumed to be able to communicate and exchange information across both channels before response decision. Mordkoff and Yantis (1991) proposed that if the conditional probabilities of target presence in both dimensions are greater than zero, it is possible that redundant targets benefit more from channel cross-talk than single target, leading to a processing advantage for redundant over single target trials that can be reflected in violations of the RMI. Mordkoff and Yantis (1991) provided formulas to calculate the amount of benefit of redundant over single signal trials. In a double factorial design, targets can be defined in one or two dimensions, while intensity (strength of feature contrast in the present case) may be high or low. Different processing architectures were shown to provide distinctive interaction patterns when comparing reaction times for redundant targets composed of different levels of feature contrast (Townsend & Nozawa, 1995). As a redundant target is always

composed of one component being defined in one and another component defined in a second dimension, it is possible to examine the interaction of a feature contrast manipulation, independently for both dimensions of redundant targets. Serial models predict that such a feature contrast manipulation is additive. Parallel self terminating as well as co-activation models lead to super-additive and parallel exhaustive models to sub-additive interactions. The present study reports a super-additive interaction, ruling out serial models of any stopping rule. Further violations of the RMI reported here substantiate the exclusion of parallel independent models.

In the second experiment of Chapter 3, the amount of cross-talk beneficial for redundant targets was manipulated by changing the conditional probabilities of target presence in both dimensions in a between-subjects design. By varying the ratio of target present to target absent trials crossed with a variation of the ratio of single to redundant target trials it was possible to examine four conditions, which varied in the amount of benefit for redundant targets, if cross-talk were possible. An interactive race model predicts that the RSE as well as the amount of violations of the RMI (Colonius & Diederich, 2006) should be positively correlated to the amount of cross-talk beneficial for redundant targets. However, no such correlation was found. Based on the exclusion of two alternatively possible processing architectures - serial exhaustive as well as interactive parallel models - evidence converges that indeed co-activation models, such as salience summation models are responsible for processing of pop-out targets in visual search.

1.5.3 Chapter 4

Two questions were addressed in Chapter 4: is pre-attentive vision penetrable by top-down intentions, and are there two different or one common processing routes for detection and compound-type of tasks? In order to test the pre-attentive nature of previously reported benefits from dimensional cues (Müller et al., 2003; Theeuwes et al., 2006; Müller & Krummenacher, 2006) as well as of dimension intertrial costs (e.g.

Found & Müller, 1996), a signal detection paradigm with very brief display durations and accuracy as the dependent variable was employed. If post-selective processing stages were responsible for such cueing benefits, they should not be observable under brief display durations (Huang & Pashler, 2005; Prinzmetal, McCool, & Park, 2005). Conversely, if both effects of trial history (i.e. dimension inter-trial effects) and/or effects of intention (i.e. benefits from dimensional cues) are observable in accuracy under brief display durations, these effects can safely be assumed to arise from pre-attentive processing stages. In the first experiment of Chapter 4, indeed cueing benefits as well as intertrial effects are reported in accuracy (as well as in unspeeded reaction times), supporting models like the dimension weighting account, which assume that both effects arise from modulations of dimension-specific feature contrast signals before integration into a salience master map.

In order to reconcile recent null findings of dimensional cueing benefits for other than detection, especially for localization and compound tasks, the second experiment investigated cueing benefits in a reaction time paradigm. Reasoning that maybe in localization tasks performance is at ceiling for high contrast targets, rendering cueing benefits unobservable, a manipulation of feature contrast was carried out. Indeed, the absence of cueing benefits for high contrast targets was replicated, but substantial benefits were observed for low feature contrast targets. That is, rather than a qualitative difference in processing between detection and localization tasks, as proposed by the dual-route account (Mortier et al., 2007; Chan & Hayward, 2007; Treisman, 1988), a quantitative difference seems to exist. However, even though a quantitative difference between detection and localization tasks exists, which yet needs explanation, both tasks are qualitatively processed by largely the same architecture. In terms of the dimension weighting account, both tasks depend on activity of the salience master map.

1.5.4 Chapter 5

Chapter 5 finally investigates salience from a decision perspective, combining cognitive, computational and neurophysiological theories of salience computations with mathematical theories of decision making (Ratcliff, 1978). A rather abstract implication of decision models, such as diffusion models is examined in the context of visual pop-out search: sequential sampling models of decisions predict that modulations of sensory processing, reflected in different drift rates, evolve into larger differences in decision (and consequently observed reaction) times, the longer the decision takes. This is independent of whether the decision process is prolonged by a more conservative criterion or by a lower base drift rate. As the RSE, dimension switch costs, as well as benefits from dimensional cues are supposed to modulate activity on a saliency master map (Müller et al., 1995; Found & Müller, 1996), these effects should be reflected in differences of drift rates, when diffusion models are used to describe decisions, which are based on such a saliency map. Therefore, applying the rather abstract property of diffusion processes to these effects, the observable reaction time differences for redundant vs. single targets, for same vs. different inter-trial transitions, and for validly vs. invalidly cued trials, should be larger the longer the decision process takes - independent of how the decision process is prolonged (criterion or base drift rate), and independent of what task is carried out using signals of the saliency map (e.g. detection, localization, or compound tasks).

Reviewing previous experiments, re-analysis of previous studies, as well as two new experiments support exactly that view: the size of the RSE, of dimension switch costs, and of benefits from dimensional cues are larger for low than for high feature contrast targets (due to lower base drift rate), as well as for a more conservative criterion (achieved by manipulating the frequency of present vs. absent responses). These effects can be found for localization, detection, as well as compound tasks. Such findings support the dimension weighting account in its generality and, when

combined with a decision perspective, provide a rather unifying account of several up till now contradicting findings.

1.6 Conclusion and outlook

This thesis provides a theoretical, an empirical, and a methodological contribution to the current debate about visual early processing. First, it contributes to closing the gap between precise models of visual processing and of how decisions and responses are generated by applying sequential sampling models (especially the RDM, Ratcliff, 1978) to saliency summation models. Thereby, evidence presented in this thesis suggests that the early processing of several tasks, specifically detection, coarse localization, and compound tasks, largely overlaps. In terms of saliency map models this overlap reaches along computations of feature contrast, including modulations by trial history and intention, to the salience master map.

Second, this thesis empirically strengthens support for the core assumptions of the dimension weighting account: the effect of dimensional weighting derives from pre-attentive rather than post-selective processing stages (Chapter 4 and 5), these pre-attentive weights can be influenced by top-down intentions (Chapter 4), theoretically possible alternatives to feature contrast summation (i.e. co-activation) models, such as serial exhaustive or parallel interactive models can be excluded (Chapter 3), and several tasks (detection, localization, and compound) seem to be based on activations of an overall salience map (Chapter 5).

Third, in this thesis reports one of the first (cf. Thornton & Gildea, 2007) application of sequential sampling models in general, and the first application of the RDM in specific to visual search paradigms. That way, it is possible to empirically discern, whether manipulations affect for instance the quality of stimuli, observers' response criteria, or non-decision (e.g. motor) related processing stages. Aided by the improved handling of the RDM by the Matlab toolbox provided by Vandekerckhove and Tuerlinckx (2007a), this thesis provides for instance evidence that dimension weighting

rather affects drift rates (i.e. the quality of stimuli) than motor processing times, whereas manipulation of the stimulus-response congruency has a post-decisional effect. Investigating salience from a decision perspective can be the fruitful basis for further research, as it generates several questions and testable hypothesis. The salience map has been theorized to guide overt and covert attention (Koch & Ullman, 1985; Wolfe, 1994; Itti & Koch, 2000), but how can these decisions with multiple choices be modeled? How well can for instance leaky, competing accumulator models (Usher & McClelland, 2001) describe the process of deciding, what location will be selected next, based on activity of the salience map? How are localization (4-choice, 2-choice) and detection decisions computed from activity on the salience map, as they need less to none spatial information?

When applied to dynamic situations in multiple sensory modalities with various effectors, the decision perspective makes specific hypothesis about how the size of pre-attentive modulations (e.g. by trial history or by intention) depends on the duration of the underlying decision process. Finally, the decision perspective can be applied to other modulations of visual processing, such as spatial attention (Posner, 1980), weighting of modalities (Töllner, Gramann, Müller, & Eimer, 2007a), weighting of effectors (Töllner, Gramann, Müller, Kiss, & Eimer, 2007), or priming of positions (Maljkovic & Nakayama, 1996; Geyer, Müller, & Krummenacher, 2007).

2. WHAT THE REDUNDANT-SIGNALS PARADIGM CAN REVEAL

In their seminal work, Livingstone and Hubel showed that visual processing operates in separate and parallel 'channels' from the retinal level onwards (e.g. Livingstone & Hubel, 1984, 1987, 1988; Hubel & Livingstone, 1985, 1987). 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., Sincich & Horton, 2005, 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: Treisman & Gelade, 1980; Treisman, 1988), Guided Search (GS: Wolfe, 1994), Dimension Weighting (DW: Müller et al., 1995; Found & Müller, 1996), and the neuro-computational models of Koch and Ullman (1985) and Itti and Koch (2000). 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 Wolfe, 1998b; Wolfe & Horowitz, 2004, for reviews). The models differ, however, in the assumptions about later stages that follow the initial parallel coding stage. In this paper, I summarize recent evidence from studies of redundancy gains in visual pop-out search, which can help to deter-

mine 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 (Wolfe, 1998b). While there is general agreement that features are initially registered in parallel in dimensionally segregated pathways, there are various models about the subsequent processing stages required for successful target detection. These models differ in a number of respects: (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), (iv) dependency of target detection on the prior allocation of focal attention (pre-attentive vs. post-selective), and (v) top-down penetrability of the further processing stage(s). First, I will describe the basic redundant-target paradigm and its applicability to visual pop-out search. Second, I will elaborate each of the above questions and review relevant studies that show how analysis of redundancy gains can contribute to answering them. Third, I will review evidence from single-cell recordings in monkeys and fMRI studies in humans concerning the neural implementation of pre-attentive search processes.

2.1 The redundant-signals 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. Todd, 1912; Miller, 1982, 1986; Giray & Ulrich, 1993; Mordkoff & Yantis, 1991, 1993; Krummenacher et al., 2001; Krummenacher, Müller, & Heller, 2002; Katzner, Busse, & Treue, 2006; Marzi et al., 1996; Corballis, 2002). 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 (1962) 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) (see Figure 2.1). 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 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 statistically 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 (1982) 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

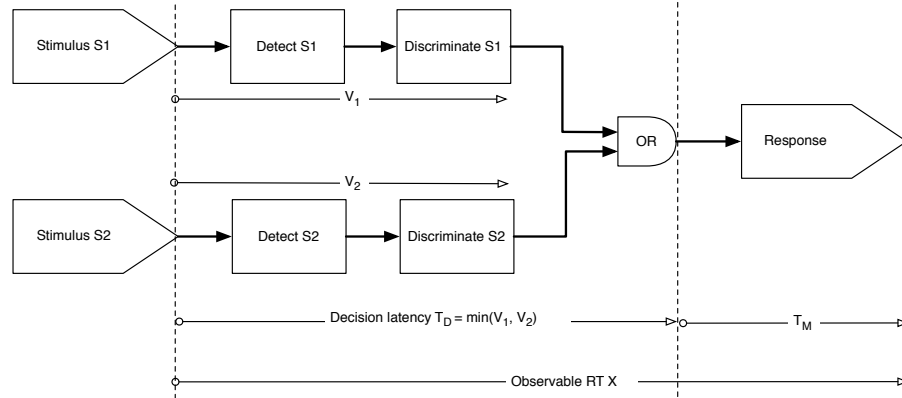


Figure 2.1. The observed reaction time in a race model is the sum of T_M , the time necessary for response and motor processing, and the minimum of the two detection times T_D for signals S_1 and S_2 . Each single signal has its own decision unit that can trigger the response.

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 (Colonius & Diederich, 2006) – 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:

$$P(RT < t|S_{12}) \leq P(RT < t|S_1) + P(RT < t|S_2), \quad (2.1)$$

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 (1982, 1986) are co-activation models (see Figure 2.2). 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.

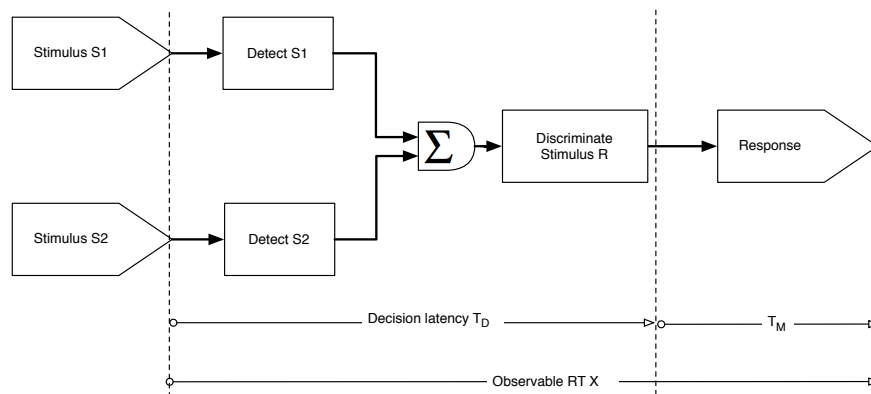


Figure 2.2. The observed reaction time in a co-activation model is the sum of T_M , the time necessary for response and motor processing, and the decision latency T_D . Both single signals are red into a common decision unit.

The redundant target paradigm has been applied in a variety of areas, such as bi-modal (e.g., visual and auditory) processing (Miller, 1982, 1991; Diederich & Colonius, 2004), divided attention (between two locations: Mordkoff & Yantis, 1991, 1993; Mordkoff & Miller, 1993), and neuro-psychological research (Pollmann & Zaidel, 1999; Savazzi & Marzi, 2004; Corballis, 2002). 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 dimension (e.g., orientation) can be denoted as a single stimulus S1, and feature contrast in another dimension (e.g., color) as single stimulus S2. Redundant targets are then the combination of both S1 and S2, 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 are elicited: if feature contrast in both dimensions can independently elicit a 'detection' response, then redundancy gains are expected without violations of the RMI. If the RMI is violated, such a parallel architecture 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 (e.g. Wolfe, 1994; Itti & Koch, 2000)? 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. Grossberg, Mingolla, & Ross, 1994; Treisman, 1988), or in an exhaustive manner (i.e., all dimensions are serially checked even if a target is found in one; e.g. Townsend & Nozawa, 1995)? For divided-attention paradigms, in which a target may be presented at two possible

locations, Mordkoff and Yantis (1991, 1993) 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 et al. (2001) 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 2.3 for an illustration of the stimuli).




Figure 2.3. Example displays with targets defined in a single dimension or redundantly in two dimensions. (a) presents a color, (b) an orientation, 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.

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 (2007a), combining a redundant-signals paradigm with a double-factorial design (Townsend & Nozawa, 1995). 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 (Townsend & Nozawa, 1997) or interactive-race models (Mordkoff & Yantis, 1991, 1993). 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 (1969a)'s 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 (1995) 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 (1995) 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 (2007a) 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

¹The 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.

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 (Mordkoff & Yantis, 1991, 1993). 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 (Krummenacher et al., 2001, 2002), (Zehetleitner, 2007a) 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 (2007a) 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 (2007a) 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 Krümmenacher et al. (2001) as well as of Zehetleitner (2007a) 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 (1994; see also Müller et al., 1995) and Itti and Koch (2000) assume that the integration stage is topographically organized that is, the integration is spatially specific (see Figure 2.4). 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, Treisman & Gelade, 1980; Treisman, 1988), 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 2.5. That is, crucial with 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 (2007) and Mortier et al. (2007). 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 (2002), which is based on the cross-dimensional response selection model

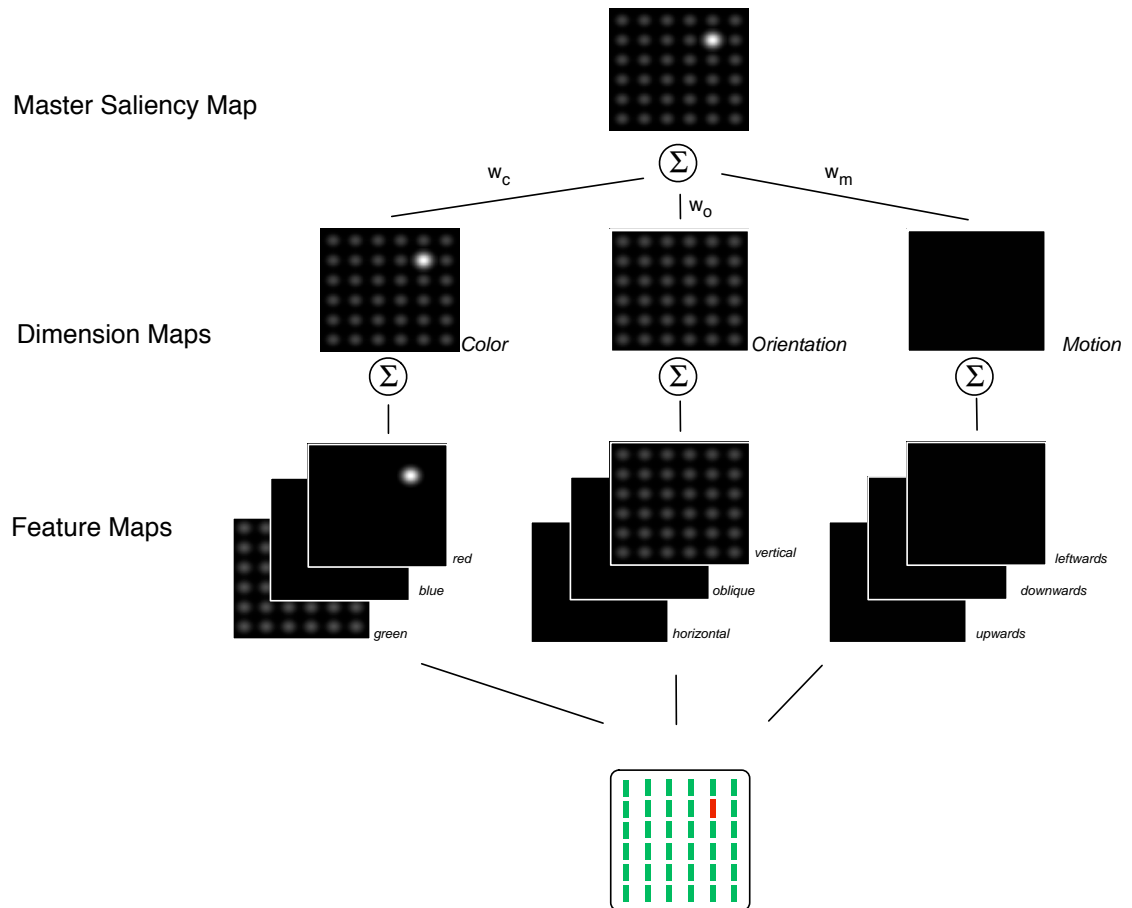


Figure 2.4. Example saliency summation model The display is first analyzed by spatio-topically organized feature analyzers for different colors, orientation, motion directions, etc. Each map is a topographical representation of the display, with black representing no activity and white strongest activity. 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.)

of Cohen and Magen (1999), see also Cohen and Shoup (1997, 2000). The DA model assumes that there are dimension-specific feature analyzer units as well as multiple response selection units, one per visual dimension (Cohen & Feintuch, 2002, 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

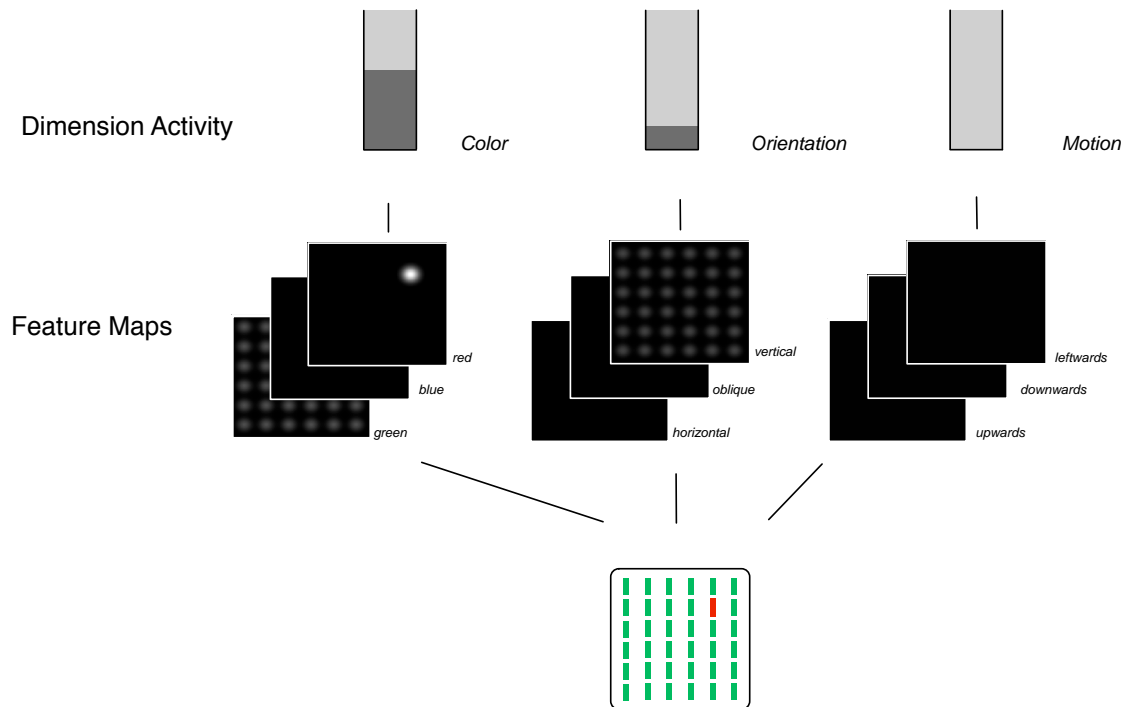


Figure 2.5. Example model in which presence of feature contrast in one dimension is represented in a non-spatial fashion . The display is analyzed in terms of color, orientation, and motion contrast. The possible activity is represented by the light grey boxes. Large activity is represented by a tall dark gray bar, low activity by a small dark gray bar.

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), 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 Cohen and Feintuch (2002, 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 (e.g., Wolfe, 1994; Itti & Koch, 2000) 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. (2002) 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 posi-

tion, 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. (2002) 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 in 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 pre–attentive process that is, independent of the locus of focal attention – a view advocated by Krummenacher et al. (2002). 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 (e.g., Itti & Koch, 2000). Although spatially specific,

spatial attention is not necessary for integration: integration of redundant pop-out targets is a pre-attentive process.

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 pre-attentive. Cognitive and neuronal models that are based on an overall-saliency map (e.g., Wolfe, 1994; Müller et al., 1995; Itti & Koch, 2000) additionally assume that feature contrast signals are dimensionally pooled before being integrated into the overall-saliency map: as illustrated in Figure 2.4, 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 Wolfe, 1998) are first summed into a dimension-specific map. Dimension-specific signals may be weighted prior to being transferred to the master saliency map (Müller et al. 1995). 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 (e.g., Itti & Koch, 2000) 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 (Krummenacher et al., 2002). But in order to assure that all dimensions contribute equally to overall-saliency, dimensional signals are normalized (to values between 0 and 1; Itti & Koch, 2000, 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 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. (2002) 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 (e.g., Müller et al., 1995; Itti & Koch 2000).

2.5 Weighting or priming?

Maljkovic and Nakayama (1994) 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 distractors 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/distractor feature swap. They found that, even when the sequence was made perfectly predictable (i.e., the target definition changed regularly every two trials in AABB manner) and observers were informed about this rule, change of the target/distractor 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. (1995). 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 (1994), Müller et al. (1995) 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, Müller and Found (1996) examined the relative contributions of the dimensional and feat-

ural 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 (Müller et al., 1995; Found & Müller, 1996) 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 2.4). In principle, there are 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. 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 weak-

ened. Consequently, responses to a color target on trial n to be 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. (2001) 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 (Müller et al. 1995; Found & Müller, 1996).

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. Bundesen, Habekost, & Kyllingsbaek, 2005), the lateral intraparietal area (e.g. Gottlieb, 2002), and the frontal eye fields (FEF; e.g. Thompson & Bichot, 2005; Bichot, Rossi, & Desimone, 2005; Bichot & Schall, 1999) – 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 (Fecteau & Munoz, 2006). In the following section, I 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 (Bruce, Goldberg, Bushnell, & Stanton, 1985; Robinson & Fuchs, 1969; Kastner et al., 2007). The featureless and feature–contrast dependent response characteristics of FEF neurons have been demonstrated by Sato, Murthy, Thompson, and Schall (2001). 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 distracters 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 signalling 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 (e.g., Wolfe, 1994; Müller et al. 1995; Itti & Koch, 2000) 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 (2002) observed a similar pattern using a pop-out search paradigm adapted from Maljkovic and Nakayama (1994), 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 tar-

get 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, Weidner, Müller, & Cramon, 2000) 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 (Pollmann, Weidner, Müller, Maertens, & Cramon, 2006; Pollmann, Weidner, Müller, & Cramon, 2006; Weidner, Pollmann, Müller, & Cramon, 2002) 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, Weidner, Müller, Maertens, & Cramon, 2006), 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 Kristjansson, 2007). 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 (2002). 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 (Treisman & Gelade, 1980; Itti & Koch, 2000; Müller et al. 1995; Wolfe, 1994). 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’ (e.g. Knierim & Essen, 1992). 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 (e.g. Krummenacher et al., 2001). 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, Thompson, Liu, Zhou, and Ault (1995) 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 (Johnson, Hawken, & Shapley, 2001; Friedman, Zhou, & Heydt, 2003; Gegenfurtner, Kiper, & Fenstemaker, 1996). 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 distracters, 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 (2007) 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 (Horwitz & Albright, 2005; Hubel & Wiesel, 1959; Ts'o & Gilbert, 1988), 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 to support of 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 (2007) null-result, Krummenacher and Müller (2007) found pop-out targets redundantly defined by a combination of color and motion to significantly violate the RMI (see also Katzner et al., 2006), 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 Leventhal et al., 1995, reported finding such cells, others, such as Horwitz and Albright, 2005, 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 (2002) and Koene and Zhaoping (2007). In summary, it is not yet possible to unequivocally decide between the summation saliency (Wolfe, 1994; Itti & Koch, 2000) and the V1 maximum models (Li, 2002), 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, I have summarized research on several critical questions concerning the nature of early visual processing and I have shown that 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 Itti & Koch, 2000, model). 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 pre–attentive 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 Müller et al., 1995) 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 (2002) requires further behavioral and neuro-physiological research to permit an unequivocal decision to be made between the alternative models.

3. CO-ACTIVATION VS. SERIAL AND PARALLEL MODELS

The redundant-signals effect (RSE) is a reaction time benefit for signals that simultaneously have two response related properties over signals that have only one. This reaction times benefit for the combination of two signals has often been reported in the literature, (e.g. Miller, 1982, 1986; Giray & Ulrich, 1993; Mordkoff & Yantis, 1991, 1993; Krummenacher et al., 2001, 2002; Turatto, Mazza, Savazzi, & Marzi, 2004; Katzner et al., 2006; Marzi et al., 1996; Corballis, 2002). These studies have different scopes and come from different areas of research: bimodal (e.g. visual and auditory) processing, divided attention between two locations or two features, visual pop-out search, or neuro-psychological research. Although coming from different fields and being used to answer different kinds of questions, the basic paradigm used for examining the RSE is very similar: The observer's task in this kind of paradigms basically is to react as soon as a stimulus that is element of a predefined set of target stimuli appears. Each target is mapped to the same response. Performance for trials in which redundant signals are presented (redundant signal trials, or RSTs) is then compared to the performance for trials in which only a single signal is presented (single signal trial or SST). The RSE is the reaction time benefit of RSTs over SSTs. Regardless of the specific field of interest or the specific nature of the two signals, the RSE can be accounted for in several mutually exclusive ways. Raab (1962) explained the RSE with statistical facilitation. He proposed that redundant targets are processed similar to a horse race: the signal that first can trigger the response, wins the race and determines the observed reaction time (see Figure 3.1¹). So on each RST the two single signals are processed in parallel by independent processors and build up activation that, if a threshold is exceeded, triggers the response. If the

¹Figures 3.1, 3.2, and 3.3 adapted from Schönwälder (2006), Townsend and Nozawa (1995), and Mordkoff and Yantis (1991).

processing times of both single signals in a RST are drawn from the two SST reaction time distributions, it is probable that for redundant targets there is always one faster and one slower processing time, leading to a mean reaction time gain - the RSE - because processing is terminated as soon as one channel detects a target. Formalized the minimum of two distributions (the processing times of both single signals in a RST) is less than each of both distributions. In summary, the RSE was explained by assuming an independent parallel race between the two signals in a RST, which leads to a reaction time distribution which is shifted to the left of the reaction time distribution of both SSTs.

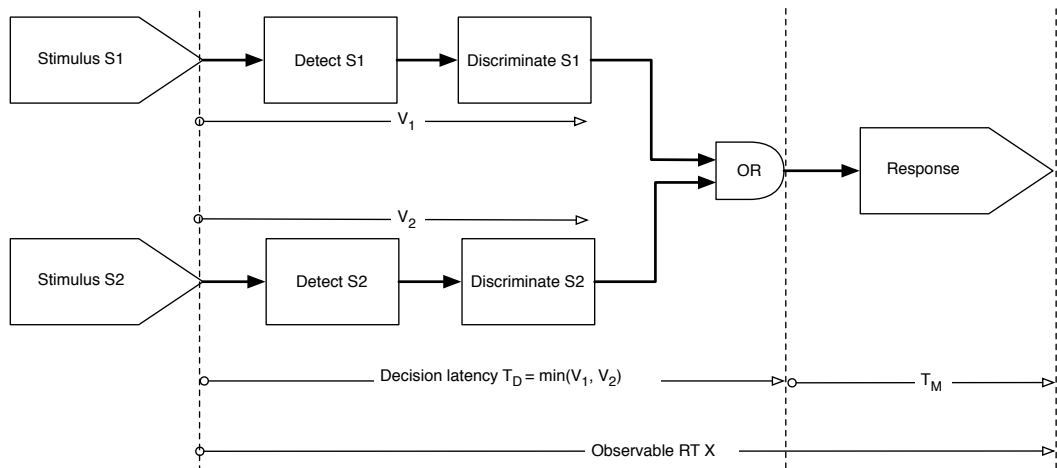


Figure 3.1. The observed reaction time in a race model is the sum of T_M , the time necessary for response and motor processing, and the minimum of the two detection times T_D for signals S_1 and S_2 . Each single signal has its own decision unit that can trigger the response.

Miller (1982) showed that this explanation of statistical facilitation in a parallel race model yields an upper bound to the RSE. If the redundancy gain exceeds a certain boundary, defined in the race model inequality (RMI), statistical facilitation of a parallel race cannot account for this gain anymore. Miller proposed a co-activation

model as an alternative explanation, in which the single signals of a RST are not processed in parallel, but are integrated before response processing. Figure 3.2 displays the architecture of a co-activation model. In a co-activation model the activation of both signals present in a RST are somehow pooled and lead to a faster triggering of the response, because there are two sources of activation for the response in a RST as compared to only one source of activation in a SST. In summary, Miller incorporated a taxonomy of cognitive processes based on the distinction between separate vs. integrated/co-active architectures of processing.

Thus the redundant signal paradigm is relevant for distinguishing between funda-

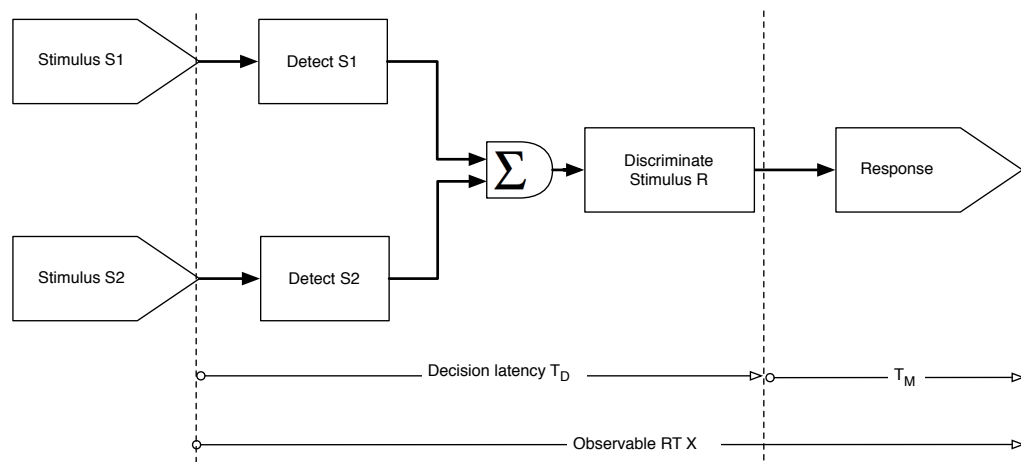


Figure 3.2. The observed reaction time in a co-activation model is the sum of T_M , the time necessary for response and motor processing, and the decision latency T_D . Both single signals are fed into a common decision unit.

mentally different modi of processing - parallel and independent versus co-active - which are relevant for a large set of questions: are signals from two modalities (e.g. auditory and visual) processed in parallel or are they being integrated? If attention is divided between two locations or two features of one objects, are signals processed separately or are they integrated?

The redundant signal paradigm has also been employed to infer about neural processing structures (Turatto et al., 2004; Marzi et al., 1996; Corballis, 2002).

The paradigm for examining the RSE as relevant for the present study is visual feature search (Krummenacher et al., 2001, 2002). Targets in pop-out search are feature singletons, i.e. the target is the only item in the display that contains a certain feature (e.g. a red among green bars). Thus feature contrast between target and distracters is much larger than between distracters themselves. Redundancy in pop-out search can be introduced in two ways (a) by creating search displays in which two separate items differ from all distracters (e.g. two red targets among green distracters: Krummenacher et al., 2002), or (b) the single target differing from the distracters in two dimensions (e.g. a red horizontal among green vertical bars: Krummenacher et al., 2001). The questions that have been targeted by these two studies are related to the dimension weighting account (DWA; Müller et al., 1995; Found & Müller, 1996). The DWA is a model of visual search that proposes parallel and independent initial processing of the visual display in a way similar to e.g. Guided Search (Wolfe, 1994) or the computational models of Itti and Koch (2000) and Koch and Ullman (1985). It assumes that initially feature contrast is calculated over the whole display. That is for each location of the display a signal is calculated, which indicates, how different this location is from its surrounding in terms of a specific color (e.g. red or blue), of a specific orientation (e.g. horizontal or vertical), or of a specific motion direction. This calculation of feature contrast happens in parallel for all features. Feature contrast is represented in topographic maps, on which neighboring locations are representing neighboring locations of the visual display. Feature contrast signals are first pooled in a dimension specific way, e.g. all color signals are integrated in a color contrast map, and secondly the dimension specific contrast signals are pooled into a saliency master map of locations. Activity at a specific location on this maps indicates that this location is different in some feature(s) on the visual display, without providing information in what way it is different (in what dimension or feature). Attention is then assumed to select the location with the highest activation on the master map

for further processing.

The crucial assumption of the DWA is that signals from the dimension maps are modulated with dimensional weights before they are summed into the master map. These weights can be changed in two ways: (i) by bottom-up, stimulus based, passive processes (e.g. Müller et al., 1995), as well as (ii) by intention (i.e. top-down, e.g. Müller et al., 2003; Müller & Krummenacher, 2006). Regardless of how the weights are changed, it is not possible to increase the weight of one dimension and leave the other weights as they are, i.e. the weights can be understood as a limited resource. The DWA applied to e.g. a search for a red pop-out target amongst green distracters predicts that activity on the master map will build up faster, if the weight of the color dimension is high, than if it is low. This faster build-up of activity on the master map based on higher weights for the target dimension is responsible for faster search times in that case.

Krummenacher et al. (2001, 2002) employed the redundant signal paradigm in order to further specify the DWA. They asked whether integration into the master saliency map happens in a pre-attentive or post-selective manner, whether the integration is spatially specific or unspecific, and whether integration occurs for dimension-based or feature-based signals. For all these questions the implicit assumption of the DWA is that the master map acts as a summator of dimensional contrast signals, i.e. the assumption is that a co-activation-model underlays processing. The question then was not so much whether co-activation was the underlying model at all, but under what circumstances co-active processing happens, and under which it does not happen.

The authors observed violations of the RMI for single targets, redundantly defined in two dimensions, as well as for two pop-out targets (dual targets). Dual targets lead to violations of the RMI only if they were defined in different dimensions (e.g. a red vertical and a green horizontal among green vertical bars) but not if they were defined as two features of one dimension (e.g. a green left tilted and a green right tilted among green vertical bars). If defined in two dimensions, the distance between both pop-out targets was also a relevant variable for integration: dual targets only lead

to violations of the RMI, if they are in close spatial proximity. This speaks for the dimensional specificity of the DWA as opposed to feature based models, as well as for a spatially limited integration process. Additionally, violations of the RMI occurred at precued (likely) and uncued (unlikely) locations, which hints to the preattentive interpretation of the DWA.

Thus in these studies the main question was not whether a co-activation model can account for the RSE at all, but under what circumstances does and does not co-active processing take place. The theoretical assumption of a saliency master map, into which dimensional contrast signals are summed is a co-activation model and has been proposed frequently in the literature (Koch & Ullman, 1985; Wolfe, 1994; Müller et al., 1995; Itti & Koch, 2000).

Still, this theoretical assumption of co-active summation of dimensional signals can be put into question itself. Then the RMI in isolation can not be decisive in that question: Townsend and Nozawa (1997) have shown that there is an additional class of models that can lead to violations of the RMI: serial models with an exhaustive search rule (i.e. search is only terminated, when all items have been processed, even if the target has been detected before the last item). An example for such a serial model with an exhaustive search would be the strategy of participants to check both channels in every trial. This double-checking would lead to an increase of reaction times in SSTs, if the time necessary for determining the absence is longer than the time for determining the presence of a signal. Mordkoff and Yantis (1991) introduced a variant of the parallel race model, in which channels may exchange information before a response is selected: the interactive race model. They showed that if there are contingencies between the channels that favor redundant over single targets, then violations of the RMI may occur, which are not due to a co-activation model. As a result the RMI as a solitary tool loses some of its decisive power. Though violations of the RMI indeed exclude independent parallel models, serial exhaustive and interactive race models from may have to count as an alternative explanation, in addition

to co-activation models.

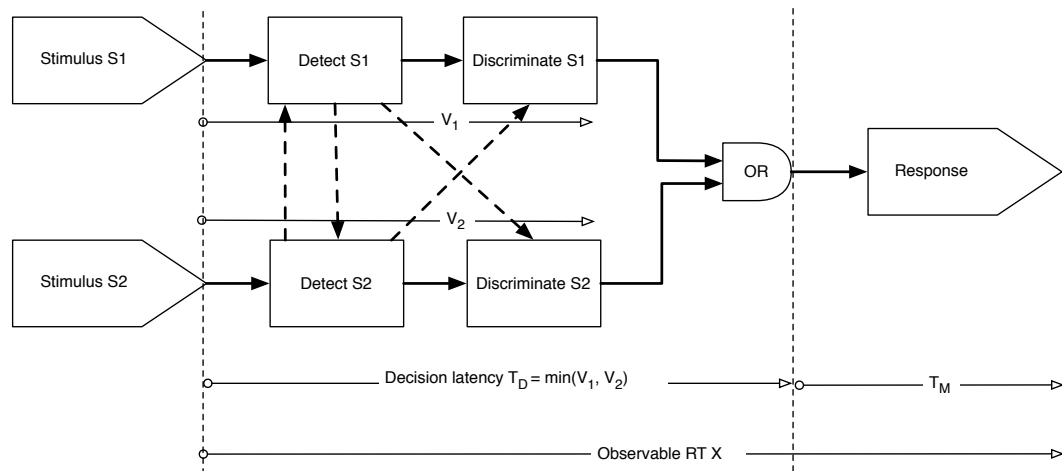


Figure 3.3. The interactive race model is similar to the parallel race model (cf. Figure 3.1) with the addition that both processing channels can exchange information e.g. about presence or absence of a signal in the respective channel to the other channel. Two possible routes for exchange of information are denoted by the dashed arrows.

Townsend and Nozawa (1995) set Miller's taxonomy of separate vs. co-active processing in the context of notions of architecture (serial or. parallel), capacity (unlimited, limited, or super-capacity), and stopping rules (exhaustive or self-terminating). They showed that if there are local (i.e. for some reaction times) violations of the RMI, the system is also locally super capacity. If a system is super capacity for all reaction times, the RMI will be violated at least locally. With measures stemming from systems factorial methodology (Sternberg, 1969b; Townsend & Ashby, 1983), Townsend and Nozawa (1995) make it possible to distinguish between different architectures and stopping rules. Using a double factorial design in combination with a redundant signal paradigm, a thorough analysis of underlying cognitive architecture is possible. If processing of visual search displays can be explained by interactive race models, then manipulations of contingencies that favor redundant over single targets

should affect the size of observed redundancy gains and violations of the RMI.

In the present study I employ a double factorial redundant target paradigm and manipulate the contingencies favoring redundant over single targets in two experiments. Participants perform a visual search detection task, in which a pop-out target can be either present or absent. The pop-out target is defined in either one dimension (orientation or luminance) or in both. That is the two signals of the redundant signal paradigm, which are mapped to the same (target-present) response are feature contrast in the orientation and in the luminance dimension. We ask what cognitive model underlies this kind of processing, in terms of architecture (serial, parallel), capacity (limited, unlimited, super capacity), stopping rule (exhaustive, self-terminating) and separate or integrated/co-active, as well as interactive processing. For each channel the two factors of the double factorial design are absence/presence and intensity. Intensity is varied in the terms of orientation contrast (e.g. 6° or 45° tilt between targets and distracters) and luminance contrast between targets and distracters. In Experiment 1 the double factorial design is implemented with the goal of determining whether serial exhaustive models - and not co-activation models - can in fact explain the RSE in visual search pop out tasks. Experiment 2 aims at verifying that targets of low intensity used in Experiment 1 lead to efficient search (i.e. slope of reaction time with respect to set size in a display is small, e.g. less than 5 ms). In Experiment 3 I manipulate the ratio of present and absent targets (1:1 vs. 3:1), as well as the ratio of single and redundant targets (1:1:1 vs. 1:1:2 for orientation:luminance:redundant), and thereby the amount of benefit for redundant targets over single targets if crosstalk between channels were possible. Our hypothesis, according to the DWA is that the RSE in visual pop-out search can be only accounted for by co-activation, neither by exhaustive serial, nor by interactive race models.

3.1 Experiment 1

Experiment 1 employs a double factorial design combined with a redundant target paradigm as proposed by Townsend and Nozawa (1995). The two visual pop-out dimensions of orientation and luminance are regarded as channels, and a signal is either present or absent, or of low or high intensity in each channel. In RSTs the varied intensity allows analyses and model tests based on systems-factorial technology: mean interaction contrast. In combination with capacity analysis and the RMI Experiment 1 aims at distinguishing between serial exhaustive and co-activation models as an explanation of the RSE. Our hypothesis based on the DWA is that serial exhaustive models, as well as parallel race models can be excluded, and co-activation models can explain the RSE.

The double factorial design is derived from Sternberg's (1969) 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. Strength of feature contrast can be manipulated by varying the similarity between targets and distracters (Duncan & Humphreys, 1989). Townsend and Nozawa (1995) proved that looking at 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 targets in such a 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 a tilt 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 low feature contrast. If lowering feature contrast in one dimension has a larger effect when the feature in second dimension is of low contrast than when contrast is high, then super-additivity is said to occur. Townsend and Nozawa (1995) 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.² Based on the DWA (Müller et al., 1995; Found & Müller, 1996) I predict violations of the RMI, and a super additive interaction contrast that is evidence for a summation of dimensional signals that lead to the reaction time benefit of redundant over single signals in visual pop-out search.

3.1.1 Method

Participants

16 observers took part in Experiment 1 (5 male, 1 left handed). Age ranged from 19 to 46 years (median: 24.5). Observers were paid with a rate of 8 Euro per hour.

Apparatus

Observers viewed the stimuli in front of a Sony Multiscan E250 17" monitor driven by personal computers with Windows XP operating system. The personal computer was placed in a sound isolated room with black interior. There was dim background light in order to prevent reflections on the monitor. Viewing distance was about 62 cm and observers were instructed to maintain constant distance to the monitor. The screen refresh rate was 85 Hz, the screen resolution was set to 1024x768 pixels.

²The 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.

Participants responded by pressing the right button of a mouse with the index finger of their right hand. Reaction times were recorded online by the computer. After each block the participants were informed about their mean reaction time and error rate of the previous block.

Stimuli and Timing

The display consisted of a 6x6 array of filled upright rectangles (bars) on a black background with 0.6 cd/m^2 . The bars were either dark (11.6 cd/m^2) or light grey; the luminance of the light grey bars were adjusted individually for each participant and ranged from 43.6 to 97.2 cd/m^2 . The bars subtended approximately 1.7° of visual angle of height and 0.35° of width. They were arranged with a vertical and horizontal distance between each other of approximately 1.6° with a jitter of about 0.2° . There were four single and four redundant target conditions. Orientation targets could differ from the vertical distracters by a tilt to the left or to the right of 6° or 45° . Luminance targets could be more or less brighter than the dark grey distracters. An adaptation phase previous to the main experiment was used in order to find two intensity levels of luminance targets, for which reaction time performance was comparable to the high and low feature contrast orientation targets. Targets were always placed in such a way that they were surrounded on all sides by distracters. They were never placed at the border of the array, i.e. they were presented in the inner 4x4 array only. Thus all targets were surrounded by eight distracters. Participants were not informed about this restriction. A target was present in 60% of all trials.

Trials started with simultaneous onset of all stimuli. The stimuli were presented until the observer responded, but no longer than 1000 ms. After 1000 ms the next trial started with a new intertrial interval. On target present trials the correct response was a right mouse button press. On target absent trials the correct action was to wait for 1 s until the next trial started. The intertrial interval was 1000 ms with a temporal jitter of 200 ms. On erroneous trials the interstimulus interval was increased

by 2 s as a feedback signal. In the practice and adaptation phase of the experiment on target present trials there were an equal number of orientation and luminance trials. In the main phase of the experiment for each SST a coin was tossed to determine its dimension. There were two times as many SSTs as RSTs.

Design and Procedure

Experiment 1 consisted of two sessions of one hour each. Both sessions started with a practice block of 30 trials, which were not included in the analysis. All other blocks consisted of 60 trials each. In session 1 the saliency of luminance targets was adapted to the saliency of orientation targets for both intensities in the six blocks after the practice block. The adaptation phase of the experiment aimed at adjusting the brightness of luminance targets in such a way that reaction times for both target types (orientation and luminance) had statistically equal means. The first half of the adaptation phase consisted of only target absent and orientation target trials. During these trials the presentation computer calculated the median reaction time for orientation targets. In the second phase of the adaptation phase, brightness of luminance targets was adjusted using an adaptive staircase procedure (after Johnston, Cumming, & Parker, 1993). In each target present trial the reaction time for a luminance target was compared to the median reaction time for orientation targets. If the reaction time was faster than this median, the luminance of the upcoming target was decreased, if it was slower, the luminance was increased. Luminance was controlled with 6 bit RGB values, and step size decreased from 8 to 1 with each reversal of the staircase. For the rest of each experiment brightness of luminance targets was kept constant at the value generated by the adaptation procedure. The task (go/no-go detection) and target present absent ratio (60:40) was kept constant during all blocks of the experiments. In summary, each session consisted of three subparts: a practice block, several adaptation blocks, and the main part, in which the experimental manipulations were presented.

3.1.2 Data Analysis

Data analysis involves three parts: (a) mean reaction time performance, (b) reaction time distribution measures, and (c) statistical testing of the distribution measures. Distribution analysis will evaluate the violations of the RMI and capacity. In the following subsections I will describe in more detail the methods employed for data analysis. For the all reaction time analysis, target absent trials, the first 4 trials of each block, and trials with errors have been excluded.

Mean Reaction Time Analysis

Mean reaction times analysis consisted of repeated measures analysis of variance (ANOVA) with different within-subject factors (target dimension and intensity). Additionally in Experiment 1 the mean interaction contrast was calculated, taking into account only RSTs:

$$\overline{RT}(l, l) - \overline{RT}(l, h) - \overline{RT}(h, l) + \overline{RT}(h, h), \quad (3.1)$$

where 'l' stands for low (i.e. a dim luminance or a 6° orientation target), and 'h' for high intensity (i.e. bright luminance or a 45° orientation target). $\overline{RT}(h, l)$ thus denotes the mean reaction time to redundant targets, defined with a high level of intensity in the orientation, and a low level of intensity in the luminance dimension. The RSE is the RT benefit of redundant targets relative to the corresponding single targets (e.g. high and low intensity orientation and luminance single targets for a redundant target with 45° tilt and low lumiannce contrast). In order not to overestimate the redundancy gains, we adapted a procedure by Miller and Lopes(1988, see also Krummenacher et al. (2001, 2002)). Miller and Lopes (1988) proposed a method to determine for each participant, whether one of the single targets is favoured using a two-sided t-test with an alpha level of 0.1. If no single target dimension is favoured, the mean RTs for both single target dimensions are averaged; if one of the single tar-

get dimensions is favoured, the mean RT of the favoured dimension (which is faster) is compared to RSTs. We adapted the test for the analysis in Experiment 1. For each participant, the mean RT of the faster single signal was compared to the mean RTs of redundant signals (i.e. the procedure proposed by Miller and Lopes is used with an alpha level of 1).

The Race Model Inequality

Miller (1982) formalized the RMI as follows:

$$P(RT < t|S_{12}) \leq P(RT < t|S_1) + P(RT < t|S_2), \quad (3.2)$$

where S_i denotes the channel i for a SST, and S_{ij} denotes a RST presented in channels i and j . So the RMI basically states that the probability that RTs on a RST that is smaller than some time t is always less or equal than the sum or the respective probabilities of both SSTs. This provides an upper bound of how much RSTs can benefit from the simultaneous presentation of S_1 and S_2 under the assumption of a race model. The RMI is usually evaluated by calculating the cumulative density functions (CDFs) of reaction times on RSTs and of both SSTs. The sum of both single CDFs from the CDF of RSTs. If this difference is less than zero, the observed RSE is in accordance with a race model. Remember that in a parallel race model (like in Figure 3.1) the two components of a redundant signal are processed in parallel, until one of both signal triggers the response. Due to random variation the processing times of both signals vary from trial to trial. That way on each trial one signal of both that are processed in parallel reaches the response trigger faster than the other.

Testing the predictions of the RMI usually involves the calculation of the cumulative density functions (CDFs) of RTs obtained in RSTs and in SSTs (separately for each type of single signal). The sum of the two SST CDFs is then related to the CDF of RSTs. If the difference between the cumulative probabilities is smaller than zero

any observed RSE is in accordance with the race model assumption. To deal with the unintuitive property of the sum of two single signal CDFs to be 2 (instead of 1 as in density functions) we adopt an alternative formulation of the RMI proposed by Colonius and Diederich (2006). They demonstrated that the minimum of the sum of the two CDFs and 1, $\min(1, P(RT < t|S_1) + P(RT < t|S_2))$, is also a density function. If the CDFs are given by $P(RT < t|S_1) = G_1(t)$, $P(RT < t|S_2) = G_2(t)$, and $P(RT < t|S_{12}) = F(t)$, the null hypothesis of the parallel race model is

$$H_0 : d(t) = F(t) - \min(1, G_1(t) + G_2(t)) \leq 0, \quad (3.3)$$

where $d(t)$ denotes the Kolmogorov distance of the two distributions $F(t)$ (i.e. the distribution stemming from RSTs), and $\min(1, G_1(t) + G_2(t))$ (i.e. the distribution which is the maximal possible benefit of redundant signals over single signals under assumption of the race model). The race model predicts that this distance $d(t)$ is less than zero.

This formulation of the RMI also allows to quantify the amount of violations of the RMI (Colonius & Diederich, 2006), which is relevant for evaluating whether there are more or less violations of the RMI in two or more experimental conditions. In short, Colonius and Diederich (2006) suggested to use the area under $d(t)$ as a measure of the amount of violations of the RMI. In order to test, whether $d(t)$ significantly differs from zero for a particular point in time t , we employed the method of vincentizing (e.g. Miller, 1982; Kiesel, Miller, & Ulrich, in press). The group distribution is calculated by evaluating $d(t)$ for each observer at a defined number of quantiles. For each quantile $d(t)$ is evaluated for each observer. The statistical significance of potential violations is examined using t-tests for each quantile q_i . Methodologically, however, the procedure has the drawback that the values of successive quantiles of the vincentized distribution are not independent. The dependence of data points may result in an overestimation of RMI violations since, given that at particular quantile the RMI is violated, violations at neighbouring points are likely to occur as well

(Zandt, 2002; Kiesel et al., in press). Artificial overestimations of RMI violations can be reduced if a large number (more than 20) of observations is used to estimate the ECDF, and if the inequality is tested within a limited range of quantiles, e.g. ranging from 0.05 to 0.20 (Kiesel et al., 2008). In Experiment 1 of the present study the number of observations per condition ranged from 94 to 205 with a median of 141.1. To agree with the second limitation, the range of quantiles where the RMI was tested was restricted to four (0.05 to 0.2).

3.1.3 Results

Trials with reaction times of less than 200 ms were considered anticipations (< 0.1% of all trials). Overall error rates were low: with 2.5% false alarms and 2.6% misses.

First, we analyzed only SST data, in order to test whether the manipulation of feature contrast (which was necessary for the double-factorial design) was successful. The SST data were subjected to a 2 x 2 repeated measures ANOVA with within-subject factors dimension (orientation, luminance) and feature contrast (low, high). The analysis revealed a significant main effect of feature contrast ($F(1,13) = 82.5$, $p < .001$), with targets of low feature contrast (477.3 ms) being processed considerably slower than targets of high feature contrast (371.5 ms). The RT difference between orientation (417.5 ms) and luminance (431.2 ms) was not significant ($F(1,13) < 2.4$), nor was the interaction between the two factors ($F(1,13) < 0.3$). In order to test, whether low feature contrast targets did still pop-out, we performed a set size experiment (see Appendix) and found low feature contrast targets to have an effective search function (with slope $\leq 4\text{ms/item}$). Thus, we can assume that both targets of high and low feature contrast are processed in the same fashion and the intensity manipulation necessary for the double-factorial design was successful. The staircase adaptation of the RGB values of luminance targets in the adaptation phase of Experiment 1 also was successful: it yielded a significant difference between the levels of

feature contrast while at the same time no significant difference between dimensions was created.

The RT data of RSTs (45° /bright, 45° /dim, 6° /bright, and 6° /dim) of Experiment 1 were subjected to a 2x2 repeated measures ANOVA with the within-subject factors of representing the intensity variation of the double-factorial design; orientation (6° , 45°) and luminance (dim vs. bright) of the redundantly defined signal. Both main effects were significant: orientation-intensity ($F(1,13) = 150.4$, $p < .001$) and luminance-intensity ($F(1,13) = 21.53$, $p < .001$). Redundant targets with a low-intensity component were processed slower (392.8 ms and 383.6 ms for the orientation and luminance component, respectively) than redundant targets with a high-intensity component (355.3 ms and 364.3 ms for the orientation and luminance component, respectively). Further, the interaction between the factors was significant ($F(1,13) = 33.50$, $p < .001$). Redundant target trials with two high-intensity components were responded to fastest (353.9 ms), followed by redundant targets with only one high-intensity component (356.4 ms and 375.4 ms for high intensity in the luminance and orientation component, respectively). Redundant targets with two low-intensity components however were responded to much slower (410.6 ms). In order to quantify that interaction, the mean interaction contrast (see Eq. 1) was calculated for each participant and subjected to a two-tailed one-sample t-test. The t-test revealed the mean interaction contrast of 32.8 ms to be significantly greater than zero, $t(13) = 5.8$, $p < .001$, indicating an over-additive interaction of the factor intensity for both dimensional components of redundant targets. The significant over-additive interaction rules out both any type of serial and exhaustive parallel processing of stimulus signals of the different dimensions.

The mean RSE for each combination of intensities for the orientation and luminance dimensions (45° /bright, 45° /dim, 6° /bright, and 6° /dim) were 8.3, 5.7, 5.6, and 44.1 ms, respectively. The RSE is significantly greater than zero ($t(13) > 2.52$) in all conditions. In order to test whether statistical facilitation in parallel models can actually explain the observed RSE, violations of the RMI were analyzed. Figure 3.4 presents

$d(t)$ for all four conditions (45°/bright, 45°/dim, 6°/bright, and 6°/dim) averaged for each 5%-quantile of the distribution of RSTs. Values of $d(t)$ greater than zero are in discrepancy to the assumption of parallel race architectures. In order to avoid generating artificial violations of the RMI (Kiesel et al., 2008), only the four quantiles between 0.05 until 0.20 were tested. In the condition with low intensities in both dimensions (6°/dim condition), the RMI in the range of 0.10 to 0.20 is significantly violated (all $t(13) > 2.0$, $p < .03$). A significant violation was found in the 45°/bright condition (at the 0.05 quantile, $t(13) = 1.8$, $p < .05$).

In summary, we found violations of the RMI in both conditions, where the intensity-levels of both dimensional components are equal (45°/bright, 6°/low). In the other cases, one dimension can be processed much faster than the other, and it thus is likely that co-activation, although in principle possible, does not take place. This finding suggests that parallel race models of processing of dimensional signals can be ruled out.




Figure 3.4. The estimate $d(t)$ for all participants and combinations of intensity.

3.1.4 Discussion

Experiment 1 implemented a double factorial design combined with a redundant signals paradigm (after Townsend & Nozawa, 1995), in order to determine architecture (serial, parallel, or co-active), stopping rules (self-terminating or exhaustive) and capacity of the underlying model of visual pop-out search. We manipulated the intensity of feature contrast in two dimensions: orientation (tilted 6° or 45° relative to vertical distracters) and luminance (dim vs. bright) of targets, which could also occur redundantly defined in both dimensions. The manipulation of feature contrast lead to a reaction time difference of approximately 100 ms between targets of high and low intensity. All four types of redundant targets (6° -dim, 6° -bright, 45° -dim, and 45° -bright) were found faster than the fastest single target (of the respective combination of dimension and intensity).

The mean interaction contrast of the intensity levels in redundant targets was significantly positive (approximately 30 ms), excluding parallel exhaustive models (which would have predicted an additive effect without any interaction), as well as serial models of any stopping rule (which would have predicted no significant interaction). Hence serial exhaustive models are not a valid alternative for explaining the observed gains. A super additive interaction contrast can be observed for parallel race models with a self-terminating search, as well as for co-activation models. Significant violations of the RMI occurred over a wide range of quantiles in the 6° -dim condition, and a tendency of violation was found in the 45° -bright condition. Taking into account only those participants who exhibited a significant redundancy gain, additional significant violations of the RMI were observed in the mixed intensity conditions (6° -bright and 45° -dim). Violations of the RMI exclude parallel race models as an explanation of the observed redundancy gains.

Thus exclusion of parallel race models as well as serial exhaustive models accumulate evidence that indeed a co-activation model underlays processing of pop-out targets in visual search, as assumed by e.g. the DWA. One possible objection may be that dim

luminance or 6° orientation targets are not pop-out targets at all, and that the observed effects are not valid for parallel pop-out search, but for inefficient serial search (i.e. items have to be serially looked at).

3.2 Experiment 2

Experiment 2 is a set size experiment with pop-out targets of low intensity and aims at verifying that the low intensity targets used in Experiment 1 actually were pop-out targets, i.e. were found efficiently. Orientation targets vary from distracters in 6° , and luminance targets are presented with such an intensity increase in comparison to distracters that they yield similar reaction times to orientation targets. The question is, whether targets defined at low intensity pop out, i.e. have search slopes of less than 5 ms/item.

3.2.1 Method

Participants

Eight observers took part in Experiment 2 (3 male, 1 left handed), with age ranging from 19 to 29 (median 22.5). Observers were paid 4 Euro for one session of half an hour.

Apparatus

The apparatus was the same as in Experiment 1.

Stimuli and Timing

Stimuli were grey bars as in Experiment 1. Stimuli were presented with set sizes 36, 25, and 16 (in 6x6, 5x5, and 4x4 arrays, respectively).

Design and Procedure

Participants received a written instruction of the task and response requirements. The experiment consisted of 10 blocks with 60 trials each and lasted about half an hour. In half of the trials no target was present. Target present trials consisted of 50% orientation and 50% luminance targets. The session for each observer consisted of one practice block, two blocks, in which the intensity of luminance targets was adjusted in order to produce reaction times as fast as for the 6° orientation targets, and seven blocks in which the set size varied randomly from trial to trial. In the second block only orientation targets were present and the median of target present response times was calculated by the stimulus presentation program. In the third block target present trials were luminance only. In this block the intensity of luminance was adjusted with an adaptive staircase procedure.

During the last 7 blocks the luminance was kept constant with the last value the staircase produced, and set size was randomly varied between 36, 25, and 16 stimuli.

Data Analysis

In Experiment 2, intercept and slope of the reaction time/set size functions are estimated with linear regression for each participant. Both intercept and slope are then averaged over all participants and target types (orientation and luminance).

3.2.2 Results and Discussion

Figure 3.5 presents the RT results for luminance and orientation targets. Search slopes are 3.1 *ms/item* (luminance) and 1.7 *ms/item* (orientation) with intercepts of 403.1 ms and 407.2 ms respectively. An analysis of variance (ANOVA) with factors target type (luminance and orientation) and set size (16, 25, and 36 items) revealed both main effects of dimension ($F(1, 7) = 5.68, p < .049$) and set size ($F(2, 14) = 16.3, p < .001$), but not the interaction of both, to be significant. Two separate paired

t-tests on intercepts and slopes of the set size function for both dimensions revealed no significant effects: $t(7) = -1.8, p < 0.12$ (slope), and $t(7) = 0.21, p < .8$ (intercept). The shallow search slopes of below 5 ms show that targets of low intensity are searched in an efficient way. Thus an orientation contrast between targets and distracters of 6° and a corresponding luminance contrast can be used as a manipulation of intensity, while still leading to efficient visual pop-out search.




Figure 3.5. Mean reaction times with standard error with respect to set size for both dimensions. The search slope for orientation is 1.7ms/item, and for luminance 3.1ms/item.

3.3 Experiment 3

Experiment 3 aims at investigating, whether interactive race models can account for violations of the RMI in redundant pop-out search. Experiment 1 confirmed that violations of the RMI observed earlier (Krummenacher et al., 2001, 2002) cannot be explained with a serial exhaustive model. If a serial exhaustive model would be

responsible for violations of the RMI, the mean interaction contrast (see Eq. 3.1) would have predicted to be zero (proposition 4 in Townsend & Nozawa, 1995), but it was significantly greater than zero. Still, an interactive race model (Mordkoff & Yantis, 1991, 1993) would also predict violations of the RMI, if there were inter-channel contingencies that favor RSTs over SSTs. The basic additional assumption that distinguishes interactive race models from independent race models (e.g. Raab, 1962), is that perceptual channels may exchange information. In an interactive race model, both processing channels can exchange information about e.g. presence or absence of a target before selection of a response (i.e. within the time T_D in Figure 3.3). For example if the display contains an orientation target, information about the presence of feature contrast signals in the orientation channel could be made available to the luminance channel. This feature distinguishes the interactive race model from co-activation models: not only unspecific activation is pooled, as in co-activation models, but specific information about e.g. presence or absence of targets can be exchanged in the perceptual stage, before response selection. This information would have no effect, if the presence or absence of feature contrast signals in one channel are uncorrelated to presence or absence of feature contrast in the other channel. In the following section I will present quantification of two types of contingencies that could lead to a benefit of redundant over single target trials, as described by Mordkoff and Yantis (1991).³

In order to calculate interstimulus contingencies that can favor RSTs over SSTs it is necessary to recapitulate, how displays are analyzed in a parallel architecture (with or without crosstalk). Targets in pop-out search are defined by local feature contrast, i.e. locations, at which some features are different from the surrounding locations. At regions, where only distracters are presented, feature contrast is low, because all neighboring items share the same features. At the location of a pop-out target, e.g. a red vertical among green vertical bars, feature contrast for red, and consequently

³Mordkoff and Yantis (1991) quantified contingencies for a paradigm in which a target could appear at either one or two locations. We will present the contingencies for the situation of the redundant pop-out paradigm.

feature contrast for color is high. In a redundant-signals paradigm displays can contain no target, a target defined in one dimension (e.g. orientation), a target defined in another dimension (e.g. luminance), or a target defined in both dimensions. In a parallel architecture, there are two processors, one for each of the possible dimensions. If there is no target presented, both channels signal absence of feature contrast. If there is e.g. an orientation target in the display, one processor signals presence of feature contrast, whereas the other channel signals absence of feature contrast.

If I denote a display that produces a feature contrast signal in the orientation dimension (i.e. an orientation or a redundant target) as T^O , and a display that produces no feature contrast signal in the orientation dimension as N^O , the conditional probability that a display produces a feature contrast signal in the orientation dimension, but not in the luminance dimension is given by $P(T^O|N^L)$. If the probability of a display with feature contrast in the orientation dimension, under the condition that there is no feature contrast in the luminance dimension, is higher than the probability of a display containing orientation feature contrast (i.e. $P(T^O|N^L) < P(T^O)$), then crosstalk from the luminance channel to the orientation channel facilitates the identification of a target. Presented in Table 3.1 are the frequencies of all four target types of Experiment 1. It is clear that $P(T^O) = 0.2$, which is less than $P(T^O|N^L) = 0.33$. So the information that no feature contrast is present in the luminance channel is beneficial for detection of feature contrast in the orientation channel. Conversely, if $P(T^O|N^L)$ would be greater than $P(T^O)$, then crosstalk between both channels would inhibit detection of feature contrast in the orientation channel.

Mordkoff and Yantis quantified this relationship as the interstimulus contingency (ISC):

$$ISC(N \Rightarrow T) = P(T^O|N^L) - P(T^O),^4 \quad (3.4)$$

where N stands for a target absent display, and T for a target present display. If $ISC(N \Rightarrow T)$ is positive, crosstalk facilitates target detection, if it is negative, target

⁴In this and the following equations, O and L can be exchanged. For simplicity only one of both cases is presented, because both would yield the same results.

detection is inhibited.

In a similar fashion there can be contingencies between channels that both contain targets:

$$ISC(T \Rightarrow T) = P(T^O|T^L) - P(T^O), \quad (3.5)$$

where $P(T^O|T^L)$ is the probability that feature contrast is present in the orientation dimension, under the condition that there is feature contrast in the luminance condition.

If the benefit for target detection under the condition that target presence has been detected in the other channel ($ISC(T \Rightarrow T)$) is greater than under the condition that target absence has been determined in the other channel ($ISC(N \Rightarrow T)$), the interstimulus contingencies favour redundant over single signal trials. In other words SSTs benefit from $ISC(N \Rightarrow T)$, and RSTs benefit from $ISC(T \Rightarrow T)$. If $ISC(T \Rightarrow T) > ISC(N \Rightarrow T)$, then RSTs have an advantage over SSTs. This benefit of interstimulus contingencies that favor redundant over single signal trials can be quantified as:

$$ISCB(N) = ISC(T \Rightarrow T) - ISC(N \Rightarrow T). \quad (3.6)$$

Mordkoff and Yantis (1991) introduced a second possibility, how crosstalk between channels could lead to a benefit of RSTs over SSTs: Nontarget-response contingencies (NRC). Here the probability of a 'present'-response is compared to the probability of a 'present'-response given that one channel has already determined that no feature contrast is present. If that conditional probability is higher than the baseline probability of a 'present'-response, than interchannel crosstalk could facilitate detection of a target. If the probability of a 'present'-response is greater than the probability of a 'present'-response under the condition that absence of feature contrast in one dimension has been detected, than SSTs become inhibited, whereas RSTs would not suffer from such a inhibition. The benefit of RSTs over SSTs due to NRC can be formalized as:

$$NRCB(N^O) = P(+)-P(+|N^O), \quad (3.7)$$

where $P(+)$ is the probability of a 'present'-response, and $P(+|N^O)$ is the probability of a 'present'-response under the condition that no feature contrast has been detected in the orientation dimension.

Mordkoff and Yantis (1991) showed that in a divided attention paradigm, in which target letters could appear at one of two or at both locations, contingencies were responsible for violations of the RMI: in situations, where there were either interstimulus, nontarget-response, or both types of contingencies, violations of the RMI could be observed. No violations of the RMI occurred, when contingencies were zero. In order to have zero contingencies, Mordkoff and Yantis introduced two non-target items, which made it possible to manipulate both types of contingencies independently.

In a visual pop-out search, there is only one type of non-target: a display that contains only distracters. Thus it is impossible to create a situation, in which there are zero contingencies in a visual pop-out search. In order to test whether an interactive race model is responsible for violations of the RMI in pop-out search, I manipulated strength of contingencies. If indeed pop-out targets are processed in an architecture of an interactive race model, redundancy gains and violations of the RMI should be greater, the stronger the contingencies are. We manipulated the ratio of target-present vs. target-absent trials (1:1 vs 3:1, target:absent) and the ratio of redundant vs. single targets if a target was present (1:1:1 vs. 1:1:2, orientation:luminance:redundant). This led to four different combinations of present-absent and single-redundant ratios, which are presented in Tables 3.2 to 3.6. ISCB, NRC and the sum of both is presented in Table 3.7 for the four possible conditions. The four conditions were varied as a between-subjects factor.

If an interactive race model were responsible for redundancy gains and violations of the RMI, than those conditions with greater contingencies should result in greater redundancy gains and violations of the RMI compared to those conditions with lower contingencies. We predict based on the DWA (Müller et al., 1995) that redundancy gains and violations of the RMI are insensitive to manipulations of the contingen-

cies, because I assume violations of the RMI to be due to summation of different dimensional feature contrast signals into a master map of saliency.

Table 3.1 Occurrence of target types in Experiment 1 per 120 trials with a present:absent ratio of 60:40 and a orientation:luminance:redundant ratio of 1:1:1.

Orientation	Luminance	
	T^L	N^L
T^O	24	24
N^O	24	48

Table 3.2 Occurrence of target types in Experiment 3 per 120 trials of condition 1, with a present:absent ratio of 1:1 and a orientation:luminance:redundant ratio of 1:1:1. ISCB is 0.25 and NRC is 0.25.

Orientation	Luminance	
	T^L	N^L
T^O	20	20
N^O	20	60

Table 3.3 Occurrence of target types in Experiment 3 per 120 trials of condition 2, with a present:absent ratio of 1:1 and a orientation:luminance:redundant ratio of 1:1:1. ISCB is 0.25 and NRC is 0.25.

Orientation	Luminance	
	T^L	N^L
T^O	20	20
N^O	20	60

Table 3.4 Occurrence of target types in Experiment 3 per 120 trials of condition 3, with a present:absent ratio of 3:1 and a orientation:luminance:redundant ratio of 1:1:1. ISCB is 0.4 and NRC is 0.15.

Orientation	Luminance	
	T^L	N^L
T^O	30	30
N^O	30	30

Table 3.5 Occurrence of target types in Experiment 3 per 120 trials of condition 4, with a present:absent ratio of 1:1 and a orientation:luminance:redundant ratio of 1:1:2. ISCB is 0.47 and NRC is 0.3.

Orientation	Luminance	
	T^L	N^L
T^O	30	15
N^O	15	60

Table 3.6 Occurrence of target types in Experiment 3 per 120 trials of condition 1, with a present:absent ratio of 3:1 and a orientation:luminance:redundant ratio of 1:1:2. ISCB is 0.59 and NRC is 0.17.

Orientation	Luminance	
	T^L	N^L
T^O	45	22.5
N^O	22.5	30

3.3.1 Method

Participants

64 observers took part in Experiment 3 (20 male, 3 left handed). Age ranged from 19 to 47 years (median: 24). Observers were paid with a rate of 8 Euro per hour.

Table 3.7 Contingencies favoring redundant over single targets for the four conditions of Experiment 3.

Target Present-Absent Ratio	Single-Redundant Ratio			
	1:1:1		1:1:2	
	ISCB	NRC	ISCB	NRC
1:1	0.47	0.3	0.25	0.25
1:2	0.59	0.17	0.4	0.17

Apparatus

The apparatus was the same as in Experiment 1.

Stimuli and Timing

Stimuli and timing are the same as in Experiment 1, except that only high intensity targets were displayed. Orientation targets were tilted either 45° to the left or to the right relative to the vertical distracters. In Experiment 1 intensity of luminance was adjusted on a per subject basis in such a way that orientation targets and luminance targets produced statistically equal reaction times. In Experiment 3 I used the mean luminance value that in Experiment 1 lead to reaction times equivalent to 45° orientation targets. Displays could either contain no target, or a target defined in one dimension (orientation or luminance), or a redundant target (both tilted and bright).

Design and Procedure

Experiment 3 consisted of one session of approximately one-hour duration. The session started with a practice block of 30 trials that was excluded from analysis. The main experiment consisted of 23 blocks with 60 trials each, leading to a total of 1410 trials. We varied the ratio of target-present and target-absent displays (present-absent ratio: 1:1 vs. 3:1, present:absent) crossed with the ratio of single to

redundant targets, if a target was present (single-redundant ratio: 1:1:1 vs. 1:1:2, orientation:luminance:redundant). Each participants was randomly assigned to one of the four possible combinations of present-absent and single-redundant ratios. The resulting contingencies for each condition is presented in Table 3.7.

3.3.2 Results

Trials with reaction times of less than 200 ms were considered anticipations (< 0.1% of all trials), and were excluded from further analysis.

Error Analysis

Overall error rate was low: with 0.2 % misses, and 4.4 % false alarms. All miss rates were below 1%, so I only further analyzed false alarms. False alarms were subject to a 2 x 2 ANOVA with factors present-absent ratio (1:1 vs. 3:1) and single-redundant ratio (1:1:1 vs. 1:1:2). Only the main effect of present-absent ratio was significant ($F(1,60)=33.9, p < .0001$, with 1.5 % false alarms in the 1:1 and 7.2 % in the 3:1 condition.

Reaction Time Analysis

The reaction time data of Experiment 3 from SSTs are presented in the left panel and from RSTs in the right panel of Figure 3.6. The mean reaction times were entered into a 3 x 2 x 2 ANOVA with within-subjects factor target type (orientation, luminance, and redundant), and between-subjects factors present-absent ratio and single-redundant ratio. The analysis revealed two significant main effects of target type ($F(2,120)=159.85, p < .0001$) and present-absent ratio: Targets in the 3:1 present-absent ratio condition were detected 18.0 ms faster than targets in the 1:1 condition ($F(1,60)=4.76, p < .033$). There was no additional significant main effect nor interaction. Comparing the different target types post hoc, employing Tukey's

Honest Significant Difference (HSD), revealed all pairwise comparisons to be significant. Orientation targets had approximately 4 ms shorter latencies than luminance targets, and redundant targets were found faster by approximately 22 ms and 27 ms than luminance and orientation targets respectively.




Figure 3.6. Mean reaction time for single targets at the two levels of intensity.

Figure 3.7 presents the mean RSE for each combination of present-absent ratio and single-redundant ratio. The RSE is significantly greater than zero as confirmed by the post-hoc test in the previous analysis.

Figure 3.8 displays the test of the RMI, i.e. $d(t)$ (Eq. 3.3) for all four conditions. In order not to overestimate violations of the RMI because of multiple t-tests I only tested the 0.05 until 0.2 quantiles (Kiesel et al., in press). In both conditions with a present-absent ratio of 1:1, there were significant violations for all tested quantiles between 0.05 and 0.2. In the 3:1 present-absent and 1:1:1 single-redundant ratio condition, there were significant violations for the 0.1 and 0.15 quantile. For the 3:1 and 1:1:2 present-absent and single-redundant ratio condition, there was the tendency of a violation at the first quantile ($p < .1$).

So far I found substantial redundancy gains and violations of the RMI. The purpose of Experiment 3 was to test, whether an interactive race model (Mordkoff & Yantis, 1991) could account for this data pattern. The interactive race model pre-




Figure 3.7. Mean RSE for each combination of present-absent ratio and single-redundant ratio.

dicts both the size of the RSE as well as the amount of violations of the RMI to increase, as the contingencies favoring redundant over single targets increase. In Table 3.7 the contingencies for all four experimental conditions are presented. We analyzed the RSE as well as the amount of violations of the RMI (as the area under $d(t)$, Colonius & Diederich, 2006) in dependence of the total amount of contingencies, i.e. the sum of both NRC and ISCB. Figures 3.10 and 3.9 present both measures as a function of the sum of NRC and ISCB. On visual inspection there is no hint of neither redundancy gains, nor the area under $d(t)$ to increase as the total amount of contingencies increase, as an interactive race model would predict. A linear regression of both redundancy gains and the area under $d(t)$ on the sum of NRC and ISCB confirmed that there was no significant slope effect. Also a linear regression on the area under $d(t)$ and on redundancy gains with NRC and ISCB as predictors revealed no significant slope effects. Hence, no interactive race model seems to be responsible




Figure 3.8. The estimate $d(t)$ for all combinations of present-absent ratio and single-redundant ratio.

for the observed redundancy gains and violations of the RMI.

Finally performed an ANOVA on redundancy gains and the area under $d(t)$ with the factors of present-absent ratio and single-redundant ratio. In both analysis the only significant effect was the main effect of present-absent ratio, with both redundancy gains, $F(1,60)=4.45$, $p < .04$, as well as the amount of violations of the RMI (i.e. the area under $d(t)$) being smaller for the 3:1 compared to the 1:1 ratio, $F(1,60)=4.9$, $p < .03$ (cf. Figs. 3.7 and 3.11).

3.3.3 Discussion

Experiment 3 replicated the finding of significant redundancy gains of redundant pop-out targets defined by feature contrast in two dimensions simultaneously compared to single targets, defined in only one dimension. These violations could not be




Figure 3.9. The area under $d(t)$ as a function of the sum of ISCB and NRC.

explained by statistical facilitation (Raab, 1962), as there were significant violations of the RMI (Miller, 1982). Additionally neither the size of redundancy gains nor the amount of violations of the RMI (i.e. the area under $d(t)$) increased, as contingencies favoring redundant over single targets increased. Although theoretically interactive race models could explain violations of the RMI, they would predict that the amount of violations increases, as the contingencies increase.

In summary, this finding supports the co-activation model of a master saliency map, into which dimension-based feature contrast signals are pooled before triggering a response (Müller et al., 1995; Wolfe, 1994).

3.4 General Discussion

The question of this study was, whether indeed co-activation models, such as the summation of dimensional feature contrast signals into a saliency master map




Figure 3.10. Mean RSE as a function of the sum of ISCB and NRC.

(e.g. Müller et al., 1995; Wolfe, 1994) were responsible for redundancy gains and violations of the RMI in visual pop-out search as opposed to two theoretical alternative models: serial models with an exhaustive search across dimensions (Townsend & Nozawa, 1997), or interactive race models (Mordkoff & Yantis, 1991), in which cross-talk would allow the exchange of information between both channels. Repeatedly authors reported violations of the RMI and interpreted them as support for co-activation models (Krummenacher et al., 2001, 2002; Turatto et al., 2004). The DWA is an example of a co-activation model, in which feature contrast of each dimension is summed into a master map of saliency (e.g. Itti & Koch, 2000; Wolfe, 1994). This sum can be modulated by weights assigned to each dimension (Found & Müller, 1996; Müller et al., 1995). The detection process in this model relies on the summed signal of both dimensions (i.e. channels).

An example for a serial exhaustive model would be the strategy of observers to check both dimensions (orientation and luminance) in both SSTs and RSTs. This proce-

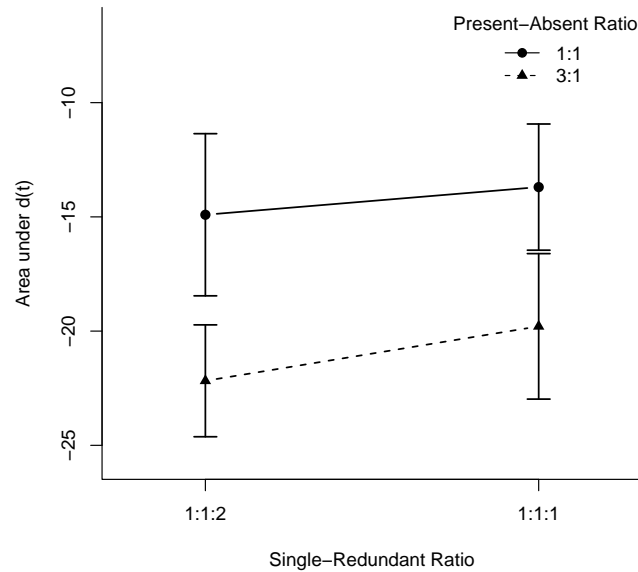


Figure 3.11. Mean RSE for each combination of present-absent ratio and single-redundant ratio.

dure would lead to prolonged reaction times in SSTs, because even if the target has been found to be defined in one dimension, the other dimension has to be checked and search terminates only when absence of a target in this dimension has been determined. Thus violations of the RMI, according to such a model, are not due to co-activation or summation of signals from both channels, but due to a stronger slowing of SSTs compared to RSTs because of the exhaustive search across both dimensions. Conversely in interactive race models, both channels can communicate before the discrimination of target presence or absence. If this communication speeds up the detection process of redundant targets to a larger extent than single targets, violations of the RMI can occur. Communication between both channels is only useful, if the experimental design leads to contingencies between the channels. For example, if the probability that a target is present in one channel (A) under the condition that there is no target present in the other channel (B) is higher than the overall probability of a target being present in channel A, then interactions between both

channels can speed up target detection. Mordkoff and Yantis (1991) formalized two types of contingencies that can favor redundant over single targets: NRC and ISCB (cf. Eqs. 3.7 and 3.6, as well as Table 3.7).

Although no authors have proposed one of both models to underlay visual processing in visual search, they are theoretical alternatives which also could be responsible for violations of the RMI in pop-out search (Krummenacher et al., 2001, 2002; Turatto et al., 2004). Experiment 1 implemented a double factorial redundant target design (Townsend & Nozawa, 1995) with the goal of testing whether serial exhaustive models indeed are an alternative explanation of reported RSEs and violations of the RMI in visual pop-out search. Pop-out targets in Experiment 1 could either be orientation or luminance defined and differed in similarity to distracters, leading to two different levels of feature contrast. In addition to the four types of single targets (orientation or luminance targets of high or low feature contrast, each), there were four possible redundant targets (two dimensions x two levels of feature contrast). For instance a redundant target could differ from distracters strongly in the orientation dimension and less in the luminance dimension. Employing Sternberg's additive factors logic (Sternberg, 1969a), Townsend and Nozawa (1995) proved that in redundant targets serial models of any stopping rule the intensity (i.e. feature contrast) manipulation would exhibit no interaction. If on RSTs both dimensions are indeed processed in a serial fashion, slowing of one component (e.g. low feature contrast in the orientation dimension) should be independent of slowing of the other component (e.g. low feature contrast in the luminance dimension). The mean interaction contrast (Eq. 3.1) would be zero, if dimensions were checked serially (independent of the stopping rule).

Mean reaction times of SSTs in Experiment 1 showed that the manipulation of feature contrast was successful: search performance was slower for single signal pop-out targets of low salience (dim or 6° orientation deviance), than for single signal pop-out targets of high salience (bright or 45°). The analysis of RSTs alone revealed the level of intensity of the luminance-component and the orientation-component to interact. The mean interaction contrast was significantly above zero, i.e. the interaction was

super additive, i.e. reducing feature contrast in one component of redundant targets had a larger effect, when the other component was of low feature contrast, also. Super additive interactions are not in accordance with serial models of any stopping rules, nor with exhaustive parallel models. Putting reaction times in SSTs into relation to reaction times in RSTs I found a significant benefit of RSTs over SSTs, i.e. I observed a significant RSE in all four conditions.

The super additive interaction serves as a constraint on what models can explain the observed RSE. The advantage of pop-out targets defined in two dimensions simultaneously over pop-out targets defined in a single dimension cannot be accounted for by serial exhaustive models, but only by self-terminating parallel, parallel interactive, or co-active processing (Townsend & Nozawa, 1995; Nozawa, Reuter-Lorenz, & Hughes, 1994; Townsend & Ashby, 1983). In summary, the observed pattern of RSE and an over additive interaction are consistent with both a parallel self-terminating race and a co-activation model.

In order to test parallel race models the RMI (Miller, 1982) has been evaluated: As in recent studies, violations of the RMI were observed (Krummenacher et al., 2001, 2002), ruling out parallel race models as the underlying processing architecture. Analysis of the capacity coefficient $C(t)$ indicated that the underlying system was super capacity for those processing times, where the RMI is violated, as well as at some additional times. This finding is in accordance with the prediction of Townsend and Nozawa (1995), who propose that local violations of the RMI always go along with local super capacity. At those processing times, where the system is super capacity, increasing the load (i.e. presenting a signal in two dimensions simultaneously, instead of in just one) leads to higher system performance. For an unlimited capacity system, performance would be independent of load, and in a capacity limited system performance decreases with increasing load. In Experiment 2 search slopes for targets of low intensity (e.g. 6° deviance between targets and vertical distracters) have been observed to be below 4 ms/item. So these targets could be searched efficiently in a pop-out manner.

In summary, the main finding of Experiment 1 was: serial exhaustive models cannot account for violations of the RMI in visual pop-out search. Before the question of interactive race model was targeted in Experiment 3, Experiment 2 aimed at verifying that the low feature contrast conditions of Experiment 1 indeed lead to efficient (i.e. pop-out) search. The exclusion of serial models (of any stopping rules) is especially important for recent studies reviving feature integration theory (FIT: Treisman & Gelade, 1980; Treisman, 1988; Treisman & Sato, 1990): Chan and Hayward (2007) and Mortier et al. (2007) based substantial arguments of their studies on the idea that for detection, feature contrast signals are pooled per dimension in a location unspecific manner. So unlike the signals that are transmitted to the saliency master map, such dimensional representations do not contain any information about the location at which the feature contrast occurred. The detection process is assumed to either serially check all dimensions (Chan & Hayward, 2007) or is not further specified (Mortier et al., 2007) and could be either serial or parallel. However, findings of Experiment 1 exclude the possibility of models, in which the search process serially checks dimensions for presence of a target, because such models would predict that for redundant target manipulations of the intensity of feature contrast signals in the two respective dimensions behave in an additive fashion. In contrast to this assumption I found the factorial manipulation of feature contrast to lead to a super-additive interaction. Also parallel checking of dimension based spatially pooled signals is problematic: as this is basically a parallel race situation, in such a case redundant signals would lead to redundancy gains, but the RMI should hold. Nevertheless violations of the RMI have been repeatedly shown for redundant pop-out targets (Krummenacher et al., 2001, 2002; Turatto et al., 2004; Koene & Zhaoping, 2007, the present study), which excludes the possibility of such parallel race models. Even if FIT would be equipped with a co-activation mechanism, such that the dimension based signals would be integrated before the detection process, the assumption of dimensional signals being spatially unspecific is at variance with the finding that redundant signals only lead to violations of the RMI, if they are present at the same location, or at

least in close spatial proximity (Krummenacher et al., 2002). Given these findings both assumption of FIT about detection of pop-out targets - spatial pooling of dimensional contrast signals and serial search over dimensions - are highly questionable.

Can interactive race models (Mordkoff & Yantis, 1991) be responsible for violations of the RMI observed in the present and in other studies? In order to answer this question in Experiment 3 I manipulated the ratio of a target being present or being absent crossed with a manipulation of the ratio of single to redundant target trials if a target was present. The present-absent ratio could be 1:1 or 3:1, and the ratio of orientation:luminance:redundant targets could be 1:1:1 (twice as many SSTs than RSTs) or 1:1:2 (as many SSTs as RSTs). Each of these four conditions lead to a different amount of contingencies that could favor redundant over single targets (NRC and ISCB, Eqs. 3.7 and 3.6). The prediction of the interactive race model would be that the size of redundancy gains and the amount of violations of the RMI (i.e. the area under $d(t)$, Colonius & Diederich, 2006) would be positively correlated to the amount of contingencies. As in Experiment 1, in Experiment 3 redundant targets had shorter latencies than single targets, and the RMI was found to be violated. However, neither the size of redundancy gains nor the amount of violations of the RMI could be predicted by the amount of contingencies. Neither NRC, nor ISCB, nor the sum of both measures lead to an observable benefit in redundancy gains or violations of the RMI. Hence, I argue, interactive race models also cannot explain the observed violations of the RMI in visual pop-out search.

Based on the evidence from Experiments 1 and 2 of the present study, both theoretical alternative models - serial exhaustive and interactive race models - have to be rejected as explanations of observed violations of the RMI in visual pop-out search. This is strong support for summation of saliency models, such as Guided Search (Wolfe, 1994), the computational model of Itti and Koch (2000), or the DWA

(Müller et al., 1995; Found & Müller, 1996), as already argued by Krummenacher et al. (2001, 2002).

So far I have argued that the reaction time benefit of redundant over single pop-out targets and violations of the RMI are based on a pre-attentive summation/pooling of feature contrast signals into a master map of saliency. There are results of Experiments 1 and 3, which I have not yet discussed that also support the perceptual origin of the redundancy gains. In Experiment 1 redundancy gains were reliable but moderate (approximately 5 ms) for all redundant targets that involved a high feature contrast component. Redundant targets defined by low feature contrast in both dimensions in contrast lead to a reaction time benefit of about 50 ms over single targets defined by low feature contrast. That is redundancy gains were larger for a condition, in which latencies were slower. In Experiment 3 redundancy gains were significantly increased in both conditions where target and non-target displays were equally likely compared to the conditions, where targets were present in 75 % (i.e. present-absent ratio of 3:1) of all displays. Simultaneously, reaction times were faster in the 3:1 than in the 1:1 present-absent ratio conditions. So again redundancy gains were larger in conditions, in which processing times were increased. In both experiments, redundancy gains were larger in those conditions, where processing times were slower than in those conditions, where processing times were faster.

In Experiment 1, strength of feature contrast was responsible for the slowing of reaction times. Targets which differed more strongly from distracters (i.e. produced high feature contrast signals), were found approximately 100 ms faster than targets which differed less from distracters. Experiment 2 confirmed that low intensity targets still were found efficiently. In terms of saliency summation models, activity on the master map builds up faster and/or to a higher level for targets of high than for targets of low feature contrast. Thus a choice between absence or presence of a target based on the level of activity on the master map can be taken faster for more salient targets than for less salient ones. So the manipulation of Experiment 1 was a perceptual one that affected the decision latency, when observable reaction time is modeled as the

sum of a decision latency T_D and a response processing time T_M (see Figure 3.2).

In Experiment 3 reaction times were faster in those conditions where the present response was more frequent. Smith and Ratcliff (2004) argued that in a decision process the criterion for a certain response is lowered, the more frequent this response occurs. A lower criterion to answer target-present is supported by the finding of significantly more false alarms in the 3:1 present-absent ratio conditions. A lower criterion would also lead to less misses, but this effect cannot be observed due to overall low miss rates of less than 1%. Latency differences in observable reaction times that are caused by a manipulation of decision criterion also affect the decision, and not the response preparation and execution component. In terms of saliency summation models, the manipulation of target frequency shifted the criterion, what level of activity on the master map had to be exceeded in order to trigger a response. In both Experiments 1 and 2 size of redundancy gains varied with manipulations that can be explained in terms of the decision process based on activity of a master map of saliency. Models such as Guided Search, the computational model of Itti and Koch, and the DWA assume that the decision to deploy of attention to a certain location is guided by activity on the master saliency map. The manipulation of feature contrast in Experiment 1 directly affected the level of activity on the master map, whereas the manipulation of target frequency in Experiment 3 affected the criterion, based on which a shift of attention is triggered. As both manipulations modifying the size of redundancy gains affect pre-attentive processing stages, it is safe to argue that the processes which produce redundancy gains also are pre-attentive. This is in line with e.g. the DWA. Other models, such as the dimension-action model of Cohen and colleagues (Cohen & Shoup, 1997; Cohen & Magen, 1999; Feintuch & Cohen, 2002; Cohen & Feintuch, 2002) propose that co-activation can happen only at locations which have been selected by attention. If that were the case, than manipulations which affect the speed with which attention can be deployed (such as manipulations of feature contrast or of criterion) would not affect the size of observed redundancy gains.

The results reported here are important for the entire category of models assuming integration of multiple dimensional signals into a common representation. Further it is related to all theories and models, which assume violations of the RMI to be indicative of co-activation. Also in such cases, the theoretically plausible alternatives of serial exhaustive and parallel interactive models have to be excluded. In addition to the double-factorial design of Townsend and Nozawa (1995) as well as to the introduction of two target-absent signals (Mordkoff & Yantis, 1991), the present study contributes a further method to evaluate interactive race models. If the experimental paradigm does not allow for the introduction of two target-absent signals, we demonstrated that controlled manipulations of contingencies beneficial for redundant over single targets are a valid alternative. Interactive race models predict the RSE and the amount of violations of the RMI to be positively correlated with the strength of such contingencies. In the present study we could show that in a redundant-target visual search paradigm both the RSE as well as the amount of violations of the RMI were independent of such contingencies, ruling out the possibility of interactive race models of processing in visual search. Together with the exclusion of serial exhaustive models as an alternative explanation for violations of the RMI, the present study strongly supports saliency map models (an instance of co-activation models) of processing in visual search. The DW account (Müller et al., 1995; Found & Müller, 1996) is such a co-activation model that proposes the existence of a master saliency representation into which dimension-based feature contrast signals are fed in a weighted fashion. The benefit of pop-out targets defined in two dimensions over targets defined in only one is stemming from a faster build-up of activation of an integrated saliency representation in case of redundant targets, because activity at the target location is driven by feature contrast signals from two dimensions rather than from only one. This perceptual/pre-attentive locus of co-activation is supported by the fact that manipulations of saliency as well as of response criterion affected the size of the RSE, as well as of the amount of violations of the RMI, that is of co-activation.

4. INTENTION AND TRIAL HISTORY IN LOCALIZATION

Whether pre-attentive vision is penetrable by top-down prior knowledge or expectation is under debate. While some authors argue that top-down modulation of pre-attentive vision is almost impossible (Maljkovic & Nakayama, 1994; Theeuwes, 1992; Theeuwes et al., 2006), others argue that it is (Müller et al., 2003; Müller & Krummenacher, 2006).

What can be understood as pre-attentive vision? Prominent theories of visual search assume the existence of topographical representations of feature contrast, in separate feature maps (Wolfe, 1994; Treisman & Gelade, 1980; Koch & Ullman, 1985; Itti & Koch, 2000). Feature contrast is a measure of how different a specific location in the visual field is relative to its surrounding locations regarding a certain feature. For example, for a red vertical bar that is surrounded by green vertical bars feature contrast for red is high and feature contrast for vertical is low. Feature contrast signals are further assumed to be pooled into dimension specific maps (for an overview about dimensions in visual search see Wolfe, 1998a; Wolfe & Horowitz, 2004) and then summed into a saliency master map. Müller et al. (1995) proposed that the dimension specific feature contrast signals can be modulated via dimensional weights. Activation at any location of the master map indicates local differences in features, without any information about features or dimensions. Activation on this map can then be used to prioritize deployment of attention for more detailed analysis. This conception has support from neuronal findings. There are topographical maps in the brain which have the main features of a saliency map: they represent the strength of center-surround contrast for the different locations of the visual scene and are feature-unspecific at the same time. One neural network with these properties is the oculo-motor network (Fecteau & Munoz, 2006). Especially the frontal eye field (FEF) is a structure that most likely implements a saliency map (Thompson & Bichot, 2005;

Bichot & Schall, 1999).¹ In the present context, pre-attentive vision will be understood as the processing that occurs including integration of feature contrast signals in the master saliency map before the deployment of attention. In contrast, processing which takes place after attentional selection of a salient location will be referred to as post-selective vision. The term 'top-down' in the present paper refers to effects of prior explicit knowledge, strategies, or intentions of participants. These modulations could for instance take place on the level of computation of feature contrast, or on the level of summation of dimensional contrast signals into the master map.

One prominent argumentation for top-down influences on pre-attentive vision is based on reaction time differences for pop-out targets in intra- or cross-dimensional blocks (Treisman, 1988). That is: when the target dimension (e.g. size) is constant across trials, reaction times are faster, then when the target dimension can change from trial to trial (size targets randomly intermixed with e.g. color target trials). The finding of reaction times being substantially faster in intra-dimensional as compared to cross-dimensional blocks led Treisman (1988) to the conclusion that top-down knowledge of the upcoming target can speed the pre-attentive perceptual search process. More precisely, the fact that size targets are detected faster in blocks in which targets are always size defined than in a block in which targets could either be defined by size or color is, according to Treisman (1988), due to prior (top-down) knowledge of the observer about the target's dimension.

A different explanation for this finding was provided by Müller et al. (1995) and Found and Müller (1996) within the framework of the dimension weighting account. This account assumes that the signals stemming from dimensional feature contrast maps are not plainly summed into the master map, but that these signals can be modulated by dimension-specific weights. These weights are assumed to act as a limited resource that is, if the weight of one dimension is increased, the weights of other dimensions are automatically decreased. So, if for example the weight for the color dimension

¹Other structures, such as the lateral intraparietal area (Gottlieb, 2002) also show properties of a saliency map. There is consensus that there is more than one area in the brain that implements a master saliency map.

is increased, contrast signals from the color dimension contribute more to activation of the master saliency map than other dimensions. The increased weight leads to a faster and/or higher build-up of activation on the master map, which results in faster reaction times and higher accuracy for color defined targets (relative to targets defined in other than the color dimensions). Simultaneously, if instead of a color target an orientation target is encountered, feature contrast from the orientation dimension has less (and/or slower) impact on the master map, because by the prior increase of weights for color, the modulatory weight for orientation had been decreased. This weaker or slower impact on the master map can then be observed in slowed reaction times.

Applied to the intra- and cross-dimensional conditions, in the intra-dimensional condition, in which the target always is defined in one dimension, dimensional weight for that dimension will be maximal, and weights for the other dimensions will be minimal. Therefore reaction times are fast. When the target can be defined in different dimensions, dimensional weights will be distributed equally across the relevant dimensions (i.e. they will be less than maximal). Consequently, reaction times are slow.

In addition, in the a cross-dimension condition Found and Müller (1996) analyzed the trial-history of short sequences of trials with regard to reaction times for a given dimension in trial n depending on whether the dimension in trial $n - 1$ was the same or different. They found substantial reaction time costs (around 25 – 45 ms), when the dimension changed over two consecutive trials compared to when the dimension stayed the same. This finding substantiates the dimension weighting account, which presumes that the dimensional weights are automatically shifted to the dimension, in which a target is defined at a present trial. If for example in trial $n - 1$ the target is defined in the orientation dimension, weights of the orientation dimension are increased, while weights of all other dimensions are decreased. In that case, on trial n an orientation target will be found faster than when the target of trial $n - 1$ would have been defined in the color dimension: the dimensional weights of color would

have been increased in trial $n - 1$, but the weights of orientation would have been decreased.

According to the dimension weighting account, this shifting of weights, is responsible for the reaction time difference between intra- and cross-dimensional search. However, assumptions of top-down processes (Treisman, 1988) are not necessary to explain these findings. Thus, it is probable that the reaction time difference between intra- and cross-dimensional search is not due to top-down penetration of pre-attentive search, but due to an automatic process that leads to a slowing of search performance if the dimension changes on consecutive trials. Does that consequently mean that there is no top-down penetrability of pre-attentive search at all? In addition to the automatic process, which shifts dimensional weights from trial to trial, the dimension weighting account includes another way, in which the dimensional weights can be changed: top-down intention. It assumes that participants can prepare themselves for a specific dimension, which leads to an increase of the respective weight. This top-down impact on dimensional weights leads to modulated build-up of activation on the master map that is, it has a pre-attentive effect.

This view is supported by the effect of symbolic dimensional cues in a pop-out detection task (Müller et al., 2003). In this study the authors presented a trial by trial cue of the dimension of upcoming target. Valid cues led to faster reaction times than neutral or invalid cues, in accordance with the predictions of the dimension weighting account. The dimension weighting account explains the facilitated search for valid cues in the following way: observers use the cue and increase attention to the cued dimension. This process increases the dimensional weights of that dimension, which modulate the signals that are fed into the master saliency map. Thus, activation on the master map accumulates faster and/or higher for the cued dimension.

Theeuwes et al. (2006) replicated the finding of faster reaction times for valid compared to invalid dimensional cues in a present/absent pop-out search task. However, their reasoning was somewhat different: instead of top-down modulation of pre-attentive vision (Müller et al., 2003), Theeuwes and colleagues attributed this

reaction time facilitation to purely post-selective processes. They based their argumentation on findings from dimensional cueing in a compound task that is, in a task, in which a feature singleton was present in every trial and the observers had to report the orientation (left or right tilted) of a line segment, which is placed inside the targets. Here Theeuwes et al. (2006) failed to find reaction time facilitation for valid dimensional over invalid cues. Because the one and only difference between the detection (Müller et al., 2003) and compound task (Theeuwes et al., 2006) were observer's response requirements, Theeuwes et al. attributed the cueing effect to post-selective, response-based processes. However, Müller and Krummenacher (2006) did indeed find significant cueing effects using the same compound task and stimulus material as Theeuwes et al. (2006). They reasoned that perhaps participants in Theeuwes et al. (2006) did not have enough incentive to make use of the cue, because the dimensional identity of the target was related to its response-relevant property. In order to increase commitment to use the cues in Müller and Krummenacher (2006), after each block of trials made the rate how much they had 'attended' to them. With this additional manipulation, valid-dimension cues were found to significantly facilitate the compound-task response (albeit to a small extent only; but see Töllner, Gramann, Müller, Kiss, & Eimer, 2007, for other biases determining compound-task responses).² Mortier et al. (2007) have recently examined this discrepancy of cueing effects in detection and absence thereof in tasks requiring spatial information tasks. They directly compared cueing benefits for detection and localization tasks and found cueing benefits only in the former, but not in the latter task, regardless whether the response had to be given in a manual or ocular fashion. Therefore, they concluded that detection tasks and tasks that require spatial information (i.e. localization and compound tasks) are processed via different processing routes - one (non-spatial) in which top-down

²Note that Mortier et al. (2007) argued that this cueing facilitation might be an artifact induced by the dual task situation, in which observers simultaneously had to use the cue as well as judge how well they make use of it. In other words, according to Mortier et al. (2007), the finding of cueing benefits under dual task conditions does not necessarily argue for the pre-attentive nature of the effect. They did not provide arguments, why this might be the case and concede that the argument is somewhat post hoc.

weighting is possible, and one (spatial) in which no top-down weighting is possible. This dual-route proposal for detection and spatial tasks has also been put forward by Chan and Hayward (2007), who compared dimensional intertrial costs between detection, localization, and compound tasks. They found dimensional intertrial costs only in detection, but not in spatial tasks and concluded that detection tasks are not processed via a master salience map, but via dimensional modules, which signal only presence of dimension-specific feature contrast, but contain no spatial information (a concept revived from Feature Integration Theory, Treisman & Gelade, 1980; Treisman, 1988).

The present study aims to shed new light on the origin of performance benefits induced by top-down dimensional cues in two regards: do they arise at a pre-attentive or at a post-selective processing stage, and are the processing routes for detection and localization tasks common or separate? The task employed in the present study is a localization task, in which observers are presented with a pop-out target in every trial and have to indicate, whether the target is located in the left or the right half of the display. Hence, this task is basically a compound task (Duncan, 1985; Bravo & Nakayama, 1992), in which the target defining and response defining features are different from each other. In Experiment 1, displays were presented only briefly and followed by a mask. This manipulation should minimize post-selective processing: result patterns that can be found both in reaction time measures for long display durations as well as in accuracy measures for brief display durations are likely to arise from pre-attentive, not from post-selective or response based processing stages (e.g. Prinzmetal et al., 2005). If the cueing effects arise from a pre-attentive stage, (i) repetition of target dimension should lead to higher accuracy compared to change of the target dimension, and (ii) valid dimensional cues should increase accuracy over neutral and/or invalid cues.

Two other studies have been examining trial sequence effects in visual search using brief display durations: Huang and Pashler (2005) and Sigurdardottir, Kristjansson,

and Driver (2007). Still, unlike the present study they did not examine intertrial dimension repetitions, but employed a paradigm, in which target and distracters could change their roles from trial to trial (Bravo & Nakayama, 1992; Maljkovic & Nakayama, 1994), as well as a conjunction search in the case of Sigurdardottir et al. (2007). Huang and Pashler (2005) specifically address the question of top-down expectancy in pre-attentive vision. However, they manipulate expectancy in a block-wise fashion by manipulating the predictability of changes compared to repetitions. In contrast in the present study, I employed a trial-by-trial cueing procedure, to discern the effects of correct and incorrect expectations, not only the effect of presence vs. absence of expectation as in Huang and Pashler (2005). We will relate the present findings to these studies in the General Discussion.

Experiment 2 measured speeded reaction times under unlimited viewing conditions in order to examine the conditions under which cueing effects can be found in a speeded localization task.

4.1 Experiment 1

In Experiment 1 the task was to report the location (left vs. right visual hemifield) of a fatural singleton that could be defined either by orientation (i.e. left or right tilted bar amongst vertical bars) or luminance (light grey bar amongst dark grey bars) dimension. The display was presented briefly and followed by a mask. After reporting the location of the target, observers indicated confidence in their response (i.e. how sure they were their response was correct). The target dimension could change randomly from trial to trial. There were three different types of symbolic cues (neutral, valid, and invalid) which before each trial indicated the dimension of the target.

4.1.1 Method

Participants

12 observers (7 male) with a median age of 24.5 (ranging from 20 to 30), two left-handed, participated in Experiment 1. They received 8 Euro (11 USD) per hour.

Apparatus

Observers viewed the stimuli in front of a Sony Multiscan E250 17" monitor driven by personal computers with Windows XP operating system. The experimental software was purpose written in C++. The personal computer was placed in a sound isolated room with black interior. There was dim background light in order to prevent reflections on the monitor. Viewing distance was about 60 cm and observers were instructed to maintain constant distance to the monitor. The screen refresh rate was 85 Hz, the screen resolution was set to 1024x768 pixels. Participants reported target location by pressing the right or left button of a mouse with the index or middle finger of their right hand and confidence of their response with the keys 1, 2, and 3 of the numeric keypad with their left hand. Reaction times and accuracy were recorded online by the computer. After each block the participants were informed about their mean reaction time and error rate of the previous block.

Stimuli and Timing

The display consisted of a 6x6 array (subtending $15.3 \times 15.3^\circ$ of visual angle) of filled upright rectangles (bars) on a black background with 0.6 cd/m^2 . The bars were either dark (11.6 cd/m^2) or light grey; the luminance of the light grey bars were adjusted individually for each participant and ranged from 5 to 51 cd/m^2 . The bars subtended approximately 1.7° of visual angle of height and 0.35° of width. They were arranged within invisible cells of 2.55° by 2.55° with a jitter of about 0.2° around the (invisible) center of the array cell. Distracters always were medium grey (5 cd/m^2)

and vertical bars. The target was either defined in the luminance (light grey), or in the orientation (45° tilted to the left or right of vertical) dimension. Targets were always placed in such a way that they were surrounded by distracters (i.e. they appeared in the inner 4x4 cells of the matrix). Participants were not informed about this restriction. A target was present in 100% of the trials and was randomly placed in the left or the right half of the array. There were three types of cues, neutral, valid, and invalid. Neutral cues were presented in separate blocks. Trials started with either the German words for 'neutral', 'orientation', or 'luminance', which displayed for 800 ms, followed by a fixation cross which stayed onscreen for 800-1000 ms. Neutral cues were presented in blocks. After this period of time, there was the simultaneous onset of all stimuli. The stimuli were presented for a brief period of time (47-82 ms) and then replaced by a mask (see Fig. 4.1 for details). For the mask at each stimulus location ten lines of 0.1° of width and a randomly chosen length between 0.14 – 1.7°. They were placed uniformly distributed with their center no further away than 1.7° from the center of the (invisible) array cell. Orientation of each line segment was also determined randomly.

The mask stayed on the screen until participants reported the location of the pop-out




Figure 4.1. The sequence of displays during one trial.

target (left or right half of the display) and went blank until they rated the confidence in their response, ranging from 1 (unsure) to 3 (sure). After erroneous trials there was a blank screen presented for 2 s as a feedback signal. On correct trials the intertrial interval was 0.5 s.

Design and Procedure

Before the experimental trials, presentation time and intensity of luminance targets was adjusted by a staircase procedure for each observer individually. Each participant performed both the neutral and the cue conditions. There were sequences of five blocks, one with neutral and four with valid or invalid dimensional cues. Each participant started at a random starting point in this sequence of five blocks. Trials were presented in blocks of 60, followed by a feedback about mean error rate. Each session lasted 50 minutes. Sessions were separated by at least two hours but not more than two days.

During the first session, both presentation time and intensity of the luminance targets were adjusted for each participant individually with an adaptive staircase procedure in such a way that, on average, performance was correct in 75% of the trials (see e.g. Johnston et al., 1993). The presentation time and intensity of luminance targets obtained by the staircase procedure were kept constant throughout the experiment for each observer. No cues were presented during the adaptation phase.

During the experimental trials a cue was always present, and participants were instructed that cues predicted the dimension of the upcoming target with a probability of 80%.

Data Analysis

Following their localization response, participants rated the confidence of their response on a 3-level scale (not sure, medium sure, and sure). In signal detection theory (SDT) terms the location task can be considered a 2-alternative-forced-choice (2AFC) procedure. The main dependent variable was accuracy in terms of percentage correct, pc . Green and Swets (1964) showed that accuracy pc in a 2AFC procedure equals the area under the receiver operating characteristic (ROC) curve, which is usually taken to reflect a measure of perceptual sensitivity. Given that according to Schulman and Mitchell (1966) the 2AFC can also be interpreted as a yes/no detection task, in Exper-

iment 1 the area under the ROC curve was calculated by a non-parametric estimate, A_g (Pollack & Hsieh, 1969; Macmillan & Creelman, 2005). In short A_g approximates the area under the ROC curve and requires confidence ratings. In addition confidence ratings were used to further examine the dimension switch effects: I calculated dimension switch costs between trials $n - 1$ and n depending on the confidence rating of trial $n - 1$.

Although the response was non-speeded, I also analyzed three aspects of reaction times: (i) rt_{loc} as the time from stimulus onset to the localization response, (ii) rt_{conf} as the time from the localization response to the rating of confidence, and (iii) rt_{all} as the time from stimulus onset to rating of confidence.

The dependent variables pc , A_g , rt_{loc} , rt_{conf} , and rt_{all} were calculated for each participant and each condition. The independent variables were target dimension (orientation and luminance), validity of cue (neutral, valid, and invalid), and intertrial transition (same dimension vs different dimension in trials n and $n - 1$). Computational data analysis was done with R (R Development Core Team, 2006).

4.1.2 Results

Trials following erroneous responses (21% overall) were excluded from analysis. For the evaluation of dimensional intertrial transitions it is necessary to analyze sequences of two trials (n and $n - 1$). Therefore only tuples of correct responses were entered into analysis. Further, responses faster than 150 ms and slower than 3 s (less than 3.5% of experimental trials) were excluded from reaction time analysis.

Mean percent correct, pc , was 79%. A repeated measures analysis of variance (ANOVA) with the factors of target dimension (orientation, luminance), cue validity (neutral, valid, invalid), and intertrial transition (same, different dimension) revealed the main effects of cue validity ($F(2,22)=10.60$, $p<.0005$) and intertrial transition

($F(1,11)=10.96$, $p<.007$) to be significant. No further effects were significant.

So trials that were preceded by the same target dimension were responded to more correctly than trials that were preceded by a target of different dimension (81% vs. 77%; main effect of intertrial transition). Further, performance on invalidly cued trials was worst, compared to performance in neutrally or validly cued trials (75% vs. 78% and 80%; the main effect of cue validity). The estimated area under the ROC curve A_g was highly correlated with pc as predicted by the Area-Theorem (Green & Swets, 1964). Across all observers the mean correlation coefficient between pc and A_g was 0.90 ($p < .001$). pc for the different types of cues are presented in Table 4.1 and pc for the different types of intertrial transitions in Table 4.2.

Overall reaction times, rt , were analyzed in a separate, repeated measures ANOVA with the factors target dimension, cue validity, and intertrial transition, which revealed the main effect of cue ($F(2,22)=16.65$, $p < 10^{-5}$), and of intertrial transition ($F(1,11)=9.41$, $p < .01$ to be significant). Additional ANOVAs on both components of rt , rt_{loc} (reaction time from stimulus onset until localization response) and rt_{conf} (reaction time from localization response until rating of confidence) both revealed significant a main effect of cue (rt_{loc} : $F(2,22)=10.80$, $p < .0005$, rt_{conf} : $F(1,11)=9.3$, $p < .001$). For rt_{loc} also the main effect of intertrial transition was significant: $F(1,11)=8.16$, $p < .01$. For rt_{conf} the three-way interaction was significant ($F(2,22)=6.45$, $p < .006$). No further effects were significant in these three ANOVAs. Reaction times for the various types of cues and intertrial transitions are also presented in Table 4.1 and Table 4.2. In summary, overall reaction times, rt , were faster, when target dimension was repeated relative than when it changed over two consecutive trials (main effect of intertrial transition). This reaction time benefit mainly stemmed from the localization response rt_{loc} (significant main effect for rt_{loc} , but not for rt_{conf}). Overall reaction times, rt , were also substantially faster for valid and slower than invalid cues, when compared to neutral cues. This cueing benefit stemmed from both reaction times of localization responses, rt_{loc} , and reaction times of confidence ratings, rt_{conf} .

Table 4.1 Accuracy in percent correct (pc in %) and reaction times in ms for valid, neutral and invalid cues. Reaction times are presented as overall reaction times (rt), reaction time for localization (rt_{loc} , and reaction time for rating of confidence (rt_{conf}).

Cue	pc	rt	rt_{loc}	rt_{conf}
Valid	80	1084	760	324
Neutral	78	1095	766	328
Invalid	75	1178	824	353

Table 4.2 Sensitivity (A_g), percent correct (pc in %), and reaction times in ms for intertrial transitions of same and different dimension. reaction times are presented as overall reaction times (rt), reaction time for localization (rt_{loc} , and reaction time for rating of confidence (rt_{conf}).

Intertrial transition	pc	rt	rt_{loc}	rt_{conf}
Same dimension	79	1099	767	338
Different dimension	76	1138	799	331

Finally I took into account observer's confidence ratings in order to examine, whether accuracy in localizing a same or different dimension target was influenced by observer's confidence ratings in trial $n - 1$. The question was as follows: does dimensional weighting show some modulation depending on how confident observers were about their response. Therefore, I calculated pc for the same and different dimension conditions (data combined across the variables target dimensions and cue validity) separately for the three confidence levels in trial $n - 1$. Figure 4.2 displays the difference of pc between same dimensional and different dimensional intertrial transitions for the three types of confidence ratings in trial $n - 1$. Taking data at face value, only if the previous response was rated 'sure', there was an accuracy benefit of repetition over change of target dimensions. A repeated measures ANOVA on the accuracy differences (same dimension target minus different dimension target accuracies) with the factor of observer's confidence ratings revealed this effect to be

significant: $F(1,11)=12.146$, $p < .0002$. Paired t-tests showed accuracy costs differed significantly between previous rating sure and the other previous ratings ($t(11)=6.7$, $p < .0001$; sure vs. medium sure, and $t(11)=4.0$, $p < .006$; sure vs. unsure). Note that p-values were adjusted with the Bonferroni correction for multiple comparisons.




Figure 4.2. This figure displays the cost in accuracy due to changes compared to repetition of dimension across two trials depending on the confidence rating of the previous trial.

4.1.3 Discussion

In Experiment 1 participants had to perform a compound task, in which they indicated on which half of the display (left vs. right) a pop-out target was presented. Before each trial they were provided with either a neutral or a dimensional valid or invalid cue. In order to minimize post-selective processing, I shortened the presentation time on an individual basis, so that mean performance was about 75% correct.

Besides accuracy I analyzed rt as the overall time from stimulus onset until the report of confidence, as well as its subcomponents rt_{loc} , i.e. the time from stimulus onset until report of location, and rt_{conf} the time from report of location until rating of confidence. Participants were not instructed to respond as quickly as possible, so reaction times have to be interpreted with caution.

We found improved accuracy for valid dimensional cues over neutral or invalid cues. Additionally performance was significantly better if the dimension repeated over two trials, than when it changed. The same pattern of results was found for reaction times. Participants localized targets faster, when the target dimension was validly cued, as well as when the target was preceded by a target of the same dimension.

These findings are in line with the dimension weighting account, according to which dimensional weights can be shifted passively in a bottom-up fashion (Müller et al., 1995; Found & Müller, 1996) as well as actively by top-down intention (Müller et al., 2003; Müller & Krummenacher, 2006). The weights modulate dimension-based feature contrast signals when they are pooled into the saliency master map which, according to the dimension weighting account, is used in both detection, as well as in compound tasks when the target has to be localized. 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. (2000) 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 (Pollmann, Weidner, Müller, Maertens, & Cramon, 2006; Pollmann, Weidner, Müller, & Cramon, 2006; Weidner et al., 2002) 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. This may explain the effect of confidence rating in the previous trial on dimension change effects reported here. We found

substantial dimension switch costs only, if the preceding trial was rated sure. Recent studies propose that correct guesses (i.e. correct responses with low confidence) can be made based on a feedforward sweep of processing, whereas correct responses with high confidence depend on the establishment of long distance feedback loops between the occipito-temporal and parieto-frontal network (Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006). Therefore, it is possible that correct responses with unsure or medium sure rating lead to activation of the fronto-parietal network, but without establishment of feedback connections/loops. Even if the fronto-polar regions triggered a change in dimensional weights, this signal may not be propagated back to the extrastriate areas in order to take effect. So in contrast to correct responses with high confidence, where feedback connections are established and the weighting signals from fronto-polar signals can be propagated back, there are no change of dimensional weights triggered by trials with lower than medium confidence.

Still, the results of Experiment 1 are in contrast to previous studies, which failed to find cueing benefits in a localization tasks, employing a reaction time paradigm (e.g. Mortier et al., 2007). They found reliable cueing effects in the detection task, replicating previous findings (Müller et al., 2003; Theeuwes et al., 2006). In the localization task in contrast, there were no reliable reaction time benefits for valid over neutral or invalid cues. This raises the question, why Experiment 1 found evidence for cueing effects in a localization task whereas others (Mortier et al., 2007) did not. Mortier et al. (2007) assume that there are different processing routes for detection and localization tasks. In their view the saliency map is involved only in processing of the localization task, as for solving that task spatial information is required, whereas detection tasks are processed via a different route that does not involve the saliency map. They further assume that weighting mechanisms leave the saliency map unaffected, but only modulate processing of the second (detection) processing route. A similar view has been recently proposed by Chan and Hayward (2007), who base their argument on differences in dimension switch costs between detection and localization tasks. However, such a view is inconsistent with the finding of cueing benefits in the

localization task of Experiment 1.

A different reason for these discrepant findings may be the amount of time observers had to localize the target. While in Experiment 1 the display was masked after a short period of time, in Mortier et al. (2007) had unlimited time to discriminate the target position. However, cueing benefits are perhaps obscured by ceiling effects as localization of a highly salient pop-out target is a very efficient process. But ceiling effects are unlikely to cloud effects of cues in Experiment 1, as accuracy was limited to 75%. Nevertheless, given that Experiment 1 employed a SDT paradigm it is possible that, with reaction times ceiling effects may diminish cueing benefits. The latter possibility was examined in Experiment 2, in which the relative saliency of targets was manipulated, increasing difficulty of the task. If with reaction times cueing effects are diminished by ceiling effects due to highly salient targets, then cueing effects should reemerge under conditions in which the singleton target's saliency is low.

4.2 Experiment 2

Strength of feature contrast is known to modulate discriminability of targets from distracters in FEF neurons (Sato et al., 2001; Schall & Thompson, 1999). For targets that produce stronger feature contrast (e.g. a red among green disks), the firing patterns of FEF neurons for targets separate earlier in time than for distracters, as compared to targets that produce weaker feature contrast (e.g. a red among orange disks). FEF neurons, which most likely implement a saliency map, can discriminate targets from distracters faster and more accurately, the greater the dissimilarity between them. Applied to the present question, if dimensional cueing effects are perceptual effects, maybe obscured by ceiling effects in speeded localization tasks (Mortier et al., 2007), then they should reappear, if saliency of pop-out targets is reduced.

4.2.1 Method

Participants

15 observers (5 male) with age ranging from 19 to 47 with a median of 24, all right-handed, participated in Experiment 2 and were paid 8 Euro (ca. 11 USD) per hour.

Apparatus, Stimuli, and Timing

The apparatus in Experiment 2 was the same as in Experiment 1. Distracters were 34 grey bars tilted 45° to the left. The bars were arranged on three (invisible) circles with a radius of 4.5° , 8.5° , and 12.5° of visual angle around a 0.2° white fixation spot. The inner circle consisted of 6, the middle circle of 12, and the outer circle of 16 items. Targets could be placed on positions 2, 3, 4 (right half of the display), and positions 8, 9, or 10 (left half of the display) on an imaginary clock superimposed over the middle circle. The bars extended 0.3° in width and 1.3° in height. A target was present on each trial. Targets were either orientation defined (tilted 35° to the left or 45° to the right from vertical, producing an orientation contrast of 10° and 90° against the 45° left tilted distracters, respectively) or they were luminance-defined, with targets differing from distracters either by a low or a high intensity, of 10 cd/m^2 or 47 cd/m^2 , respectively (relative to an intensity of 5 cd/m^2 for distracters).³ Presentation of cues and timing of the trials was the same as in Experiment 1.

4.2.2 Design and Procedure

In Experiment 2 I manipulated the target dimension (orientation or luminance), feature contrast of the target (low or high), and validity of dimensional cues (valid or invalid). The target was equally likely to be orientation or luminance defined, and to

³These target features were chosen based on a pilot experiment, in order to ensure that there was no reaction time difference between dimensions, only for the different levels of feature contrast.

be of low or high feature contrast. Trials were presented in 18 blocks of 72 trials each. Validity of the dimensional cues was 80%. Participants received written instructions about presentation and validity of the cues and were asked to prepare themselves to encounter a target defined in the dimension of the cue, even though dimension was irrelevant for the task.

4.2.3 Results

Overall error rates were low (2.4%). An ANOVA on the error rates with the factors dimension (luminance, orientation), contrast (low, high), and validity of cue (valid, invalid) revealed that participants performed better when feature contrast was high compared to low (1.0% vs. 3.9% errors respectively), $F(1,14)=16.02$, $p < .001$, main effect of feature contrast. Luminance targets led to significantly less errors than orientation targets (1.9% vs. 3.8% respectively), $F(1,14)=8.17$, $p < .012$, main effect of dimension. No further effects were significant.

Reaction times and cueing benefits for the luminance and contrast manipulations are presented in Figure 4.3 and 4.4, respectively. An ANOVA on reaction times with the factors of cue, dimension, and feature contrast revealed a significant main effect of feature contrast, $F(1,14)=57.68$, $p < .001$, and a tendency for cue, $F(1,14)=4.14$, $p < .06$. With a mean reaction time of 402 ms, high contrast targets were responded to faster than low contrast targets (360 ms vs. 443 ms) independently of target dimension. Most importantly the interaction between cue validity and contrast was significant: $F(1,14)=6.35$, $p < .025$. This effect indicates that there was a substantial cueing effect in the low feature contrast condition, but no such effect in the high feature contrast condition (12 ms vs. 1 ms).

4.2.4 Discussion

Experiment 2 was motivated by the question, whether dimensional cueing effects in a speeded localization task re-emerge, when saliency of pop-out targets is less than

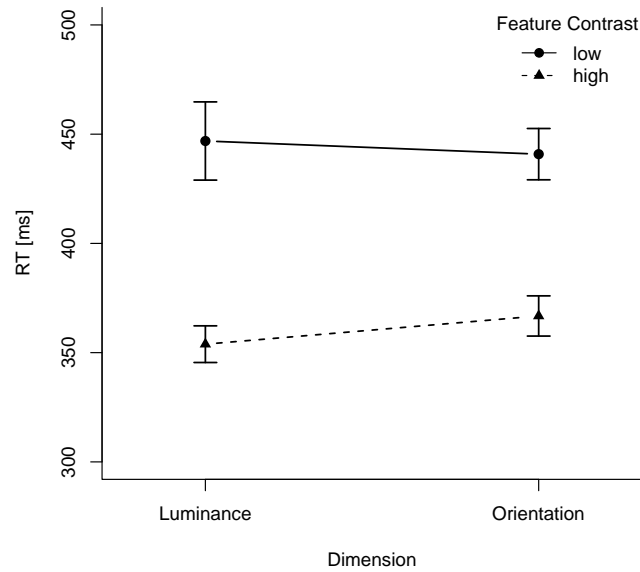


Figure 4.3. Reaction times for luminance and orientation defined targets of high (triangles and dashed line) and low (solid line and bullets) feature contrast.

maximal. We manipulated saliency of targets in varying the similarity of target and distracter features. Orientation targets could differ from distracters by an orientation contrast of 10° or 45° , and luminance targets by an intensity increment of 6 or 47 cd/m^2 , respectively. Participants were instructed to use a symbolic cue (valid in 80%) in order to prepare themselves for the dimension of the upcoming trial. The manipulation of feature contrast revealed that there was a ca. 100 ms difference in latencies for the two different levels of feature contrast. Note that this difference is unlike to be due to a serial type of search with targets of low feature contrast as Zehetleitner (2007a) have shown that with orientation contrast of even 6° between targets and distracters lead to efficient (i.e. pop-out) search.

Most importantly I found cueing effects to interact with feature contrast. For high feature contrast targets I replicated the findings of Chan and Hayward (2007) and Mortier et al. (2007) and found no cueing effects at all. For low feature contrast




Figure 4.4. Cueing benefits (i.e. reaction times of invalidly cued targets minus reaction times of validly cued targets) for targets of high and low feature contrast.

targets, however, I did find significant cueing benefits. Thus, the absence of cueing benefit in speeded localization tasks is likely not a qualitative, but a quantitative effect. That is, cueing effects are not absent per se in such a task, but reappear, if feature contrast of targets is less than maximal. Related to this, the findings from Experiment 2 also argue against the assumption of different processing routes for detection and localization tasks (Chan & Hayward, 2007; Mortier et al., 2007). This 'dual route' hypothesis is further problematic because of findings from the redundant-signals paradigm (Krummenacher et al., 2001, 2002), for a review see Zehetleitner (2007c).

4.3 General Discussion

4.3.1 Summary of findings

The present study aimed at reexamining the question whether pre-attentive vision is top-down penetrable. Regarding this, dimension weighting account (Müller et al., 1995; Found & Müller, 1996) proposes that pre-attentive vision can be modulated by previous experience (dimensional intertrial effect) as well as by top-down intentions (e.g. cueing of upcoming dimension for each trial). This view contrasts the idea of an implicit memory system that is inaccessible to intentional control (e.g. Maljkovic & Nakayama, 1994). If any, cueing effects were attributed to post-selective rather than pre-attentive processes (Theeuwes et al., 2006; Mortier et al., 2007).

In Experiments 1 and 2 a pop-out target was present in each trial and participants had to determine whether the target was present in the left or in the right half of the display (i.e. they had to localize the target). Therefore, a localization task is a kind of compound task (Duncan, 1985; Bravo & Nakayama, 1992), in the sense that the target defining and the response defining features are dissociated from each other. Thus, a localization task can be used to dissociate pre-attentive from response-related (post-selective) effects.

In order to examine effects of top-down knowledge on performance, a trial by trial cueing procedure was used. In Experiment 1 participants had to report the location of a target under brief viewing conditions. As with reaction time measures it is possible that cueing effects stem from pre-attentive as well as post-selective processing stages (e.g. Santee & Egeth, 1982; Prinzmetal et al., 2005). However with brief presentation of displays and measurement of accuracy one can measure pre-attentive effects unconfounded by post-selective processing stages.

In Experiment 1, valid cues significantly increased accuracy in comparison to invalid cues. The same pattern was observed in reaction times, although participants were not instructed to respond in a speeded manner. The repetition of target dimension

led to significantly higher accuracy than the change of target dimension. This pattern could also be observed for overall reaction times, mainly stemming from the localization response. That is, participants located the target faster for sequences in which the target defining dimensions was repeated, than when it changed.

In Experiment 2 participants had to localize targets as quickly as possible, and as in Experiment 1, a cue indicated the likely dimension of the following target. In Experiment 2 I manipulated the saliency of the pop-out targets, in order to examine whether previous studies did not find cueing effects in speeded localization tasks (Mortier et al., 2007) possibly because of ceiling effects. Indeed, I found substantial cueing effects for targets that were defined by feature contrast less than maximal, in contrast for targets with maximal feature contrast I replicated the (null-)finding of no cueing benefits found by previous studies. and thus can exclude theories that assume a qualitative difference in processing between detection and localization/compound tasks rather than only a quantitative one.

Overall, these findings are in line with the dimension weighting account (Müller et al., 1995; Found & Müller, 1996). According to this account, feature contrast signals are modulated by dimension specific weights before they are pooled into the master salience map. Dimensional weights are considered to be changed in two fundamentally different ways: passively in a bottom-up fashion (by trial history) and actively in a top-down fashion (by intention). The dimension weighting account further assumes that the salience map is involved both detection and localization of a target. Applied to the present findings, for validly cued or repetition of the target defining dimension the input signals into the salience map are stronger due to increased modulation by dimensional weights. Hence, the detection or localization process improves in performance, both in terms of latencies and accuracy. In other words, improvement of detection/localization is assumed to occur at pre-attentive processing stage.

The findings of the present study support the core assumptions of the dimension weighting account. In Experiment 1, changes of the target defining dimension across two consecutive trials led to less accurate performance than repetition of dimension,

and additionally, prior knowledge of the following target defining dimension increased accuracy. So, both ways of altering dimensional weights can be observed: passively, in a bottom-up fashion, and actively, in a top-down fashion. Secondly, the pre-attentive nature of both effects is also supported by the present findings: both the brevity of visual presentation renders the possibility of post-selective or response-based effects unlikely (Santee & Egeth, 1982; Prinzmetal et al., 2005). Thirdly, Experiment 2 verified that cueing effects also can be observed in speeded localization tasks, supporting the view that processing of both detection and localization involves the salience map.

4.3.2 Relations to post-selective accounts

There are several authors, whose theoretical positions exclude the possibility of top-down modulation of pre-attentive vision. For instance, Theeuwes et al. (2006) found cueing benefits in a pop-out detection task, replication the findings of Müller et al. (2003). However, in compound task (i.e. where the target's perceptual and response defining features are distinct) the authors found no reliable cueing benefit. The latter finding lead (Theeuwes et al., 2006) to concluded that the top-down cueing effects, as observed in Müller et al. (2003), were not pre-attentive (perceptual), but post-selective (probably response-based) effects, as suggested, for example, by the dimensional action system (DA: Cohen & Feintuch, 2002; Cohen & Magen, 1999; Cohen & Shoup, 1997). The DA model assumes that there are dimension-specific feature analyzer units as well as multiple response selection units, one per visual dimension. 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. In order to select one response module, the activity of which is further transferred to an executive stage in the DA model requires focal attention. Both Theeuwes et al. (2006) and Cohen and Feintuch (2002) agree that both the detection and compound type of tasks involve a search process.

Unlike the detection task, in compound tasks, the dimension-based response selection modules of the DA system are ineffective, because they are unrelated to target's response properties (but to its perceptual properties). As dimensional cueing effects according to Theeuwes and colleagues occur only in detection but not in compound tasks, they have to arise from the response selection process, as the search process in both tasks is similar.

However, this theoretical position is challenged by the findings of the present study. Particularly, in Experiment 1 I found substantial benefits from dimensional cues in a localization task, in which the perceptual and response properties of targets were separated. This excludes the possibility of DA-like response selection modules as being responsible for the dimensional cueing effects. Furthermore the results from Experiment 1 cast doubt on any account assuming the cueing effects to arise from a post-selective processing stage. This is because in this experiment displays were presented for only 60 ms succeeded by a mask, which rendered deployment of attention for detailed analysis (i.e. post-selective processing) impossible (Santee & Egeth, 1982; Prinzmetal et al., 2005).

Moreover, Experiment 2 showed that reaction time cueing benefits in speeded compound tasks are modulated by feature contrast. That is: although Theeuwes et al. (2006) did not find reliable benefits from dimensional cues in a compound task (even if numerically present - see their Experiment 3), Experiment 2 showed that cueing effects may have well been present but dwarfed by very short decision times due to the high saliency of the targets (see also Müller & Krummenacher, 2006). In summary the finding of reliable benefits from dimensional cues in Experiment 1 under limited viewing conditions rules out any model that assumes post-selective processes to be responsible for dimensional cueing benefits (in particular the DA system of Cohen & Feintuch, 2002).

4.3.3 Relation to further studies

There are two previous studies that have examined sequence and expectancy effects (Huang & Pashler, 2005) or only sequence effects (Sigurdardottir et al., 2007) in visual search with brief display durations. Huang and Pashler (2005) presented an orientation pop-out target for a brief period of time, followed by a mask, and participants were required to report the rough location of the target (left or right half of the display), as in our Experiment 1. Huang and Pashler argue that bottom-up factors (in terms of target feature repetition) and top-down factors (in terms of expectancy or prior knowledge) are confounded most of the time. For instance, in homogeneous trial blocks, in which the target is constant, both bottom-up and top-down effects are operating, the later because the observer knows that the upcoming target will be the same as the previous one (i.e. repetition plus expectancy). In heterogeneous blocks, in which the target-defining feature changes unpredictably across trials, expectation does not work, but there may or may not be repetition of the target-defining feature (i.e. no expectancy / no repetition or no expectancy / repetition). A fourth possibility with expectation but no repetition occurs, if the trials are arranged in an *ABAB...* sequence, such that the observer knows in every trial what the upcoming target will be, while the target definition always is different from the preceding trial. Using a Nakayama paradigm (Bravo & Nakayama, 1992; Maljkovic & Nakayama, 1994), Huang and Pashler (2005) examined all four combinations of expectancy and repetition conditions. They found performance to be comparable in three of four conditions: in heterogeneous blocks, in which top-down knowledge is ineffective, they found no difference between repetition or swap of the target defining feature (no expectation with and without repetition). Similarly, performance was at the same level as for the no expectation condition in blocks with alternating sequences, in which the target-defining feature changed predictably on every trial (i.e. expectation without repetition). The only condition, where performance is improved, is in homogeneous blocks, in which the target stays the same, and the participant can expect the target

to stay the same. The authors conclude that a perceptual advantage only arises, if both expectancy and repetition is present, and hence call this perceptual benefit the expectation-repetition effect. In a last experiment, Huang and Pashler examined the previous manipulations in a speeded reaction time localization under unlimited viewing conditions. If participants had to respond as quickly as possible, reaction times were fastest in homogeneous blocks, in which both expectation and repetition acted together. In addition there was a reaction time benefit for repeated targets over non-repeated targets in heterogeneous blocks, replicating the finding of Maljkovic and Nakayama (1994). Due to the fact that repetition effects did indeed occur without expectancy (i.e. in heterogeneous blocks) in a speeded reaction time paradigm with unlimited viewing conditions, but were absent in a signal detection paradigm with only brief stimulus presentation, Huang and Pashler (2005) follow that the target-distracter feature swap cost reported by e.g. Maljkovic and Nakayama (1994) reflects mainly post-selective and postperceptual factors. A perceptual (pre-attentive) effect in that view only occurs, if both repetition and expectation are present.

Sigurdardottir et al. (2007) on the other hand did find repetition effects in heterogeneous blocks (where top-down knowledge is ineffective) in both a conjunction search and a compound task. In their study the displays were also presented only briefly, followed by a mask. In the conjunction search task, Sigurdardottir et al. were able to discern whether repetition effects alter sensitivity or response criteria: only the SDT sensitivity measure d' was increased for streaks of repeated target definitions, whereas the response criterion was unaffected. In their compound task, a pop-out target was color defined and the target and distracter color could swap from trial to trial. The response defining attribute was the location (left or right) of a small hole in the stimuli that could be displaced between 0.05° to 0.2° from the center of each stimulus. They found accuracy to increase, as the target definition was repeated, irrespectively of the acuity demands of the task. Why did Sigurdardottir et al. (2007) find repetition effects in condition with heterogeneously defined targets under brief viewing

conditions, whereas Huang and Pashler (2005) attributed such advantages to post-selective processing stages? One critical feature of the experiments of Sigurdardottir et al. (2007) was the predictability of the target definition to change. In order to produce streaks of repeating target definition, the authors manipulated the probability of repetition ranging from 75% to over 90%. Therefore participants were safe to expect the target definition to repeat. In that respect the experiments of Sigurdardottir et al. (2007) resemble more the expectation conditions of Huang and Pashler than the non-expectation conditions (in which target repetitions occurred with a probability of 50%). And according to their account, perceptual advantages (measurable in sensitivity or accuracy under brief viewing conditions) can occur if both expectation and repetition are present. If this indeed were the case, the target-distracter feature swap effect originated from Bravo and Nakayama (1992) and theoretically deepened by Maljkovic and Nakayama (1994) would not be a perceptual, but a postperceptual effect, because it occurred in situations, in which the target/distracter swaps could occur randomly. According to Huang and Pashler (2005) in such a situation no pre-attentive effects could be found under brief display durations and consequently the reaction time effects reported by Maljkovic and Nakayama (1994) would have to arise from post-selective processing stages.

In the present study I employed a different procedure in order to manipulate top-down expectations. Instead of comparing conditions, in which expectations would help (homogeneous blocks) with conditions, in which expectation is of no use (heterogeneous blocks), I cued participants on a trial by trial basis (see Müller et al., 2003). So in addition to the general expectation condition of Huang and Pashler, I could examine performance in a situation, where top-down expectation was present - but false. Neutral cues are similar to the no-expectation heterogeneous condition of Huang and Pashler. According to their expectancy-repetition effect, they would predict no accuracy difference between repetitions and changes of the target dimension in blocks with neutral cues, because although repetition effects may play a role, there is no top-down preparation (expectation) possible. In the valid cue condition,

they would predict repetition effects to emerge, because there both repetition and expectation are present. It is unclear, what they would predict in the invalid cue condition. Although expectation is present, it is wrong, so maybe they would also predict no dimension switch costs to be measurable in accuracy. In summary they predict an interaction between dimension switch costs and type of cue. Only for valid cues, where both (correct) expectation and repetition are present, dimension switch costs should be observable, but not for invalid or neutral cues, because there, although repetitions may happen, expectation is either wrong or not present at all.

The findings of Experiment 1 clearly did not fit these predictions: I found dimension switch costs for invalid as well as for neutral cues (there was no interaction between intertrial transition and cue validity). Müller et al. (2003) reported an interaction between cue validity and intertrial transition, but dimension switch costs were larger for neutral than for dimensional cues (valid or invalid). This is in contrast to the prediction of Huang and Pashler (2005) who would predict switch costs only when repetition and expectation coincide that is only for valid, but not for invalid or neutral cues. Thus, the dependence of repetition effects on both repetition and expectancy as described by (Huang & Pashler, 2005) does not apply to dimensional uncertainty. Additionally it is open to question, how the account of Huang and Pashler (2005) would explain the difference in performance for correct and incorrect expectations (induced by valid and invalid cues). Based on this discrepancy it is probable that situations in which the target defining dimension is uncertain differ from situations, in which target and distracter features swap roles unpredictably. The former affects pre-attentive processing stages and is described best by the dimension weighting account (Müller et al., 1995), whereas the latter affects (mostly) post-selective processing stages and is described best by the account presented by Huang and Pashler (2005). This depends on whether the repetition effects found by Sigurdardottir et al. (2007) in accuracy and sensitivity measures is actually caused by the high probability of target repetitions. Kristjánsson (2007) for instance argues that the localization task in Huang and Pashler (2005) has different acuity requirements than the compound task used

by Sigurdardottir et al. (2007), and that this difference in tasks may be responsible for the null findings of Huang and Pashler. A further argument for a pre-attentive repetition effect stems from neuronal recordings, where FEF neurons could discriminate targets from distracters faster and more accurately for repeated compared to swapped target-distracter features (Bichot & Schall, 2002). However also in this study, repetition effects were examined in a situation, where expectancy could play a role: although for some blocks the probability of repetition was 50%, in some blocks, repetition occurred in 75% of trials, or every 10 trials. Consequently, the repetition- expectation account proposed by Huang and Pashler (2005) could explain the observed behavioral and neuronal pre-attentive performance of Bichot and Schall (2002) as well. The neuronal benefits of Bichot and Schall (2002) can be understood as pre-attentive, in the sense that the activity of visually responsive neurons is thought to reflect a representation of saliency, which in turn is understood to guide allocation of attention and directing of saccades. Until further research it remains at least possible that in dimensional uncertainty conditions pre-attentive processing stages are affected by expectancy and repetition effects, whereas in situations, where targets and distracters can unpredictably swap roles, post-selective processing stages are affected.

4.3.4 Summary and conclusion

In summary the present study reported performance benefits for dimensional cues in a localization task for accuracy under brief viewing conditions as well as for reaction times under unlimited viewing conditions. These findings contradict theories which assume that pre-attentive processing in visual search is inaccessible to top-down control, as well as theories that assume qualitatively different processing structures for detection tasks on the one hand and localization (and compound) tasks on the other hand. The model best in accordance with the present data is the dimension weighting

account (Müller et al., 1995; Found & Müller, 1996), which assumes that dimensional feature contrast signals are modulated by dimensional weights before they are pooled into a saliency master map. For different tasks such as detection, localization, or compound, activity of the master map is used to further process the task. The dimensional weights can be modulated by both passive bottom-up and active top-down processes and their effects under conditions of dimensional uncertainty are basically pre-attentive.

5. DECISION PERSPECTIVE ON SALIENCY

All animals are limited in their possibilities to interact with the environment. Humans and other primates have two hands to grasp or point with, one head to turn, one pair of eyes with one fovea each to look at with, while the amount of possible locations and objects in the environment to use these effectors on seems unlimited in comparison. Additionally our processing system to control behavior, the brain, is limited in its capacity to deal with the environmental complexity, as well (Tsotsos, 1990; Rensink et al., 1997). Given these limitations of animals with regard to the complexity of the environment, it is necessary to adaptively solve a problem of selection. For an adaptive control of selection, there has to be a trade-off between responding to the affordances of the environment and to choose actions based on the internal state of the organism. The effect of internal states can operate on the sensory level by enhancing certain sensory properties that are known to be more relevant over others, or on a semantic level, for which object recognition is necessary.

Both response and choice behavior considered in the present study include binary decisions about the absence or presence, or the rough location (e.g. left vs. right) of a target, as well as selection decisions. Selection decisions determine where in space to move an effector or where to commit the limited processing capacity (i.e. covert attention) to next.

The message of the present paper is the following: The size of observable performance differences due to modulations of sensory processing depends on the duration of the relevant decision process, i.e. the observed differences are larger, if the relevant decision needs more time, and are smaller, if the relevant decision process is fast. Although I apply this principle to search tasks in the visual modality with manual responses, I propose that the principle is fairly general and predict that it is valid for different sensory modalities and all kinds of tasks involving various effectors indepen-

dent on how the duration of the decision process is manipulated. In order to present evidence for this principle, I will first review cognitive models and their neuronal implementations of what decisions can be based on, second how these decision processes can be mathematically modeled, and third what kind of tasks I apply this principle to in the present study.

Both binary decisions (e.g. about presence/absence, left/right location of a target) as well as multiple decisions (where to move an effector or attention next) depend on determining where the most interesting and important location is, at any given moment - depending both on the properties of the environment, as well as on the internal state of the animal. The solution to the problem of selection convergingly assumed by cognitive (Treisman, 1986, 1988; Wolfe, 1994; Müller et al., 1995; Theeuwes et al., 2004) and neuro-cognitive/computational (Koch & Ullman, 1985; Itti & Koch, 2000; Li, 2002) models is a topographical representation of the environment that signals distinctiveness of locations irrespective from what sensory properties this distinctiveness originates from (i.e. a saliency map).

In order to serve the goal of providing a prioritization of locations in terms of relevance, a saliency map must have some general features. First of all it has to take into account properties of the environment. Locations that are very similar to its surroundings do not contain much information and should receive low priority. On the other hand, locations, which differ from its surroundings significantly should be considered as potentially interesting. But only taking into account sensory input in order to determine interestingness or relevance of locations is clearly insufficient. Relevance is not only determined by properties of the environment, but also by the current state of the organism. The current intentions and needs greatly contribute to the assignment of priority to locations. In cognitive psychology the former determinants of relevance usually are called bottom-up factors, whereas the latter are called top-down factors. The most explicit model that describes a master saliency map is the computational model of Itti and Koch (2000). They assume that initially local feature contrast is extracted in parallel for all locations and a number of features (i.e. different colors,

orientations, motion directions, intensities of luminance). At each location not only strength of feature presence is taken into account, but also the difference between presence of a feature at a given location and its immediate surroundings. That way the local distinctiveness can be calculated. The feature contrast signals are then pooled into dimension specific maps, which are again further pooled into a master saliency map. Signals from dimensions (Müller et al., 1995; Wolfe, 1994) and from features (e.g. Itti & Koch, 2000) can be modulated with dimension or feature specific weights. The topographical representation of the visual scene on the saliency map signals strength of feature contrast or distinctiveness at each location in a feature unspecific manner. If there is strong activity at one location of the saliency master map, this activity carries no information about what feature this contrast derives from, e.g. whether from a red among green spots, or one single moving among stationary items. The neural implementation of the cognitive and theoretical concept of a master saliency map is not yet fully clear. Evidence suggests that there exists more than one single neural structure, which has the features of a saliency map. Structures that are currently hypothesized to implement a saliency map focussed on the oculo-motor network (Fecteau & Munoz, 2006) are specifically the superior colliculus (McPeck & Keller, 2002a, 2002b), the lateral intraparietal area (e.g. Gottlieb, 2002), and the frontal eye fields (FEF: Thompson & Bichot, 2005; Bichot & Schall, 1999).

The typical behavior of visually responsive FEF neurons is depicted in Figure 5.1 for weak and strong pop-out targets. Importantly, for the first 100 ms the firing patterns of a distracter or a target in the receptive field of a neuron is indiscernible. Only after a certain time, earlier for strong compared to weak targets, the firing patterns diverge and the neuron is said to discriminate targets from distracters. That is, the firing patterns from highly salient targets lead to a faster and better discriminability for visually responsive FEF neurons. In terms of a saliency map these FEF neurons topographically represent the strength of feature contrast over the visual scene irre-

spective of the feature it originated from (cf Schall, Hanes, Thompson, & King, 1995; Bichot & Schall, 1999, for color and form targets).

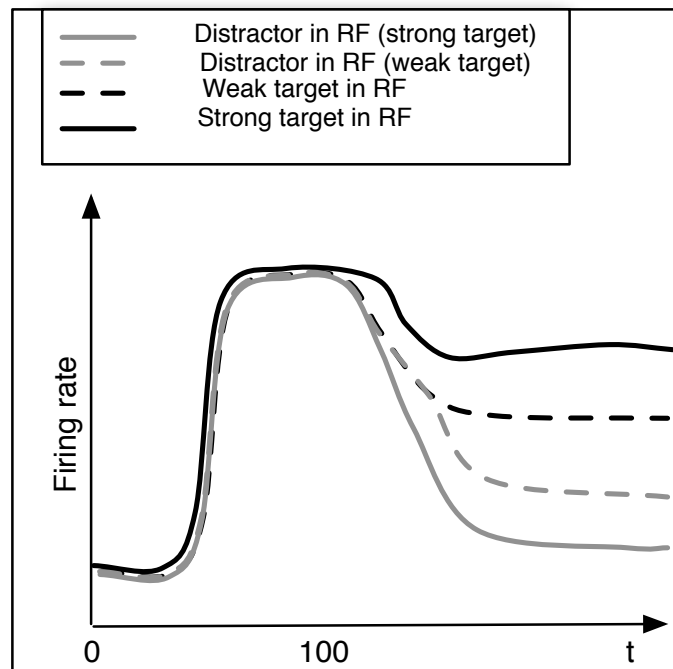


Figure 5.1. A cartoon of firing patterns from visually responsive FEF neurons for weak and strong targets or distracters in the context of weak or strong targets.

After having reviewed models of what simple decisions can be based on, I present mathematical models of the decision process itself. The relevant decisions further on are simple binary (is a target present/absent, left/right, red/green) or multiple decisions that can be made within a second. These decisions are pre-attentive in the way that they specify the next location that should be selected as a movement goal (e.g. for eyes, hand, or head) or for assignment of the limited processing capacities of the brain (covert attention). Binary decisions can be pre-attentive, if they do not need deployment of overt or covert attention (e.g. absent/present decisions).

The most prominent feature of such decisions is that they are not deterministic. The same stimuli and the same task lead to different decision latencies from trial to trial, in some cases errors are made. This provides a number of measures that describe

the performance of decision behavior: mean latencies, the distribution of latencies for correct and incorrect responses and error rates. There exist mathematical models that can capture most of these properties. The basic assumption of these models is that decisions are based on noisy signals. In contrast to e.g. models of signal detection theory the decision is not based on one single sample of evidence (which can vary from trial to trial) but on a continuously monitoring of the response relevant signals. A response is triggered (i.e. a decision is made), as soon as enough evidence for that response has accumulated. Thus these models are of the class of sequential-sampling models with two subclasses: random-walk models and accumulator (and counter) models (Ratcliff & Smith, 2004). Random-walk models have a relative stopping rule, i.e. there is one accumulator that monitors the noisy signal and evidence for one alternative simultaneously is evidence against the other alternative. A successful class of these models is the Ratcliff diffusion model (RDM: a type of Wiener diffusion model, Ratcliff, 1978). In contrast, accumulator models assume one accumulator for each decision alternative. If these accumulators are independent from each other, they cannot predict the shape of reaction time distributions or account for error responses sometimes being faster than correct responses (Ratcliff & Smith, 2004). The leaky competing accumulator model (Usher & McClelland, 2001) can account for these properties of the behavioral data in assuming that the separate accumulators for each response are not independent from each other, but connected with mutual inhibition (i.e. evidence for one alternative is evidence against the other).

We will describe the components of the RDM in more detail. First of all, observable response latencies T_{Obs} are assumed to be the sum of the time needed for making the decision T_{dec} and the time needed for stimulus encoding, response preparation, and motor execution (i.e. all other processes apart from the decision process), T_M . T_{dec} is the time needed until evidence reaches the criterion for response one or the criterion for response two (a_1 and a_2 , or zero, respectively, in Figure 5.2). In average, evidence accumulates with a certain drift rate ν (depicted as the gray line in Figure 5.2). Within each trial this drift rate is assumed to vary, leading to different

paths, how evidence is accumulated. Although stemming from the same stimulus, decision latencies can be faster, slower, or erroneous. In order to capture fast errors, additionally the starting value z is assumed to vary from trial to trial, representing e.g. small variation in the criterion of the observer. So sources of variability of decision latencies stem from within trial modulations of the drift rate, and from between trial variations of the starting value of evidence.




Figure 5.2. Exemplary sample paths with the same mean drift rate (depicted in gray). The stars mark, where the decision process would terminate with more liberal criterion a_2 .

For instance, the RDM has been applied to a huge variety of behavioral tasks: lexical decision (Ratcliff, Gomez, & McKoon, 2004; Gomez, Ratcliff, & Perea, 2007), feature discrimination (Ratcliff & Rouder, 1998; Ratcliff, 1981), or recognition memory (Ratcliff & Smith, 2004; Ratcliff, 1978). There are recent efforts to apply these types of models not only to behavioral but also to single-cell recording data (Ratcliff,

Hasegawa, Hasegawa, Smith, & Segraves, 2007; Ratcliff, Cherian, & Segraves, 2003; Mazurek, Roitman, Ditterich, & Shadlen, 2003), aiming at a mathematical model that unifies behavioral and neuronal data. In these studies, it was possible to predict the firing patterns of neurons in the superior colliculus with a RDM that has been fitted to behavioral data of monkeys. That is the same parameter set could describe behavioral as well as neuronal data patterns.

The property of sequential-sampling models particularly relevant for the present study is presented in Figure 5.3. If two drift rates differ by a value of α , the resulting latency difference in decision times depends on the duration of the decision process. If a drift rate already high receives an increase of α , this leads to a latency difference X_1 , which is smaller than the latency difference, Y_1 , resulting from a low drift rate receiving the same increase. Additionally the difference α leads to smaller differences in decision latencies, Y_2 , with a more liberal criterion, a_2 , than Y_1 resulting from the more conservative criterion, a_1 . So in summary, the same numerical difference in drift rates (α) leads to different decision latency differences, depending on the overall duration of the decision. The duration of the decision can be manipulated via stimulus properties, i.e. by making the decision more difficult (increasing the drift rate), or via properties of the observer, i.e. the readiness to respond (increasing the criterion).

This is a fairly general principle that should hold for a variety of tasks, which lead to performance that can be captured by e.g. the RDM.

In the present study, I will apply this principle to visual pop-out search tasks (Treisman & Gelade, 1980), examining assumedly pre-attentive effects of redundancy gains (Krummenacher et al., 2001, 2002; Zehetleitner, 2007a), benefits from dimensional cues (Müller et al., 2003; Theeuwes et al., 2006; Müller & Krummenacher, 2006; Zehetleitner, 2007b; Mortier et al., 2007), and dimension switch costs (DSCs: Müller et al., 1995; Found & Müller, 1996; Chan & Hayward, 2007). In visual search tasks,




Figure 5.3. The same difference in slope α leads to latency differences X_1 and Y_1 with criterion a_1 for fast or slow drift rates. With a more liberal criterion a_2 , the latency difference Y_1 is larger than Y_2 , still with the same drift rate difference of α .

participants are presented with a number of stimuli, some of which are distracters, some of which are targets. In pop-out searches, the target differs from the distracters in one or more different dimensions (cf. Wolfe & Horowitz, 2004), for example a red item among green or a moving left-ward among moving right-ward items. With pop-out targets it is possible to design several tasks: e.g. detection, coarse localization, or compound tasks. For detection, observers have to determine in each trial, whether a target is present or absent and respond to each of both possibilities. In localization tasks a pop-out target is present in every trial, and the participants are instructed to respond to the rough location of the target, e.g. left vs. right or upper vs. lower half of the display. Compound tasks are a combination of two tasks: selection of the target (with covert or overt attention), and report of a property of that target. This property can be for instance the orientation of a line segment placed inside of the

stimulus (e.g. Theeuwes, 1992), or the side (left/right) a target is chipped off (e.g. Bravo & Nakayama, 1992; Maljkovic & Nakayama, 1994), or the dimension the target is defined in (e.g. Found & Müller, 1996). All of these tasks are relatively simple and can be solved mostly within one second.

In relation to the cognitive concept of saliency maps, in a detection task activity on the master map could be used to decide whether a target is present or absent. This type of processing is assumed by e.g. Müller et al. (1995) and Found and Müller (1996), although also an alternative processing route via spatially pooled dimensional modules has been proposed (Treisman & Gelade, 1980; Treisman, 1988) and been recently revived (Chan & Hayward, 2007; Mortier et al., 2007). Especially in situations, where the target can be defined in a different dimension on each trial, the featureless representation of distinctiveness, a main property of saliency maps, may be useful to solve the task. Localization tasks also could be based on a saliency map, because for a coarse localization the exact features of the target are irrelevant, also. The saliency map is most certainly involved in compound tasks, because in order to respond in such a task, it is necessary to deploy attention (covertly or overtly) to the target location, in order to further extract the response defining feature. Hence in detection and localization tasks, only one decision process is involved, whereas in compound tasks, there are two separate decision processes. The first decision determines what item in the display is subject to further analysis, and the second decision determines the response by extracting the response defining feature of the target.

Based on the properties of sequential sampling models, I predict that modulations of pre-attentive processing will lead to response time differences of different size, depending on the duration of the decision process. The decision process here refers to the task-specific evaluation of signals on the saliency master map. For the detection task, these signals have to lead to a present/absent, in a coarse localization task to a left/right decision. If attention has to be deployed or eyes to be moved to a location (e.g. in a compound task), signals on the saliency master map are evaluated to decide what location to select next.

The pre-attentive effects relevant for the present study are the redundant signals effect (RSE), dimension switch costs (DSC), and benefits from dimensional cues. The RSE is a reaction time benefit of signals that have two target defining features (redundant targets) over signals that have only one (single targets, cf. Todd, 1912; Raab, 1962). A redundant signal in a visual pop-out search is either one target that differs from its surrounding distracters in two dimensions simultaneously (e.g. a red horizontal among green vertical bars, cf. Krummenacher et al., 2001) or two pop-out targets, defined in two different dimensions (e.g. a red vertical and a green horizontal target among green vertical distracters, cf. Krummenacher et al., 2002). Redundant pop-out targets have repeatedly been reported to lead to shorter reaction times, than single pop-out targets (Krummenacher et al., 2001, 2002; Koene & Zhaoping, 2007; Zehetleitner, 2007a). This RSE could not be explained in terms of statistical facilitation (Raab, 1962) in a parallel race. A parallel race would mean that feature contrast signals from both dimensions of a redundant target would race for triggering a response independently. In such a model, as soon as a threshold is reached for one dimension, a response is triggered. Miller (1982) has shown that parallel race models yield an upper bound to how much a redundant signal can gain over single signals, formulated in the race model inequality (RMI). If the cumulative density distribution (CDF) of redundant targets is faster (i.e. left to) the sum of both CDFs for the single targets, then a race model cannot account for the observed RSE anymore. Violations of the RMI with pop-out targets have been reported frequently (Krummenacher et al., 2001, 2002; Turatto et al., 2004; Koene & Zhaoping, 2007; Zehetleitner, 2007a) providing plenty of evidence that redundant pop-out targets are indeed not processed in a race model like architecture. Zehetleitner (2007a) could also exclude two theoretical alternative models that also can produce violations of the RMI: interactive race models (Mordkoff & Yantis, 1991) and serial exhaustive models (Townsend & Nozawa, 1997). Miller (1982) proposed that if the RMI is violated, the processing architecture would be described best with co-activation models, in which signals of parallel processing channels are summed or pooled before triggering a response. A

master saliency map, into which dimension specific feature contrast signals are pooled is an instance of co-activation models. Regarding redundant targets, feature contrast signals are present on two dimension specific maps and are pooled into the master map. That way for redundant targets, activity on the master map builds up faster and to a higher level than for single targets, leading to a reaction time benefit - the RSE. Our prediction regarding this effect is that slowing of the decision process based on the saliency master map leads to a larger observable RSE.

A second effect examined in the present study is the DSC effect. In situations, where the pop-out target can be defined in two or more dimensions with the target dimension changing randomly from trial to trial, reaction times are slower than in situations, where the target defining dimension is predictively constant over a period of time - even if the specific feature is varying, e.g. in a block of orientation targets, defined either by a slant left or right to the vertical distracters (Treisman, 1988; Müller et al., 1995). Found and Müller (1996) examined short sequences of trials and found that reaction times for a target in trial n depend on the dimension the target in the previous trial, $n - 1$, has been defined in. If the target is defined in a different dimension in trial n than in trial $n - 1$, then the reaction time is slowed compared to a repetition of the dimension: the DSC. The dimension weighting account of Müller and colleagues (Müller et al., 1995; Found & Müller, 1996) explains this effect in pre-attentive terms. The dimension weighting account assumes the existence of dimension specific weights that modulate dimension specific feature contrast signals before being pooled into the salience map. Signals from a dimension which is weighted highly has a larger impact on the master map, than signals from a dimension for which the weights are low. The dimension weighting account assumes a passive bottom-up changing of weights from trial to trial to the relevant dimension. If in a specific trial a pop-out target is color defined, then the weight for the color dimension is increased, while simultaneously the weights of all other dimensions are decreased. If in the subsequent trial the target is again color defined, decisions based on the master saliency map can be made faster and more accurately than if the target is orientation defined in the subsequent trial:

the previous target lead to an increase of the weights for color, leading to a larger impact of subsequent feature contrast signals from the color dimension as compared to the orientation dimension. We predict that the size of the DSC effect increases, when the decision process based on the saliency map is slowed.

In addition to passive, bottom-up changes of dimensional weights due to the target definition in the current trial, the dimension weighting account assumes that weights can be modulated intentionally by the observer. This prediction has been tested by use of trial-by-trial dimensional cues (Müller et al., 2003; Müller & Krummenacher, 2006; Theeuwes et al., 2006). In these experiments again the target defining dimension could change randomly from trial to trial, and participants are given a chance to prepare for the next dimension by means of a semantic cue. This cue can be unpredictable (e.g. the word 'neutral') or predictive (e.g. the words 'orientation' or 'color'). Dimensional cues can be valid, i.e. correctly announce the dimension of the next trial (mostly 80% of the cases), or it can be invalid. If participants were unable to make use of these cues, no reaction time difference between valid, neutral, or invalid cues would be expected. In contrast, Müller et al. (2003) for instance found reaction times to be fastest after valid cues and slowest after invalid cues, with reaction times for neutrally cued trials in between. The dimension weighting account explains this data pattern in terms of dimensional weights. A cue makes it possible for the observer to intentionally prepare for the dimension, the next target most probably will be defined in. This preparation leads to an increase of the dimensional weight indicated by the cue. The increase of weight for one dimension leads to a larger impact of the cued dimension on the salience map, resulting in higher reaction speed. So the effect of dimensional weights is invariant to how they have been modulated, passively by influence of the previous trial, or actively by intention of the observer. As benefits from dimensional cues according to the dimension weighting account are pre-attentive effects, I predict them to increase in size, if the decision process, which is based on the salience map, is prolonged. There exist controversies about the nature of especially DSCs and benefits from dimensional cues, as they are reported frequently for

detection tasks (Müller et al., 1995; Found & Müller, 1996; Pollmann, 2001; Meeter & Olivers, 2006; Olivers & Meeter, 2006; Mortier et al., 2007) but are under dispute for compound and localization tasks (Müller & Krummenacher, 2006; Mortier et al., 2007; Chan & Hayward, 2007). We will address this issue in the general discussion. In the present study I will review results from previous studies, report reanalysis from one previous study and present two new experiments that test our prediction that the size of DSCs, the RSE, and benefits from dimensional cues is increased, if the decision process, which is based on the salience map is prolonged.

5.1 The RSE

5.1.1 Manipulation of feature contrast

In Experiment 1 of Zehetleitner (2007a), a feature contrast manipulation was used to implement a double-factorial design (Townsend & Nozawa, 1995; Sternberg, 1969a), in order to test alternative models to co-activation models, which can explain violations of the RMI. Targets were defined by feature contrast in the luminance or in the orientation dimension for single targets, and by a combination of both in case of redundant targets. Both luminance and orientation feature contrast could be either high or low. In an array of vertical dark grey distracters, orientation targets differed from distracters by a slant of 45° (high salience) or 6° (low salience). Targets of two luminance levels were adjusted individually to match reaction times of the two levels of orientation contrast.¹ There were four possible redundant targets: bright- 6° , dim- 6° , bright- 45° , and dim- 45° . In neuronal terms, low feature contrast targets evoked firing patterns, in which target and distracter activation separates later in time and to a less extent than high feature contrast targets. The task was to detect, whether a target was present or absent. The manipulation of feature contrast leads to slower

¹An additional set-size experiment verified that targets with low feature contrast still lead to efficient search.

drift rates for low than for high feature contrast targets in a diffusion model and I am interested in the size of the RSE for fast and slow decision times. Therefore I compare the RSE for the dim-6° and the bright-45° conditions, in which feature contrast of both dimensions that make up the redundant target is the same (i.e. high-high and low-low). In the other two cases one dimension of redundant targets produces a high, and the other a low feature contrast signals. In these situations for a certain time the decision is mainly driven by the feature contrast of the high-contrast dimension, because for that dimension target and distracter activation separates faster. Only after a certain time the second, low-contrast dimension can contribute to the decision process. Hence, mainly the high feature contrast dimension of a redundant target with mixed feature contrasts determines the decision time. Still the low-contrast component has an effect that leads to a significant RSE and violations of the RMI. On the other hand for redundant targets, in which both dimensions produce feature contrast signals of the same strength (either high or low), the time in which both signals being summed/integrated into the master saliency map is longer, because both feature contrast signals have a similar time course.

In Experiment 1 (Zehetleitner, 2007a) I showed that the RSE for the bright-45° condition was 7.52 ms, whereas in the dim-6° condition the RSE was 47.25 ms. Post-hoc analysis of this difference with a paired t-test reveals this difference to be significant, $t(14)=5.5$, $p < 7 * 10^{-5}$. This finding is in accordance of the prediction issued in the present paper: the RSE becomes larger, the longer the decision process takes. In this experiment the decision time was manipulated by strength of feature contrast, influencing the drift rates of a RDM.

5.1.2 Manipulation of response bias

The question of Experiment 3 in Zehetleitner (2007a) was, whether interactive race models can explain the observed violations of the RMI and have to count as an alternative explanation in addition to the co-activation model of saliency summation.

In order to test interactive race models, I manipulated the amount of information that could be exchanged between dimensional modules in varying the present-absent ratio (1:1 vs. 3:1, present:absent), and the ratio of single to redundant targets (1:1 vs. 1:2 - redundant:single).² Relevant for the present study is the manipulation of present/absent ratio. The stimuli and procedure were similar to Experiment 1 of the same study, only that targets were always defined with high feature contrast (45°, bright, or both), and the present/absent and single/redundant ratios were manipulated as between-subjects factors. The task again was a detection task, where participants had to press a button if a target was present and withhold a response, if no target was present (go/no-go procedure). Therefore with a present/absent ratio of 1:1, a button press was as likely as withholding a response, whereas in the 3:1 present/absent ratio, a button press was three times as likely. The criteria for more likely responses are known to become more liberal, leading both to faster reaction times and more mistakes (e.g. Smith & Ratcliff, 2004). We observed evidence for such a criterion shift in Experiment 3: false alarms (i.e. button presses although no target was present) were 1.5% in the 1:1 and 7.2% in the 3:1 present/absent ratio condition. An analysis of variance (ANOVA) of error rates with factors present/absent ratio and single/redundant ratio revealed only the difference between 1:1 and 3:1 present/absent ratio to be significant: $F(1,60)=33.9$, $p < .0001$. Additionally reaction times were detected faster (18 ms) in the 3:1 compared to the 1:1 condition: $F(1,60)=4.76$, $p < .033$. Both the effects of increased reaction times and increased error rates argue for a more liberal criterion in the condition where present responses are more frequent. The RSE of 20.0 ms in the (slower) 1:1 (present:absent) condition, was found to be significantly larger than the RSE of 16.0 ms in the (faster) 3:1 condition, $F(1,60)=4.45$, $p < .04$. This small but significant 4 ms difference in RSE between the faster and the lower condition is in accordance with our prediction of pre-attentive effects being larger for longer decision times. In this experiment, the decision times were manipulated via a shift of the response criterion.

²For additional details see Zehetleitner (2007a).

5.1.3 The effect of spatial attention

A similar pattern of results can be observed in Experiment 3 (a two-choice detection task) of Krummenacher et al. (2002). The goal of that experiment was to examine the prediction of Cohen and Feintuch (2002) that for redundant targets to be integrated (i.e. to violate the RMI), focal attention is needed. Krummenacher and colleagues directly tested this by cueing a quadrant of the display, in which the next target would be presented with a probability of 80%. The cue was a centrally presented arrow and indicated an area of $2.1^\circ \times 2.9^\circ$ in which the target most likely appeared. The cueing manipulation was successful in 11 of 16 participants, who in mean reported presence of a target 14 ms faster if the target appeared at a validly cued compared to an invalidly cued location. Spatial cues are assumed to enhance processing at the cued location and suppress processing at all other locations. For neuronal data this goes along with increased activation patterns of neurons with receptive fields that cover the cued location and decreased activation patterns for other neurons (Luck, Chelazzi, Hillyard, & Desimone, 1997; Connor, Preddie, Gallant, & Essen, 1997; McAdams & Maunsell, 2000; Sengpiel & Hübner, 1999). Krummenacher and colleagues found a RSE of 15 ms at validly cued locations, and an increased RSE of 21 ms at invalidly cued locations. Although this difference is numerical in nature and statistical significance is unknown, it is in line with our prediction of increased effect size for slowed decisions. In this case the slowing of the decision process was induced by the target appearing at an uncued location.

5.2 Benefits from Dimensional Cues

Regarding benefits from dimensional cues I review Experiment 2 of Zehetleitner (2007b), in which observers had to perform a left/right localization task. One target was present in every display and differed from distracters in either the orientation or luminance dimension. We manipulated feature contrast in two levels, low and high, for both dimension each. Targets could be localized faster (360 ms) if they were more

salient and slower (440 ms) if they were less salient: $F(1,14)=57.68$, $p < .001$. For each level of feature contrast, there was no reaction time difference between both dimensions. The target defining dimension was determined randomly for each trial, and participants received a semantic cue (the German words for luminance or orientation) before each trial started. This cue was valid in 80% of the cases. The cueing benefits, reported earlier for detection tasks (Müller et al., 2003; Theeuwes et al., 2006; Müller & Krummenacher, 2006), are calculated as the reaction time benefit of validly cued over invalidly cued trials. Cueing benefits of 12 ms (significantly different from zero) were larger in the low feature contrast condition, than of 1 ms (not significantly different from zero) in the high feature contrast condition, as apparent in the significant interaction between cue and contrast, $F(1,14)=6.35$, $p < .035$. Again this finding of larger pre-attentive cueing benefits for slower decision times is in accordance with our prediction.

5.3 Dimension Switch Costs

We reanalyzed the data of both Experiment 1 and Experiment 3 (Zehetleitner, 2007a) with regard to dimension switch costs. For that analysis only sequences of two correct trials were used, i.e. trials after errors were excluded (less than 1.5% of the data). Additionally I only analyzed trials in which in trial $n - 1$ and in trial n a single target was presented, i.e. I excluded redundant and absent trials in trial n or in trial $n - 1$. For Experiment 1, an ANOVA with between-subjects factors of feature contrast (low, high), and intertrial transition (same dimension) revealed the main effects of feature contrast and intertrial transition to be significant, $F(1,14)=87.8$, $p < .0001$ and $F(1,14)=76.1$, $p < .0001$, respectively (i.e. the mean DSC effect of 19 ms was significant). Intertrial transition interacted significantly with feature contrast, $F(1,14)=6.2$, $p < .025$: DSCs were higher for low contrast targets (28 ms) than for high contrast targets (10 ms), in accordance with our hypothesis.

An ANOVA with between-subjects factors of present/absent ratio (1:1 vs. 3:1) and single/redundant ratio (1:1 vs. 2:1) as well as the within-subjects factor feature contrast (high, low), and intertrial dimension transition (same, different) revealed a significant main effect of present/absent ratio, as in the original analysis ($F(1,60)=4.9$, $p < .03$) and a significant main effect of intertrial transition, $F(1,60)=71.75$, $p < .0001$. Intertrial transition interacted significantly with present/absent ratio, $F(1,60)=6.98$, $p < .01$. DSCs were larger with 15 ms for the 1:1 than with 8 ms for the 3:1 present/absent condition.

5.4 Experiment 1

The question of Experiment 1 of the present study was twofold: (i) how are DSCs in a left/right localization task affected by a feature contrast manipulation, and (ii) is the previously reported dependence of DSCs on feature contrast due to a general slowing of responses, or due to shortening of the decision process. In addition to feature contrast, I manipulated the spatial congruency of the stimulus response-mapping. In a pro-localization task, observers responded with the hand corresponding to the side of the display a pop-out target was presented. Conversely, in an anti-localization task, participants had to press a button with the hand opposite of the target location (i.e. if the target was presented on the left half of the display, they responded with the right hand). If the modulation of DSCs were due to a general slowing of reaction times, both manipulations should lead to an increase of that effect. If, as I propose, the size of the effect depends on the duration of the decision process, which is affected by the perceptual manipulation of feature contrast, and not by the congruency of the stimulus response mappings, only the former should affect the DSCs. In order to verify that the feature contrast manipulation affects perceptual processing and the manipulation of task congruency influences post-selective processing, I fit the RDM

(Ratcliff, 1978, 1981) using a procedure recently proposed by Vandekerckhove and Tuerlinckx (2007b).

5.4.1 Method

Participants

12 observers (6 male) with a median age of 24 (ranging from 19 to 47), all right-handed, participated in Experiment 1. They received 8 Euro (ca. 11 USD) per hour for participation in the experiment.

Apparatus

Observers viewed the stimuli in front of a Sony Multiscan E250 17" monitor driven by a personal computer with Windows XP operating system. The personal computer was placed in a sound isolated room with black interior. There was dim background light in order to prevent reflections on the monitor. Viewing distance was about 62 cm and observers were instructed to maintain constant distance to the monitor. The screen refresh rate was 85 Hz, the screen resolution was set to 1024x768 pixels. Stimulus presentation and response recording was controlled with a purpose written C++ program.

Stimuli and Timing

The display consisted of grey, 45° left-tilted bars that were arranged on three concentric (imaginary) circles around a white fixation point (cf. Figure 5.4).

The (invisible) circles had a distance to the center of the screen of 4.5°, 8.5°, and 12.5° of visual angle. There were six, twelve, or 16 equidistant locations on the small, medium, and large circle, respectively, where stimuli were placed. Each bar was 0.6° wide and 2.7° high. Targets differed from distracters either in orientation




Figure 5.4. An exemplary stimulus display of Experiment 1. For demonstration there are presented two orientation and two luminance targets, of low and high saliency each. In the experiment always only one of these four possible targets was displayed.

(low feature contrast: 32° tilted left from vertical; high feature contrast: 45° tilted right from vertical) or in luminance. Luminance of distracters was $5\text{cd}/\text{m}^2$, of low feature-contrast luminance targets $11\text{cd}/\text{m}^2$, and of high-contrast luminance targets $51\text{cd}/\text{m}^2$. One target was randomly placed on one of six out of the twelve possible positions of the middle circle. If all positions of the middle circle were placed on a clock, possible target locations would be at two, three, and four, as well as eight, nine and ten o'clock.

Each trial began with the presentation of a fixation spot for 800-1000 ms, until all 34 stimuli appeared. One target, the dimension of which was orientation and the feature contrast of which was high in 50% of the cases each, was present in each trial. The stimuli remained on screen until the participant responded with a mouse click on either the left or right button with the index fingers of the left or right

hand, respectively. Trials were presented in 20 blocks of 72 trials each. The response requirements, pro- or anti-localization were constant for ten blocks and then changed. The order of tasks was counterbalanced across participants.

5.4.2 Results

Data analysis was done using R (R Development Core Team, 2006) except for fitting the RDM, which was done with the Diffusion Model Analysis Toolbox (DMAT, Vandekerckhove & Tuerlinckx, 2007b, 2007a), running on Matlab. Overall error rates were low: 2.8%. An analysis of variance (ANVOA) on error rates with the within-subjects factors target dimension (orientation, luminance), feature contrast (high, low), task (pro-localization, anti-localization), and intertrial transition (same dimension, different dimension) revealed the main effects of task, target dimension and feature contrast to be significant. Error rates were higher for the anti-localization than for the pro-localization task: 3.2% vs. 2.4%, $F(1,10)=21.2$, $p < .001$. There were more errors in the orientation than in the luminance dimension: 3.4% vs. 2.3%, $F(1,10)=31.4$, $p < .001$. Error rates were increased for low feature contrast (4.2%) compared to high feature contrast (1.4%), $F(1,10)=40.2$, $p < .001$. There was no further significant effect.

For reaction time analysis erroneous trials and trials after errors were excluded (less than 5.5% of the data). An ANOVA with the same factors as above revealed the main effects of task, contrast, and intertrial transition to be significant. Performance was faster in the pro-localization (377 ms) than in the anti-localization task (429 ms): $F(1,11)=28.77$, $p < .001$. Further, reaction times were faster for high contrast, than for low contrast targets: 389 ms vs. 438 ms, $F(1,11)=91.6$, $p < .001$. Intertrial transitions of same dimension were detected 6 ms faster than transitions of different dimension, $F(1,11)=4.9$, $p < .04$. There were two significant interactions, one between target dimension and feature contrast ($F(1,11)=7.8$, $p < .02$), and another, most importantly, between feature contrast and intertrial transition ($F(1,11)=10.7$,

$p < .007$). For low feature contrast, targets of both dimensions are virtually the same (438 ms), whereas luminance targets are detected faster (361 ms) than orientation targets (376 ms) under the high feature contrast condition. DSCs in the low feature contrast condition were 13 ms, whereas they were numerically non-existent in the high feature contrast condition. Additionally, there was no interaction between intertrial transitions and task.

We assumed that the manipulation of feature contrast affected perceptual processing, whereas the different stimulus-response mappings affected post-selective processing stages. Our previous analysis of mean reaction times revealed that both factors slow observed response latencies in an additive, non-interactive fashion. Still our assumption of the affected processing stage cannot be addressed by analysis of mean reaction times. In order to test this hypothesis I fitted the RDM using the procedures developed by Vandekerckhove and Tuerlinckx (2007b) using their Matlab DMAT (Vandekerckhove & Tuerlinckx, 2007a). The RDM has seven free parameters, which can be illustrated with help of Figure 5.2. The level of a is the boundary separation (i.e. the decision criterion). The higher a , the more accurate and slower responses are (speed/accuracy trade-off). Parameter z is the starting point, where accumulation of evidence begins at the start of each trial. z can be interpreted as a response bias. If $z = a/2$, there is no bias for any response. A value of z nearer to one of the decision boundaries, zero or a reflects a bias towards the nearer response, leading to faster but also more erroneous responses. Parameter ν represents the drift rate of the diffusion process (cf. the gray line in Figure 5.2). ν can be interpreted as the quality of the stimulus and determines the rate of accumulation of evidence for one of the two response alternatives. Observed response times in the RDM are assumed to be the sum of two components: the decision time, T_{dec} , and the non-decision time, T_{er} . T_{er} comprises time for stimulus encoding (in neuronal terms for visual stimuli the time necessary for the signal cascade started on the retina to reach cortical areas), and the time for response preparation and motor execution. In order to model variability of the decision process it is assumed that the drift rate ν varies

within one trial. Within-trial, ν is drawn from a normal distribution with a standard deviation of 0.1 (Vandekerckhove & Tuerlinckx, 2007b; Ratcliff, Zandt, & McKoon, 1999). Between trials ν is assumed to be drawn from a Gaussian distribution with standard deviation η , reflecting trial by trial variations in participants' attention or motivation. Finally, s_z is a parameter of the variability of starting point and s_t of the variability of the non-decision time, T_{er} , from trial to trial.

Across the different conditions I were interested in variations of the non-decision time, T_{er} , and the drift-rate, ν . All other parameters were assumed to be constant across conditions. The first question I wanted to answer by fitting the RDM was the following: does manipulation of feature contrast mainly affect stimulus quality (i.e. ν), and does manipulation of task (i.e. congruency of stimulus-response mapping) affect mainly non-decision processing (i.e. T_{er})? We fitted the RDM to four conditions: two levels of feature contrast crossed with two levels of response congruency. Specifically our question was, if manipulation of feature contrast affected ν and manipulation of task affected T_{er} . To this end I defined a series of three models, each an extension of the former. Model 1 was the so called null-model, where all parameters were assumed to be constant over the four conditions. Model 2 reflected the design of our experiment: I let drift rates vary according to feature contrast, with one parameter of ν for high and one for low feature contrast, and I let T_{er} vary according to task congruency. In Model 3 I tested deviations of the design, in letting both drift rates and non-decision times to vary across conditions. Model 1 has seven free parameters, Model 2 has nine, and Model 3 has eleven free parameters, with each model being nested in the next. Additionally I saw no reason to assume a bias for one over the other response and kept the starting value z fixed at $a/2$. We used the multinomial likelihood estimation based on quantiles (Vandekerckhove & Tuerlinckx, 2007b) with the 1, 2, 5, 10, 30, 50, 70 and 90% percentiles to test the fits of the models.

We collapsed data over dimensions and intertrial transitions for this analysis, because fitting the RDM depends on the estimation of reaction time distributions for both correct and erroneous responses, and overall error rates were low. All three models

were fitted for each participant. Exemplary, for one participant the statistics are presented in Table 5.1.

Table 5.1 Fit statistics from the model queue for one participant (Experiment 1).

Model	Λ	df	$\Delta\Lambda$	Δdf	p	AIC_C	BIC
1	5495.32	6				5507.38	5538.88
2	5226.48	8	268.85	2	0.000	5242.54	5284.54
3	5215.90	12	10.57	4	0.031	5240.13	5303.01

Model 2 outperformed the null-model significantly, indicating that feature contrast indeed affected drift rates and task congruency affected non-decision times. Still, Model 3 outperformed Model 2, indicating a deviation from our experimental design. Each participant showed this pattern of results. Therefore, I calculated the mean of all seven parameters of the RDM over all participants. In order to test the deviation from our design, I calculated two separate ANOVAs on the fitted ν and T_{er} parameters and the factors feature contrast (high, low) and task (congruent, incongruent stimulus-response mapping). One participant was left out of this analysis, because one parameter was badly fitted. For drift rates ν , the only significant effect was the main effect of feature contrast, $F(1,9)=34.3$, $p < .001$, with lower drift rates for low compared to high feature contrast (0.61 vs. 1.08, cf. Figure 5.5). For T_{er} , both the factors task ($F(1,9)=84.5$, $p < .001$) and feature contrast ($F(1,9)=9.5$, $p < .01$), but not their interaction was significant. In addition to T_{er} being slower for the incongruent than for the congruent task (292 ms vs. 261 ms), it was also slowed for low compared to high feature contrast targets (283 ms vs. 269 ms, cf. Figure 5.6). That means, although T_{er} also seems to be affected by feature contrast, feature contrast still had a significant effect on drift rates, and the task manipulation exclusively affected non-decision times. With respect to our main question these fitting results strengthen our hypothesis: pre-attentive effects, such as DSCs are affected by the duration of the decision process, not by the duration of non-decision processing time.

We found DSCs to be larger for low feature contrast, but unmodulated by the manipulation of stimulus-response congruency.

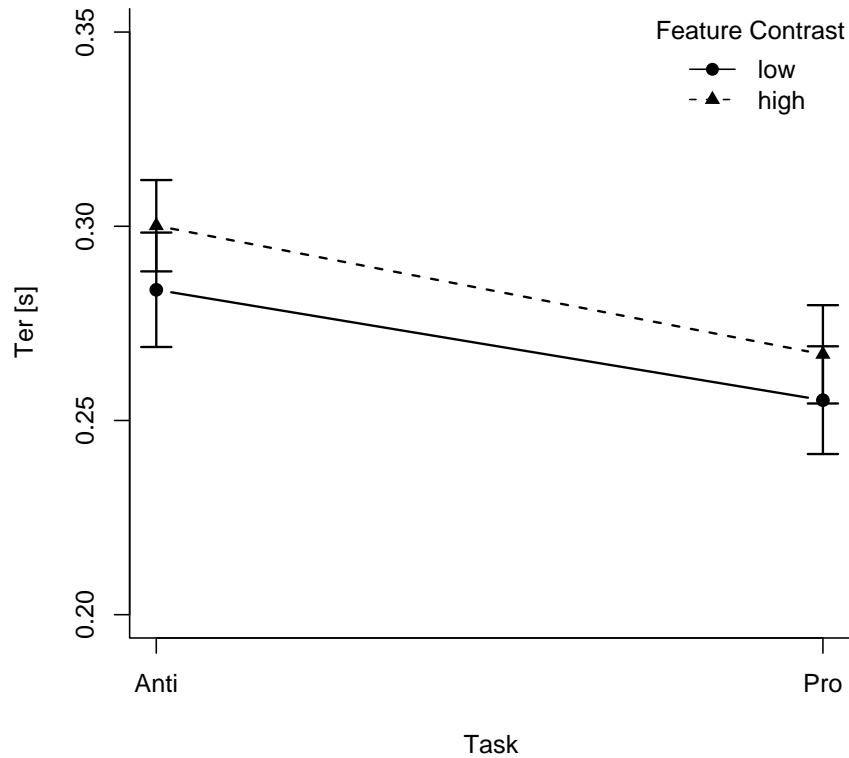


Figure 5.5. Fitted non-decision times T_{er} averaged across participants for both levels of task and of feature contrast.

Additionally to examining the question of what aspects of the processing architecture were affected by the task and feature contrast manipulation, I fitted the RDM to test effects of dimensional intertrial transitions on drift rates. For this analysis I collapsed data across dimensions and tasks, in order to have more than 300 observations per conditions. We examined the effect of the two factors contrast (high, low) and intertrial transition (same, different dimension) on drift rates. Therefore, again I tested a sequence of nested models. Model 1 was the null-model, in which drift rates were constant across all four conditions. In Model 2, I let drift rates vary with feature




Figure 5.6. Fitted drift rates ν averaged across participants for both levels of task and of feature contrast.

contrast, but not with intertrial transitions. Based on the analysis of task and feature contrast effects on drift rates, I expected Model 1 to outperform Model 2. In Model 3 and 4 I accounted for intertrial dimension changes in two ways: in Model 3 I tested, whether DSCs were equal for high and low feature contrast. That is, for the two different dimension conditions, I added one increment of drift rates for the different compared to the same dimension condition. It is also possible that the size of the DSC effect differs for high and low feature contrast targets. Therefore, in Model 4 I let drift rates vary freely across all four conditions. That way, Model 1 had six, Model 2 seven, Model 3 eight, and Model 4 nine degrees of freedom. Otherwise the procedure was equivalent to the previous fit, except that the non-decision time was

kept constant across all conditions.

The fitting procedure produced suspect results in one condition for one participant, which is consequently left out of the further analysis. In Table 5.2 I present the statistics of one participant. Model 2 outperformed the null-model, Model 1 significantly, as expected. Additionally allowing drift rates to vary with intertrial dimension changes significantly increased the fit of the model. Further, allowing for different DSCs for high and low feature contrast did not improve the fit any further. Another two participants showed the same pattern of results. For four other participants, Model 3 was not significantly better than Model 2, but Model 4 did increase the fit. That means for these four participants DSCs seemed to differ for high compared to low feature contrast. The recovered drift rates of Models 3 and 4 were subject to an ANOVA with the factors of feature contrast and intertrial transition. The main effects of feature contrast ($F(1,10)=51.9, p < .001, F(1,10)=51.3, p < .001$) and intertrial transition ($F(1,10)=8.3, p < .016, F(1,10)=5.79, p < .036$) were significant for Models 3 and 4, respectively. In addition, for Model 4 there was a significant interaction between feature contrast and intertrial transition, with DSCs being larger for low feature contrast than for high feature contrast ($F(1,10)=16, p < .002$). Post hoc tests revealed that DSCs in drift rates were significant only in the low feature condition, for Model 4.

Simulations by Vandekerckhove and Tuerlinckx (2007b) revealed that for numbers of

Table 5.2 Fit statistics from the model queue for one participant (Experiment 2).

Model	Λ	df	$\Delta\Lambda$	Δdf	p	AIC_C	BIC
1	5419.07	6				5431.13	5462.62
2	5227.86	7	191.21	1	0.000	5241.94	5278.67
3	5223.70	8	4.16	1	0.041	5239.80	5281.76
4	5223.17	9	0.53	1	0.47	5241.30	5288.50

observations similar to those in the present analysis, the probability of a 0.02 difference in drift rates to be discovered by their testing procedure was about 15%. Given

that the recovered differences in drift rate for dimensional intertrial transitions are exactly in that range, I have strong evidence for DSCs being reflected in changes of drift rates. However the distinction between Model 3 and Model 4 (i.e. the question, whether DSCs are equal for high and low feature contrast, or different) cannot be conclusively answered by this analysis.

5.4.3 Discussion

In Experiment 1, reaction times were manipulated in a perceptual way by different levels of similarity between targets and distracters, and in a response related way, by a more or less congruent stimulus response mapping. Both manipulations had a substantial effect on performance (about 70 ms for feature contrast and 50 ms for task, respectively), which is also reflected in error rates. Intertrial transitions of dimension lead to increased DSCs under the low feature contrast condition, whereas DSCs were numerically non-existent under high feature contrast conditions. At the same time, DSCs did not interact with the task manipulation. Additionally, fitting the RDM verified that the feature contrast had a perceptual effect (i.e. improved stimulus quality), whereas the task manipulation affected exclusively non-decision, probably post-selective components of processing. Further, I found effects of dimensional intertrial transitions on drift rates. These findings are in accordance with our proposal that the duration of the decision process influences the size of pre-attentive effects, such as DSCs, in contrast to non-decision components of reaction times, which I assume not to have such an effect.

Similar effects for target changes have been described by Meeter and Olivers (2006), although in a slightly different paradigm. Instead of the target being possibly defined in one of two dimensions, in the experiments of Meeter and Olivers (2006) the target could differ from distracters in one of two possible colors (i.e. features). But in contrast to experiments reported previously in the present study, target and distracter could either change roles randomly from trial to trial (Experiments 1 and 2), or were

placed among gray distracters (Experiment 3). That is, only changes of features, not dimensions have been investigated. The task of participants was to determine a feature of the target, e.g. in which direction an arrowhead is pointing (cf. Bravo & Nakayama, 1992; Maljkovic & Nakayama, 1996). Meeter and Olivers (2006) manipulated display density (i.e. distance between the display elements) in varying set size, they introduced a salient, task irrelevant singleton distracter, and varied the size of elements. Display density and an irrelevant distracter clearly influence the difficulty of the decision what item should be selected first. Accordingly, Meeter and Olivers (2006) found reaction times to be slowed by 90 ms both for sparse displays and for displays that contained an irrelevant singleton distracter. Simultaneously, this slowing of reaction times went along with an increased reaction time costs of changing the target features: change costs were 55 ms larger in case of the set size manipulation, and 40 ms larger for the singleton distracter condition. So in contrast to dimension changes, which have been topic of examination until now, these experiments either incorporate feature changes, either in combination with a swap of target and distracter features, or with constant (neutral) distracter features. Huang and Pashler (2005) argued that in reaction time measures both pre-attentive and post-selective processing stages affect performance, whereas with accuracy measures under brief viewing conditions post-selective processing is eliminated. They found no significant intertrial effects for swaps of target and distracter features under brief viewing conditions and interpreted this finding in that intertrial effects in a target/distracter swap situation may be largely due to post-selective processing. Given that in cross-dimensional situations, in which the target can be defined by one of two dimensions, Zehetleitner (2007b) have demonstrated the existence of dimension switch costs in accuracy under brief viewing conditions, arguing for the pre-attentive nature of dimensional switch costs, it is possible that DSCs and target-distracter swap effects stem from different weighting systems. Thus, taking into account the fact that the tasks employed by Meeter and Olivers (2006) either used a paradigm that possibly produced post-selective (their Experiments 1 and 2), or feature based effects (their

Experiment 3), in Experiment 2 I examine the effect of target saliency and congruency of stimulus-response mapping in a compound task with dimensional uncertainty.

5.5 Experiment 2

In Experiment 2 I manipulated feature contrast of targets and congruency of stimulus-response mapping in a compound task. At the aim of this experiment was the question, whether reduction of the decision time necessary to select a target consequently reduces the size of DSCs, as predicted by our hypothesis.

5.5.1 Method

Participants

12 observers (six male) with a median age of 25 (ranging from 19 to 47), three left-handed, participated in Experiment 2. They received 8 Euro (ca. 11 USD) per hour for participation in the experiment.

Apparatus

The apparatus was the same as in Experiment 1.

Stimuli and Timing

An exemplary stimulus display is presented in Figure 5.7. Stimuli were of same dimensions and arrangement as in Experiment 1 with the following differences: distracters were yellow instead of gray (CIE xyY coordinates 0.475, 0.475, 25) and instead of luminance targets could be color defined. Low feature contrast targets were reddish orange (0.562, 0.382, 25) and high feature contrast targets red (0.591, 0.351, 25). Additionally I manipulated orientation of distracter items randomly from block to block. In addition to distracters being tilted 45° left from vertical, they could be

tilted 45° right from vertical. In such blocks orientation targets were either 45° tilted to the left (high feature contrast), or tilted 32° to the right. In both types of blocks high orientation contrast was 90° and low feature contrast 13°.

Timing also differed from Experiment 1. In Experiment 2 at the beginning of each trial a fixation spot was visible for 800-1000 ms, which disappeared with the onset of all 34 stimuli. The stimuli were present until participants responded, followed by a blank screen of 500 ms.

Trials were presented in 14 blocks of 90 trials each. The task was to respond to the




Figure 5.7. An exemplary stimulus display of Experiment 2. Presented is a high contrast orientation target. The task of observers was to respond to the 'E' or mirror-'E' characters in the target.

'E' or mirror-'E' characters inside the target (Deubel & Schneider, 1996). Stimulus-response mapping could be congruent or incongruent. In the congruent condition, participants were instructed to respond with the hand corresponding to the side at which the character is open (e.g. the right hand for the character 'E'). For the incongruent mapping response assignment was reversed (i.e. the correct response to the

character 'E' was the left hand). The task condition was constant for seven blocks and then changed. The first task condition was counterbalanced across observers.

5.5.2 Results

Overall error rates were low (4.2%). An ANOVA on error rates with the factors dimension (color, orientation), feature contrast (high, low), task (congruent, incongruent), and intertrial transition (same, different dimension) revealed only the effect of contrast to be significant, $F(1,8)=7.5$, $p < .02$ (3.4% vs. 5.5% for high and low feature contrast, respectively).

For reaction time analysis erroneous trials and trials after errors were excluded from analysis (less than 9% of the data). An ANOVA on reaction times with the same factors as for error rates revealed the main effects of dimension and contrast to be significant and the main effect of task to be marginally significant. Reaction times were faster for color than for orientation targets (719 ms vs. 770 ms), $F(1,8)=16.9$, $p < .003$. Performance also was faster for high contrast than for low contrast targets (718 ms vs. 771 ms), $F(1,8)=81.9$, $p < .001$. The congruent stimulus response mapping in tendency was responded to faster (736 ms) than the incongruent mapping (753 ms), $F(1,8)=3.5$, $p < .09$. The contrast manipulation had a larger effect for orientation targets (90 ms) than for color targets (20 ms), $F(1,8)=22.7$, $p < .001$. Most importantly, intertrial effects interacted with feature contrast: the DSC was 0.1 ms for high feature contrast and 14 ms for low feature contrast targets. Intertrial effects did not interact with task ($F(1,8)$, $p < 0.8$). There were no further significant effects. The RDM could not be used as an analytic tool in the analysis of Experiment 2, because it is confined to binary decisions. Of course the response of participants in the compound task was based on a binary decision, in which way a character 'E' was turned, but the decision, what target should be selected, was non-binary. Our proposal of increased decision times increasing the size of pre-attentive effects refers

to the decision what stimulus to select, not on the binary response defining decision of the compound task.

5.5.3 Discussion

In Experiment 2 the manipulation of feature contrast had a significant effect on DSCs. DSCs were low for high feature contrast, and high for low feature contrast targets. This is in line with our hypothesis that as longer duration of the decision process, which is based on the saliency master map the size of pre-attentive effects, such as the DSC, is increased. In a compound task there are two decision processes: (i) selection of a stimulus and (ii) discriminating the response defining feature. The former decision is a decision with more than two alternatives, because there is more than one possible stimulus location, whereas the latter decision is binary. According to our hypothesis, only manipulation of the first, perceptual decision process influences the size of observable DSCs, while manipulation of difficulty of the later, post-selective decision leaves these effects untouched.

5.6 General Discussion

In the present paper I investigated pre-attentive visual processing from a decision perspective and proposed that the size of observable performance differences, which are caused by pre-attentive modulations of sensory processing, depend on the duration of the relevant decision processes (i.e the observed differences are larger, the slower the decision process is). Relevant decisions in that sense are decisions about presence or absence of a target (detection) or the coarse location of a target (e.g. left/right localization). A further relevant decision is where to direct the gaze or covert attention to next.³ The former binary decisions can mathematically be modeled by sequential

³Of course it is not possible to totally dissociate detection and localization or movement of an effector from deployment of attention (Deubel & Schneider, 1996; Baldauf, Wolf, & Deubel, 2006; Theeuwes, 1992). But in the present context the different task make use of activation on the saliency map for different purposes: either to detect the presence of a target, to roughly localize it, or to select it for further analysis.

sampling models, such as the Wiener diffusion model (Ratcliff, 1978) or the leaky competing accumulator model (Usher & McClelland, 2001) for binary decisions with multiple alternatives. Based on the geometry of such decisions (cf. Figure 5.3) it is apparent that differences in drift rates, with which evidence for an alternative accumulates (e.g. due to modulations of pre-attentive processing), have a stronger effect, the longer the decision process takes. The duration of the decision process can be manipulated in two fundamentally different ways: (i) by manipulating the quality of the stimulus and (ii) by manipulating the observer's response criteria. The quality of the stimulus, for instance, can be modulated by different levels of feature contrast, and the response criteria can be modulated by instruction (speed/accuracy trade-off) or by frequency of a specific response (thereby lowering the criterion for the frequent response).

In the present study I applied this proposal to visual pop-out search. Most models of visual pop-out search assume that processing depends on a topographical representation of the visual display, which signals distinctiveness of each location (i.e. a saliency map, cf. Koch & Ullman, 1985; Wolfe, 1994; Müller et al., 1995; Itti & Koch, 2000). Activity on this master map can then be used by the visual system to solve different tasks: detection (presence or absence of a target), coarse localization (target position in the left or the right half of the display), and selection of a location for deployment of attention or direction of an effector to that location. In that sense compound tasks, in which a pop-out target has to be selected in order to discriminate a response defining feature of or within the target (c.f. Bravo & Nakayama, 1992) are comprised of two decision processes: the first decision determines what stimulus to select for further analysis, and a second decision regarding the response defining feature. For the purpose of the present study the first decision of selection is most relevant, because this decision what item to select is based on activity of the salience map, whereas the second decision about the response defining feature requires featural analysis.

Regarding visual pop-out search, there are several effects that are assumed to modu-

late pre-attentive processing: (i) performance gains of targets defined redundantly in two dimensions over targets defined in only one dimension (RSE), (ii) performance decreases after change of the target defining dimension in a sequence of trials (DSCs), and (iii) performance benefits of intentional preparation for the dimension of an upcoming target. All three effects are assumed to modulate build-up of activity on the saliency map (i.e. modulate saliency), thus affecting the ease of decisions based on that map, as found in electrophysiological recordings (Bichot & Schall, 1999, 2002). The dimension weighting account (Müller et al., 1995; Found & Müller, 1996) in that context is an explicit model of how an implicit memory system (which is penetrable to top-down control) is implemented in dimensional weights, which affect processing of feature contrast signals before pooling into the salience map.

We reported evidence that performance differences (specifically the RSE, DSCs, and benefits from dimensional cues) differ in size in correspondence to the duration of the relevant decision process. This finding holds for different types of tasks (detection, localization, and compound tasks), as well as for different manipulations of the duration of the decision process (feature contrast, frequency of responses). Additionally, Experiment 1 of the present study provides evidence that in fact the duration of the decision process matters - not the overall latencies: although manipulating the congruency of the stimulus-response mapping increased latencies, similar to the manipulation of feature contrast, it did not modulate DSCs. However, DSCs were larger for targets defined by lower feature contrast. This view was substantiated by fitting the RDM, which revealed feature contrast to affect the drift rates of the decision process (and thereby its duration), whereas the congruency of the stimulus-response mapping affected processing, which was not related to the decision.

5.6.1 Relation to previous studies

We will relate the present findings to three previously proposed accounts of early visual processing: the ambiguity account (Olivers & Meeter, 2006; Meeter & Olivers, 2006), the post-selective dimension-action system (Cohen & Feintuch, 2002), and recent revivals of Feature Integration Theory (FIT: Treisman & Gelade, 1980; Treisman, 1988), who propose a dual-route processing architecture for detection and spatial (such as localization or compound) tasks.

The ambiguity account

The ambiguity account proposed in two recent studies (Meeter & Olivers, 2006; Olivers & Meeter, 2006) is especially relevant for the present findings. They described an effect similar to what I propose in a more general way for all types of pre-attentive modulations. The authors propose that ambiguity is the key defining feature of when intertrial priming effects (comparable to our DSCs) do or do not appear. They conclude their findings in an ambiguity account, which states that intertrial priming becomes functional, and therefore measurable only under circumstances of ambiguity. They understand ambiguity as the presence of uncertainty, conflict or competition at any level between stimulus and response (Olivers & Meeter, 2006, p. 3). In that way, the ambiguity account has both a broader and a narrower scope than the present study. It is broader in the sense that it takes into account ambiguity not only on the level of visual selection, but also at the level of response selection and execution. We will discuss later what our proposal can contribute to these aspects of processing. Still, the ambiguity account is simultaneously narrower in scope, as it is confined to intertrial priming effects, whereas in the present study I refer to any modulations of sensory processes that can affect a decision, especially pre-attentive modulations of a visual saliency map, based on which decisions are made. That way, our proposal is not confined to intertrial priming effects, as the ambiguity account, but specifically extends to the RSE and benefits from dimensional cues.

The key construct of Meeter and Olivers (2006) and Olivers and Meeter (2006) is ambiguity. They explicitly delineate ambiguity from saliency. In their sense, different levels of saliency do not necessarily mean different levels of ambiguity. They argue ambiguity to be related to the number of items that compete for visual selection (or the number of response alternatives that compete for response execution). In the domain of visual selection, for instance the number of distracters (i.e. a set size manipulation) in that sense affects ambiguity, because it affects the number of items that compete for visual selection. Additionally, introducing an irrelevant singleton also increases the number of items that compete for visual selection, thus increasing ambiguity. They found increased intertrial priming effects in situations where the target could be defined in one of two possible dimensions (Olivers & Meeter, 2006), or in situations where the target and distracter features could be swapped (Meeter & Olivers, 2006). Simultaneously, in addition to increased intertrial priming effects, both manipulations also led to increased reaction times. In that sense, our proposal would assume that the set size manipulation as well as the introduction of a singleton distracter increase the duration of the decision process, thus leading to larger intertrial effects. In both Olivers and Meeter (2006) and Meeter and Olivers (2006) the manipulation of set size was confounded with a manipulation of display density. All items were equidistant to the central fixation point and equidistant to each other. Therefore, in the 12-item condition, a target was in close vicinity to distracters, whereas in the 3-item condition, a target was farther away from the two distracters. Thus it is assumable that the set size manipulation in both studies was in fact a manipulation of salience, in line with the findings of Nothdurft (2000), who reported saliency of pop-out targets to decrease with decreasing display density. So it may well be that saliency - and not the number of items - is the key determining feature in modulating the size of priming effects. Regarding the second manipulation of Meeter and Olivers (2006), distraction by an irrelevant singleton is widely assumed to be a pre-attentive effect (Theeuwes, 1991, 1992; Müller, Krummenacher, & Geyer, 2007). That way it affects the time necessary to select the relevant target and, ac-

ording to our proposal, lead to increased observable priming effects. Specifically, as the duration of decision processes depends on drift rates and response criteria, a singleton distracter may affect the former, rather than the latter. The selection decision of a compound task can be understood as a leaky, competing accumulator model (Usher & McClelland, 2001), in which there exists one accumulator for each possible target location. Importantly, in that model, accumulators are not independent, but mutually inhibitive (i.e. increase of evidence in one accumulator reduces evidence in all other accumulators). That way, a task irrelevant singleton distracter may lead to an increase evidence larger than (non-singleton) distracters. Consequently, due to the mutual inhibitory connections of the accumulators, the drift rate with which evidence accumulates at the target location will be slower in the presence compared to absence of a singleton distracter. So far, both the ambiguity account and our current proposal lead to similar predictions. However, Meeter and Olivers (2006, p.206) explicitly exclude saliency being equivalent to their concept of ambiguity:

We thus propose that it is not the absolute salience of the target that determines whether intertrial priming occurs, but whether it is unambiguously the most salient element in the display.

Hence, in their view only the ratio of targets to distracters (modulated by set size or by introduction of an irrelevant singleton) affects ambiguity. This proposal, however, is in contrast to several findings reported in the present study. We reported that direct manipulations of saliency via similarity between targets and distracters did indeed modulate the size of intertrial priming effects in detection (Zehetleitner, 2007a), localization, and compound tasks (Experiments 1 and 2 of the present study). Further, this effect is not restricted to intertrial priming, DSCs in case of the studies reported here, but also affects the size of the RSE or the benefit from dimensional cues. Furthermore, in addition to manipulation of feature contrast, the same effect arose from a manipulation of the response criterion (Zehetleitner, 2007a). In terms of the ambiguity account, (visual) ambiguity is the same in both cases, as the number of

distracters compared to the number of targets stays the same. Therefore, this finding does not find a direct explanation in the ambiguity account.

In summary, although the findings of Meeter and Olivers (2006) can be explained in terms of our current proposal, the modulation of intertrial priming effects by the duration of the relevant decision process (via salience or response criterion) reported here are in contrast to the predictions of the ambiguity account.

Further, although both Meeter and Olivers (2006) and Olivers and Meeter (2006) report intertrial effects to be modulated by various manipulations, fail to give an explanation of the underlying processes, which are responsible for these modulations. Their explanation is a functional one: they speculate that in cases of less competition, when there is less need for ambiguity resolution, it would not be beneficial for the system to use information about the target definition from the last trial. However, when there is ambiguity to be resolved, priming may be beneficial. In their sense the system has adaptively evolved to applying information about the previous trial only when it may be helpful and increase performance.

On the other hand the present proposal provides an explanation of the underlying mechanism, which modulates pre-attentive effects. The modulation of intertrial priming effects by the duration of decision processes is a feature that emerges from the system structure: the longer a decision process takes, the larger the difference in decision latencies for one and the same difference in drift rates becomes.

Additionally, the conception of detection and compound tasks differ between the present proposal and the ambiguity account. Olivers and Meeter (2006) describe detection of absence or presence of a target as a signal detection task (Green & Swets, 1964): in order to discriminate noise from an actual target, the observers have to set a response criterion. In contrast, they argue that compound tasks are fundamentally different, in such that in a compound task there is no uncertainty as in a decision task, because a target is present in all cases. Additionally, in order to discriminate the response defining feature of a target, all ambiguity about the target has necessarily already been resolved at the time, when the response is generated. Therefore,

because according to the ambiguity account, priming effects are affected by uncertainty/ambiguity, intertrial priming effects can be observed in detection, but not in compound tasks. The key feature of detection tasks in their view is that evidence for or against presence of a target is subject to noise. Signal detection theory assumes that the presence/absence decision is based on one sample of evidence (which is subject to noise), whereas sequential sampling models assume that evidence for absence or presence is continuously sampled from noisy signals. In that sense both conceptions of the detection task are equivalent. However, in a compound task, from a decision perspective, there also is uncertainty - not about presence of the target, but about its location. Again sequential sampling models can be used to describe this selection decision: for all possible target locations there is an accumulator, which samples evidence from the saliency map that there is a target at this specific location. Additionally these accumulators are coupled with mutually inhibitory connections (i.e. evidence for one accumulator is evidence against others, Usher & McClelland, 2001). As soon as activity of one accumulator exceeds a criterion, attention (and perhaps the gaze) is directed to that location. That way, there is no fundamental difference between the presence/absence decision in a detection or the selection decision in a compound task. In both cases decisions are based on noisy signals (i.e. there is uncertainty about the correct decision). The fact that in both compound and detection tasks manipulations of feature contrast lead to modulations of DSCs provide further evidence that both tasks have common grounds in terms of being based on noisy signals from the salience map.

Post-selective accounts

The findings of the present study are additionally relevant for theories that assume DSCs or the RSE to be post-selective effects. The dimension-action model (Cohen & Feintuch, 2002; Cohen & Shoup, 1997) for instance assumes that there exist di-

mension specific response selection modules in addition to visual analyzers, which are also organized in a dimension specific fashion. That is, they assume that visual processing initially is organized in dimensions (Itti & Koch, 2000; Wolfe & Horowitz, 2004; Müller et al., 1995) and further propose the existence of dimension specific modules, which process selection of responses. Just like the visual processing, these response selection modules are initially activated in parallel across the visual display. Attention then serves as a gate that determines which response selection module can spread its activity further to executive processing stages and thus can trigger a response. In their view, dimensional intertrial effects (DSCs) and benefit of redundant over single targets (RSE) occur on the level of the dimension specific response selection modules. That means, both DSCs and the RSE are post-selective effects, as they depend on deployment of attention. This view has been challenged by findings of Krummenacher et al. (2002), who reported a significant RSE and violations of the race model inequality (RMI, an indicator of co-active/integrative compared to parallel processing: Miller, 1982) even if attention was deployed to a specific location and the target appeared at a different location. Additionally, in the present study I report evidence that manipulation of feature contrast that is manipulation of pre-attentive processing has an effect on the size of DSCs and the RSE. A model, in which DSCs and the RSE arise from post-selective processing stages, would predict both effects to be untouched by modulations of pre-attentive processing. Hence, in the light of evidence from Krummenacher et al. (2002) and the present study, it seems justified to favor models that understand DSCs and the RSE as pre-attentive effects (e.g. the dimension weighting account, Müller et al., 1995) over models that understand them as post-selective effects (e.g. the dimension action model of Cohen & Feintuch, 2002).

The dual-route proposal

In addition to post-selective accounts of effects such as DSCs or the RSE, Feature Integration Theory (FIT: Treisman & Gelade, 1980; Treisman, 1988) has been recently revived in two studies (Chan & Hayward, 2007; Mortier et al., 2007). Both studies compare performance for different tasks (detection, localization, and compound tasks). While Chan and Hayward (2007) focus on the effect of task on DSCs, Mortier et al. (2007) evaluate the effect of task on benefits from dimensional cues. Chan and Hayward (2007) compared DSCs for detection, localization, and compound tasks, and found them to be substantially reduced in compound and localization, when compared to the detection task. Additionally they found singleton distracters to interfere with performance in localization and compound, but not in the detection task. In order to explain these discrepant findings, they revive FIT, which assumes, in addition to a salience map (similar to Wolfe, 1994; Müller et al., 1995), there are dimensional modules, which only signal target presence, without providing any spatial information. So in contrast to signals on the salience map, activity in dimensional modules is dimension specific and spatially pooled (although feature unspecific). According to their argument, for detection tasks processing relies on the dimensional modules, because spatial information is irrelevant for detection. In contrast, for localization and compound tasks, spatial information is necessary, and thus processing relies on the salience map for these tasks. In order to account for the differences in DSCs, Chan and Hayward (2007) propose that dimensional weighting mechanisms only affect the spatially unspecific dimensional modules, but not the salience map. Therefore, DSCs are prominent in detection, but not in localization or compound tasks. However, Krummenacher et al. (2002) showed that in a detection task, integration of redundant signals from two dimensions is spatially specific, ruling out a-spatial processing in detection tasks. Further, Töllner, Gramann, Müller, Kiss, and Eimer (2007) reported modulations of the N2pc, an electrophysiological marker for allocation of attention, with both shorter latencies and higher amplitudes for repeti-

tion compared to a change of the target defining dimension. At the same time, the N2pc was invariant to response changes over a sequence of two trials. In order to account for both findings, Chan and Hayward (2007) conceded that there may be some (weak) dimensional weighting mechanism on the saliency map. Further, they proposed two possible processing routes for the detection task. For trials with a target defined in only one dimension, processing occurs via the spatially unspecific dimensional modules, whereas redundant targets are processed via the salience map. Mortier et al. (2007) on the other hand compared benefits from dimensional cues for detection and localization tasks. They replicated such benefits in the detection task (Müller et al., 2003; Theeuwes et al., 2006) but found them to be substantially reduced (and numerically non-existent) in localization tasks, irrelevant of the effector used to submit the response (eyes or hands). They hold this finding as evidence for a processing architecture similar to the revised version of FIT by Chan and Hayward (2007): weighting mechanisms modulate processing of detection tasks via spatially unspecific dimensional modules, whereas localization tasks (which require spatial information) are processed via the salience map, for which no weighting mechanisms exist. In summary, both studies argue for qualitatively different processing in detection compared to compound and localization tasks. However, in the present study I report all three tasks to show a similar behavior, when manipulating the duration of the relevant decision process (the absent/present decision in case of the detection, the left/right decision in case of the location, and the selection decision in case of the compound task): the longer this decision takes, the larger latency differences become that are induced by differences in activity on the salience map. Hence, from a decision perspective differences in the size of DSCs, the RSE or benefits from dimensional cues are not due to qualitative differences in processing, but due to quantitative differences. That is, these effects are not absent in localization or compound tasks, but they are latent, and become measurable, the slower the decision process becomes.

5.6.2 Theoretical Implications

In the present paper, I have examined pre-attentive visual processing from a decision perspective. We have analyzed, how different tasks, which require different decisions, may operate on a visual salience map, and how modulations of activity on the salience map affect observable response latencies with regard to different durations of the relevant decision process. For that purpose I have applied a very general property of such decision processes, according to which the same difference in drift rates leads to different decision latencies depending on the duration of the decision. The tasks described in the present paper comprised binary decisions about presence/absence (i.e. detection) or rough location (left/right, i.e. localization) of a target, or selection decisions in terms of compound tasks. Compound tasks in that sense are comprised of two decisions: first the decision what stimulus to select for attentional deployment, and the second decision about the response defining feature. Therefore, this decision perspective serves as a tool to determine whether a performance difference between two conditions arises from modulations of pre-attentive or post-selective processing. Only if modulations of pre-attentive processing are responsible for the observed effect, this effect should be modulated with manipulation of decision processes based directly on sensory input, such as a salience map. If on the other hand the effect under question arises from modulations of post-selective processing stages, it should be invariant to manipulations of duration of the primary decision process (e.g. via salience or response criteria), which relies on initial sensory processing.

This decision perspective makes several predictions. The same principle should apply for different weighting systems through out the processing hierarchy that is on the sensory level, on the level of rule-sets, or on the effector level. Recently weighting systems similar to the dimension weighting account (Müller et al., 1995) have been proposed for sensory modalities (Töllner, Gramann, Müller, & Eimer, 2007b) response preparation/programming (Töllner, Gramann, Müller, Kiss, & Eimer, 2007),

and stimulus-response mappings (Rangelov, Zehetleitner, & Müller, 2007). From a decision perspective, the prediction is that performance costs when switching between sensory modalities become larger, when the relevant decision process is prolonged. If for decisions, what effector or what rule set to use the recently proposed weighting mechanisms affect the rates with which evidence for one or the other alternative accumulates, than prolonging of the decision process would lead to larger observed performance differences. In summary, the decision perspective can be applied to early sensory processing, selection of task- or rule-sets, or preparation of responses.

5.7 Conclusion

In the present study I investigated precisely formulated theories of early visual processing (Itti & Koch, 2000; Müller et al., 1995) from the perspective of decision processes (Ratcliff, 1978), with the goal to close a gap between theories of visual processing and theories of decision making. In the present context, decisions could be binary about the presence or coarse location of a target, or could have multiple outcomes, in case of the decision what location to select next for deployment of covert and/or overt attention. Specifically I investigated a prediction that derives from the geometry of decisions modeled with sequential sampling processes (cf. Figure 5.3): the size of latency differences due to modulations of sensory input (reflected by different drift rates), depends on the duration of the relevant decision process. The findings support theories, which assume the benefit of redundant over single pop-out targets (RSE), the cost of changing the target defining dimension (DSCs), and the benefit from dimensional cues to pre-attentively modulate activity on a master salience map (especially the dimension weighting account Müller et al., 1995). Although applied to early visual processing, the decision perspective is fairly general and makes specific predictions about recently proposed weighting systems for sensory modalities, task/rule-sets, and response preparation. Finally I demonstrated the use-

fulness of fitting the RDM (Ratcliff, 1978; Vandekerckhove & Tuerlinckx, 2007b) for visual search tasks.

References

- Baldauf, D., Wolf, M., & Deubel, H. (2006, Dec). Deployment of visual attention before sequences of goal-directed hand movements. *Vision Research*, *46*(26), 4355–4374. Available from <http://dx.doi.org/10.1016/j.visres.2006.08.021>
- Bichot, N. P., Rossi, A. F., & Desimone, R. (2005, Apr). Parallel and serial neural mechanisms for visual search in macaque area v4. *Science*, *308*(5721), 529–534. Available from <http://dx.doi.org/10.1126/science.1109676>
- Bichot, N. P., & Schall, J. D. (1999, Jun). Effects of similarity and history on neural mechanisms of visual selection. *Nature Neuroscience*, *2*(6), 549–554. Available from <http://dx.doi.org/10.1038/9205>
- Bichot, N. P., & Schall, J. D. (2002, Jun). Priming in macaque frontal cortex during popout visual search: feature-based facilitation and location-based inhibition of return. *Journal of Neuroscience*, *22*(11), 4675–4685. Available from <http://dx.doi.org/20026410>
- Bravo, M. J., & Nakayama, K. (1992). The role of attention in different visual-search tasks. *Perception and Psychophysics*, *51*, 421–432.
- Bruce, C. J., Goldberg, M. E., Bushnell, M. C., & Stanton, G. B. (1985, Sep). Primate frontal eye fields. II. physiological and anatomical correlates of electrically evoked eye movements. *Journal of Neurophysiology*, *54*(3), 714–734.
- Bundesen, C., Habekost, T., & Kyllingsbaek, S. (2005, Apr). A neural theory of visual attention: bridging cognition and neurophysiology. *Psychological Review*, *112*(2), 291–328. Available from <http://dx.doi.org/10.1037/0033-295X.112.2.291>
- Chan, L. K. H., & Hayward, W. G. (2007). Feature integration theory revisited: Dissociating feature detection and attentional guidance in visual search. (Submitted manuscript.)
- Cohen, A., & Feintuch, U. (2002). The dimensional-action system: A distinct visual system. In W. Prinz & B. N. Hommel (Eds.), *Attention and performance: XIX. common mechanisms in perception and action* (p. 587–608). Oxford, UK: Oxford University Press.
- Cohen, A., & Magen, H. (1999). Intra- and cross-dimensional visual search for single feature targets. *Perception & Psychophysics*, *61*, 291–307.
- Cohen, A., & Shoup, R. (1997, Mar). Perceptual dimensional constraints in response selection processes. *Cognitive Psychology*, *32*(2), 128–181. Available from <http://dx.doi.org/10.1006/cogp.1997.0648>

- Cohen, A., & Shoup, R. (2000, Feb). Response selection processes for conjunctive targets. *Journal of Experimental Psychology: Human Perception and Performance*, *26*(1), 391–411.
- Colonius, H., & Diederich, A. (2006). The race model inequality: Interpreting a geometric measure of the amount of violation. *Psychological Review*, *113*(1), 148–154.
- Connor, C. E., Preddie, D. C., Gallant, J. L., & Essen, D. C. V. (1997, May). Spatial attention effects in macaque area v4. *Journal of Neuroscience*, *17*(9), 3201–3214.
- Corballis, M. C. (2002). Hemispheric interactions in simple reaction time. *Neuropsychologia*, *40*(4), 423–434.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006, May). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, *10*(5), 204–211. Available from <http://dx.doi.org/10.1016/j.tics.2006.03.007>
- Deubel, H., & Schneider, W. X. (1996, Jun). Saccade target selection and object recognition: evidence for a common attentional mechanism. *Vision Research*, *36*(12), 1827–1837.
- Diederich, A., & Colonius, H. (2004, Nov). Bimodal and trimodal multisensory enhancement: effects of stimulus onset and intensity on reaction time. *Perception & Psychophysics*, *66*(8), 1388–1404.
- Duncan, J. (1985). Visual search and visual attention. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance 5: Attention and neuropsychology* (p. 85 – 106). Hillsdale, NJ : Lawrence Erlbaum Associates, Inc.
- Duncan, J., & Humphreys, G. W. (1989, Jul). Visual search and stimulus similarity. *Psychological Review*, *96*(3), 433–458.
- Fecteau, J. H., & Munoz, D. P. (2006, Aug). Saliency, relevance, and firing: a priority map for target selection. *Trends in Cognitive Sciences*, *10*(8), 382–390. Available from <http://dx.doi.org/10.1016/j.tics.2006.06.011>
- Feintuch, U., & Cohen, A. (2002, Jul). Visual attention and coactivation of response decisions for features from different dimensions. *Psychological Science*, *13*(4), 361–369.
- Found, A., & Müller, H. J. (1996). Searching for unknown feature targets on more than one dimension: investigating a "dimension-weighting" account. *Perception & Psychophysics*, *58*(1), 88–101.
- Friedman, H. S., Zhou, H., & Heydt, R. von der. (2003, Apr). The coding of uniform colour figures in monkey visual cortex. *Journal of Physiology*, *548*(Pt 2), 593–613. Available from <http://dx.doi.org/10.1113/jphysiol.2002.033555>
- Gegenfurtner, K. R., Kiper, D. C., & Fenstemaker, S. B. (1996). Processing of color, form, and motion in macaque area v2. *Visual Neuroscience*, *13*(1), 161–172.

- Geyer, T., Müller, H. J., & Krummenacher, J. (2007, Aug). Cross-trial priming of element positions in visual pop-out search is dependent on stimulus arrangement. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(4), 788–797. Available from <http://dx.doi.org/10.1037/0096-1523.33.4.788>
- Giray, M., & Ulrich, R. (1993, Dec). Motor coactivation revealed by response force in divided and focused attention. *Journal of Experimental Psychology: Human Perception and Performance*, *19*(6), 1278–1291.
- Goldberg, M. E., Bisley, J. W., Powell, K. D., & Gottlieb, J. (2006). Saccades, salience and attention: the role of the lateral intraparietal area in visual behavior. *Progress In Brain Research*, *155*, 157–175. Available from [http://dx.doi.org/10.1016/S0079-6123\(06\)55010-1](http://dx.doi.org/10.1016/S0079-6123(06)55010-1)
- Gomez, P., Ratcliff, R., & Perea, M. (2007, Aug). A model of the go/no-go task. *Journal of Experimental Psychology: General*, *136*(3), 389–413. Available from <http://dx.doi.org/10.1037/0096-3445.136.3.389>
- Gottlieb, J. (2002, Apr). Parietal mechanisms of target representation. *Current Opinion in Neurobiology*, *12*(2), 134–140.
- Green, D. M., & Swets, J. A. (1964). *Signal detection theory and psychophysics*. New York: John Wiley & Sons.
- Grossberg, S., Mingolla, E., & Ross, W. D. (1994, Jul). A neural theory of attentive visual search: interactions of boundary, surface, spatial, and object representations. *Psychological Review*, *101*(3), 470–489.
- Horwitz, G. D., & Albright, T. D. (2005). Paucity of chromatic linear motion detectors in macaque v1. *Journal of Vision*, *5*(6), 525–533. Available from <http://dx.doi.org/10:1167/5.6.4>
- Huang, L., & Pashler, H. (2005, Jan). Expectation and repetition effects in searching for featural singletons in very brief displays. *Perception & Psychophysics*, *67*(1), 150–157.
- Hubel, D. H., & Livingstone, M. S. (1985). Complex-unoriented cells in a subregion of primate area 18. *Nature*, *315*(6017), 325–327.
- Hubel, D. H., & Livingstone, M. S. (1987, Nov). Segregation of form, color, and stereopsis in primate area 18. *Journal of Neuroscience*, *7*(11), 3378–3415.
- Hubel, D. H., & Wiesel, T. N. (1959, Oct). Receptive fields of single neurones in the cat's striate cortex. *Journal of Physiology*, *148*, 574–591.
- Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, *40*(10-12), 1489–1506.
- Itti, L., & Koch, C. (2001, Mar). Computational modeling of visual attention. *Nature Reviews Neuroscience*, *2*(3), 194–203.
- Johnson, E. N., Hawken, M. J., & Shapley, R. (2001, Apr). The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nature Neuroscience*, *4*(4), 409–416. Available from <http://dx.doi.org/10.1038/86061>

- Johnston, E. B., Cumming, B. G., & Parker, A. J. (1993). Integration of depth modules: stereopsis and texture. *Vision Research*, *33*(5-6), 813–826.
- Kastner, S., DeSimone, K., Konen, C. S., Szczepanski, S. M., Weiner, K. S., & Schneider, K. A. (2007, May). Topographic maps in human frontal cortex revealed in memory-guided saccade and spatial working-memory tasks. *Journal of Neurophysiology*, *97*(5), 3494–3507. Available from <http://dx.doi.org/10.1152/jn.00010.2007>
- Katzner, S., Busse, L., & Treue, S. (2006). Feature-based attentional integration of color and visual motion. *Journal of Vision*, *6*(3), 269–284. Available from <http://dx.doi.org/10.1167/6.3.7>
- Kiesel, A., Miller, J., & Ulrich, R. (in press). Systematic biases and type i error accumulation in tests of the race model inequality. *Behavior Research Methods*.
- Knierim, J. J., & Essen, D. C. van. (1992, Apr). Neuronal responses to static texture patterns in area v1 of the alert macaque monkey. *Journal of Neurophysiology*, *67*(4), 961–980.
- Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology*, *4*(4), 219–227.
- Koene, A. R., & Zhaoping, L. (2007). Feature-specific interactions in salience from combined feature contrasts: evidence for a bottom-up saliency map in v1. *Journal of Vision*, *7*(7), 6.1–6.14. Available from <http://dx.doi.org/10.1167/7.7.6>
- Kristjansson, A. (2007). "i know what you did last trial" - a selective review of research on priming in visual search. *Frontiers in Bioscience*. (submitted manuscript)
- Krummenacher, J., & Müller, H. J. (2007). The role of saliency signals in parallel coactive processing of dimensional information. *unpublished manuscript*.
- Krummenacher, J., Müller, H. J., & Heller, D. (2001). Visual search for dimensionally redundant pop-out targets: Evidence for parallel-coactive processing of dimensions. *Perception & Psychophysics*, *63*(5), 901-917.
- Krummenacher, J., Müller, H. J., & Heller, D. (2002). Visual search for dimensionally redundant pop-out targets: Parallel-coactive processing of dimensions is location specific. *Journal of Experimental Psychology: Human Perception and Performance*, *28*(6), 1303-1322.
- Leventhal, A. G., Thompson, K. G., Liu, D., Zhou, Y., & Ault, S. J. (1995, Mar). Concomitant sensitivity to orientation, direction, and color of cells in layers 2, 3, and 4 of monkey striate cortex. *Journal Neuroscience*, *15*(3 Pt 1), 1808–1818.
- Li, Z. (2002). A saliency map in primary visual cortex. *Trends in Cognitive Sciences*, *6*(1), 9-16.
- Livingstone, M. S., & Hubel, D. H. (1984). Anatomy and physiology of a color system in the primate visual cortex. *Journal of Neuroscience*, *4*, 309-356.
- Livingstone, M. S., & Hubel, D. H. (1987, Nov). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *Journal of Neuroscience*, *7*(11), 3416–3468.

- Livingstone, M. S., & Hubel, D. H. (1988, May). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science*, *240*(4853), 740–749.
- Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997, Jan). Neural mechanisms of spatial selective attention in areas v1, v2, and v4 of macaque visual cortex. *Journal of Neurophysiology*, *77*(1), 24–42.
- Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory - a user's guide* (2nd ed.). Mahwah, New Jersey: Lawrence Erlbaum Associates.
- Maljkovic, V., & Nakayama, K. (1994, Nov). Priming of pop-out: I. role of features. *Memory & Cognition*, *22*(6), 657–672.
- Maljkovic, V., & Nakayama, K. (1996, Oct). Priming of pop-out: II. the role of position. *Perception & Psychophysics*, *58*(7), 977–991.
- Marzi, C. A., Smania, N., Martini, M. C., Gambina, G., Tomelleri, G., Palamara, A., et al. (1996, Jan). Implicit redundant-targets effect in visual extinction. *Neuropsychologia*, *34*(1), 9–22.
- Mazurek, M. E., Roitman, J. D., Ditterich, J., & Shadlen, M. N. (2003, Nov). A role for neural integrators in perceptual decision making. *Cerebral Cortex (New York, N.Y.)*, *13*(11), 1257–1269.
- McAdams, C. J., & Maunsell, J. H. (2000, Mar). Attention to both space and feature modulates neuronal responses in macaque area v4. *Journal of Neurophysiology*, *83*(3), 1751–1755.
- McPeck, R. M., & Keller, E. L. (2002a, Oct). Saccade target selection in the superior colliculus during a visual search task. *Journal of Neurophysiology*, *88*(4), 2019–2034.
- McPeck, R. M., & Keller, E. L. (2002b, Apr). Superior colliculus activity related to concurrent processing of saccade goals in a visual search task. *Journal of Neurophysiology*, *87*(4), 1805–1815. Available from <http://dx.doi.org/10.1152/jn.00501.2001>
- Meeter, M., & Olivers, C. N. L. (2006). Intertrial priming stemming from ambiguity: A new account of priming in visual search. *Visual Cognition*, *13*, 202–221.
- Miller, J. (1982, Apr). Divided attention: evidence for coactivation with redundant signals. *Cognitive Psychology*, *14*(2), 247–279.
- Miller, J. (1986). Timecourse of coactivation in bimodal divided attention. *Perception & Psychophysics*, *40*, 331–343.
- Miller, J. (1991, Feb). Channel interaction and the redundant-targets effect in bimodal divided attention. *Journal of Experimental Psychology: Human Perception and Performance*, *17*(1), 160–169.
- Miller, J., & Lopes, A. (1988, Dec). Testing race models by estimating the smaller of two true mean or true median reaction times: an analysis of estimation bias. *Perception & Psychophysics*, *44*(6), 513–524.

- Mordkoff, J. T., & Miller, J. (1993). Redundancy gains and coactivation with two different targets: The problem of target preferences and the effects of display frequency. *Perception & Psychophysics*, *53*(5), 527-535.
- Mordkoff, J. T., & Yantis, S. (1991). An interactive race model of divided attention. *Journal of Experimental Psychology: Human Perception and Performance*, *7*(2), 520-538.
- Mordkoff, J. T., & Yantis, S. (1993). Dividing attention between color and shape: Evidence of coactivation. *Perception & Psychophysics*, *53*(4), 357-366.
- Mortier, K., Zoest, W. van, Meeter, M., & Theeuwes, J. (2007). No top-down control of selection in singleton search: Evidence from manual and eye movement responses in singleton detection and localization tasks.
(Submitted manuscript.)
- Müller, H. J., Heller, D., & Ziegler, J. (1995, Jan). Visual search for singleton feature targets within and across feature dimensions. *Perception & Psychophysics*, *57*(1), 1-17.
- Müller, H. J., & Krummenacher, J. (2006). Locus of dimension weighting: Preattentive or postselective? *Visual Cognition*, *14* (4/5/6/7/8), 490-513.
- Müller, H. J., Krummenacher, J., & Geyer, T. (2007). Attentional capture by salient color singleton distractors is modulated by top-down dimensional set. *Journal of Experimental Psychology: Human Perception and Performance*. (In press.)
- Müller, H. J., Reimann, B., & Krummenacher, J. (2003, Oct). Visual search for singleton feature targets across dimensions: Stimulus- and expectancy-driven effects in dimensional weighting. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(5), 1021-1035. Available from <http://dx.doi.org/021>
- Nothdurft, H. C. (2000). Salience from feature contrast: variations with texture density. *Vision Research*, *40*(23), 3181-3200.
- Nozawa, G., Reuter-Lorenz, P. A., & Hughes, H. C. (1994). Parallel and serial processes in the human oculomotor system: bimodal integration and express saccades. *Biological Cybernetics*, *72*(1), 19-34.
- Olivers, C. N. L., & Meeter, M. (2006). On the dissociation between compound and present / absent tasks in visual search: Intertrial priming is ambiguity-driven. *Visual Cognition*, *13*, 1-28.
- Pollack, I., & Hsieh, R. (1969). Sampling variability of the area under the roc-curve and of d'_e . *Psychological Bulletin*, *71*, 161-173.
- Pollmann, S. (2001, Jul). Switching between dimensions, locations, and responses: the role of the left frontopolar cortex. *NeuroImage*, *14* (1 Pt 2), S118-S124. Available from <http://dx.doi.org/10.1006/nimg.2001.0837>
- Pollmann, S., Weidner, R., Müller, H. J., & Cramon, D. Y. von. (2000, May). A fronto-posterior network involved in visual dimension changes. *Journal of Cognitive Neuroscience*, *12*(3), 480-494.

- Pollmann, S., Weidner, R., Müller, H. J., & Cramon, D. Y. von. (2006). Neural correlates of visual dimension weighting. *Visual Cognition*, *17*, 877-897.
- Pollmann, S., Weidner, R., Müller, H. J., Maertens, M., & Cramon, D. Y. von. (2006, Mar). Selective and interactive neural correlates of visual dimension changes and response changes. *NeuroImage*, *30*(1), 254-265. Available from <http://dx.doi.org/10.1016/j.neuroimage.2005.09.013>
- Pollmann, S., & Zaidel, E. (1999, Apr). Redundancy gains for visual search after complete commissurotomy. *Neuropsychology*, *13*(2), 246-258.
- Posner, M. I. (1980, Feb). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3-25.
- Prinzmetal, W., McCool, C., & Park, S. (2005, Feb). Attention: reaction time and accuracy reveal different mechanisms. *Journal of Experimental Psychology: General*, *134*(1), 73-92. Available from <http://dx.doi.org/10.1037/0096-3445.134.1.73>
- R Development Core Team. (2006). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Available from <http://www.R-project.org> (ISBN 3-900051-07-0)
- Raab, D. H. (1962). Statistical facilitation of simple reaction times. *Transactions of the New York Academy of Sciences*, *24*, 574-590.
- Rangelov, D., Zehetleitner, M., & Müller, H. J. (2007). Two different weighting systems for dimensional pre-attentive processing and dimension-specific task sets. (Manuscript in preparation.)
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59-108.
- Ratcliff, R. (1981). A theory of order relations in perceptual matching. *Psychological Review*, *88*, 552-572.
- Ratcliff, R., Cherian, A., & Segraves, M. (2003, Sep). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of Neurophysiology*, *90*(3), 1392-1407. Available from <http://dx.doi.org/10.1152/jn.01049.2002>
- Ratcliff, R., Gomez, P., & McKoon, G. (2004, Jan). A diffusion model account of the lexical decision task. *Psychological Review*, *111*(1), 159-182. Available from <http://dx.doi.org/10.1037/0033-295X.111.1.159>
- Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., & Segraves, M. A. (2007, Feb). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology*, *97*(2), 1756-1774. Available from <http://dx.doi.org/10.1152/jn.00393.2006>
- Ratcliff, R., & Rouder, J. N. (1998). response times for two-choice decisions. *Psychological Science*, *9*(5), 347-356.
- Ratcliff, R., & Smith, P. L. (2004, Apr). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, *111*(2), 333-367. Available from <http://dx.doi.org/10.1037/0033-295X.111.2.333>

- Ratcliff, R., Zandt, T. V., & McKoon, G. (1999, Apr). Connectionist and diffusion models of reaction time. *Psychological Review*, *106*(2), 261–300.
- Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: the need for attention to perceive changes in scenes. *Psychological Science*, *8*, 201–215.
- Robinson, D. A., & Fuchs, A. F. (1969, Sep). Eye movements evoked by stimulation of frontal eye fields. *Journal of Neurophysiology*, *32*(5), 637–648.
- Santee, J. L., & Egeth, H. E. (1982, Aug). Do reaction time and accuracy measure the same aspects of letter recognition? *Journal of Experimental Psychology: Human Perception and Performance*, *8*(4), 489–501.
- Sato, T., Murthy, A., Thompson, K. G., & Schall, J. D. (2001, May). Search efficiency but not response interference affects visual selection in frontal eye field. *Neuron*, *30*(2), 583–591.
- Savazzi, S., & Marzi, C. A. (2004). The superior colliculus subserves interhemispheric neural summation in both normals and patients with a total section or agenesis of the corpus callosum. *Neuropsychologia*, *42*(12), 1608–1618. Available from <http://dx.doi.org/10.1016/j.neuropsychologia.2004.04.011>
- Schall, J. D., Hanes, D. P., Thompson, K. G., & King, D. J. (1995, Oct). Saccade target selection in frontal eye field of macaque. i. visual and premovement activation. *Journal of Neuroscience*, *15*(10), 6905–6918.
- Schall, J. D., & Thompson, K. G. (1999). Neural selection and control of visually guided eye movements. *Annual Review of Neuroscience*, *22*, 241–259. Available from <http://dx.doi.org/10.1146/annurev.neuro.22.1.241>
- Schönwälder, B. (2006). *Mechanisms of target selection and feature binding in visual object recognition: Evidence from the asynchronous presentation of target features*. Unpublished doctoral dissertation, Ludwig-Maximilian-University, Munich.
- Schulman, A. I., & Mitchell, R. R. (1966, Aug). Operating characteristics from yes-no and forced-choice procedures. *Journal of The Acoustical Society Of America*, *40*(2), 473–477.
- Sengpiel, F., & Hübner, M. (1999). Visual perception: Spotlight on the primary visual cortex. *Current Biology*, *9*(9), R318–R321.
- Sigurdardottir, H. M., Kristjansson, A., & Driver, J. (2007). Repetition streaks increase perceptual sensitivity in visual search of brief displays. *Visual Cognition*, *9999*(1), 1350–6285. (Retrieved September 04, 2007, from <http://www.informaworld.com/10.1080/13506280701218364>)
- Sincich, L. C., & Horton, J. C. (2005). The circuitry of v1 and v2: integration of color, form, and motion. *Annual Review of Neuroscience*, *28*, 303–326. Available from <http://dx.doi.org/10.1146/annurev.neuro.28.061604.135731>
- Smith, P. L., & Ratcliff, R. (2004, Mar). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, *27*(3), 161–168. Available from <http://dx.doi.org/10.1016/j.tins.2004.01.006>

- Sternberg, S. (1969a). Attention and performance. In W. G. Koster (Ed.), (p. 276-315). North-Holland.
- Sternberg, S. (1969b). Memory-scanning: mental processes revealed by reaction-time experiments. *American Scientist*, *57*(4), 421-457.
- Theeuwes, J. (1991, Aug). Cross-dimensional perceptual selectivity. *Perception & Psychophysics*, *50*(2), 184-193.
- Theeuwes, J. (1992, Jun). Perceptual selectivity for color and form. *Perception & Psychophysics*, *51*(6), 599-606.
- Theeuwes, J., Kramer, A. F., & Belopolsky, A. V. (2004, Aug). Attentional set interacts with perceptual load in visual search. *Psychonomic Bulletin & Review*, *11*(4), 697-702.
- Theeuwes, J., Reimann, Brit, & Mortier, K. (2006, December). Visual search for featural singletons: No top-down modulation, only bottom-up priming. *Visual Cognition*, *14*(4-8), 466-489. Available from <http://dx.doi.org/10.1080/13506280500195110>
- Thompson, K. G., & Bichot, N. P. (2005). A visual salience map in the primate frontal eye field. *Progress In Brain Research*, *147*, 251-262. Available from [http://dx.doi.org/10.1016/S0079-6123\(04\)47019-8](http://dx.doi.org/10.1016/S0079-6123(04)47019-8)
- Thornton, T. L., & Gildea, D. L. (2007, Jan). Parallel and serial processes in visual search. *Psychological Review*, *114*(1), 71-103. Available from <http://dx.doi.org/10.1037/0033-295X.114.1.71>
- Todd, J. W. (1912). *Reaction to multiple stimuli*. New York: Science Press.
- Töllner, T., Gramann, K., Müller, H. J., & Eimer, M. (2007a). The anterior n1 component as an index of modality shifting.
(Unpublished manuscript.)
- Töllner, T., Gramann, K., Müller, H. J., & Eimer, M. (2007b). Electrocortical correlates of modality changes between vision and touch.
(Submitted manuscript.)
- Töllner, T., Gramann, K., Müller, H. J., Kiss, M., & Eimer, M. (2007). Electrophysiological markers of visual dimension changes and response changes. *Journal of Experimental Psychology: Human Perception and Performance*. (In press.)
- Townsend, J. T., & Ashby, F. G. (1983). *The stochastic modeling of elementary psychological processes*. Cambridge, UK: Cambridge University Press.
- Townsend, J. T., & Nozawa, G. (1995). Spatio-temporal properties of elementary perception: An investigation of parallel, serial, and coactive theories. *Mathematical Psychology*, *39*, 321-359.
- Townsend, J. T., & Nozawa, G. (1997). Serial exhaustive models can violate the race model inequality: Implications for architecture and capacity. *Psychological Review*, *104*(3), 595-602.

- Treisman, A. (1986). Features and objects in visual processing. *Scientific American*, 255, 144B-125.
- Treisman, A. (1988, May). Features and objects: the fourteenth bartlett memorial lecture. *Quarterly Journal of Experimental Psychology A*, 40(2), 201-237.
- Treisman, A., & Gelade, G. (1980). A feature-integration theory of vision. *Cognitive Psychology*, 12, 97-136.
- Treisman, A., & Sato, S. (1990, Aug). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception and Performance*, 16(3), 459-478.
- Ts'o, D. Y., & Gilbert, C. D. (1988). The organization of chromatic and spatial interactions in the primate striate cortex. *Journal of Neuroscience*, 8, 1712-1727.
- Tsotsos, J. K. (1990). Analyzing vision at the complexity level. *Behavioral and Brain Sciences*, 13, 423-469.
- Turatto, M., Mazza, V., Savazzi, S., & Marzi, C. A. (2004). The role of the magnocellular and parvocellular systems in the redundant target effect. *Experimental Brain Research*, 158, 141-150.
- Usher, M., & McClelland, J. L. (2001, Jul). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, 108(3), 550-592.
- Vandekerckhove, J., & Tuerlinckx, F. (2007a). Diffusion model analysis with matlab: a smart primer. *Behavior Research Methods*. (In press.)
- Vandekerckhove, J., & Tuerlinckx, F. (2007b). Fitting the ratcliff diffusion model to experimental data. *Psychonomic Bulletin & Review*. (In press.)
- Weidner, R., Pollmann, S., Müller, H. J., & Cramon, D. Y. von. (2002, Mar). Top-down controlled visual dimension weighting: an event-related fmri study. *Cerebral Cortex (New York, N.Y.)*, 12(3), 318-328.
- Wolfe, J. M. (1994). Guided search 2.0 - a revised model of visual search. *Psychonomic Bulletin and Review*, 1, 202-238.
- Wolfe, J. M. (1998a). Attention. In H. Pashler (Ed.), (p. 13-73.). Psychology Press.
- Wolfe, J. M. (1998b). What can 1 million trials tell us about visual search? *Psychological Science*, 9(1), 33-39.
- Wolfe, J. M., & Horowitz, T. S. (2004, Jun). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews. Neuroscience*, 5(6), 495-501. Available from <http://dx.doi.org/10.1038/nrn1411>
- Zandt, T. van. (2002). Analysis of response time distributions. In *Stevens' handbook of experimental psychology (3rd edition), volume 4: Methodology in experimental psychology* (pp. 461-515). New York: Wiley Press.
- Zehetleitner, M. (2007a). Co-activation - neither serial exhaustive nor interactive - models can explain violations of the race model inequality in visual pop-out search. (Chapter 3 of the present thesis.)

Zehetleitner, M. (2007b). Intention and trial history modulate dimensional weights in localization of pop-out targets.

(Chapter 4 of the present thesis.)

Zehetleitner, M. (2007c). What the redundant-signals paradigm reveals about pre-attentive visual processing.

(Chapter 2 of the present thesis)

VITA

Michael Zehetleitner

Born on 1. February 1977 in Kempten

Education

- 1997 Abitur, Gymnasium Donauwörth
- 1998 Mathematics, University of Heidelberg
- 1999-2001 Mathematics and computer science,
 Ludwig-Maximilians-Universität Munich (Vordiplom)
- 2001-2004 Psychology, Ludwig-Maximilians-Universität Munich
- 2004-2006 Neuro-cognitive psychology,
 Ludwig-Maximilians-Universität Munich (M.Sc.)
- 2006-2007 Ph.D. in psychology, Ludwig-Maximilians-Universität Munich

Professional Experience

- 2000-2006 IT system development, self-employed
- 2006-2007 Research Fellow, Department of psychology,
 Ludwig-Maximilians-Universität Munich