Guiding Features in Visual Search

A Dissertation Presented

by

Robert George Alexander

to

The Graduate School

in Partial Fulfillment of the

Requirements

for the Degree of

Doctor of Philosophy

in

Psychology

Stony Brook University

December 2013

Stony Brook University

The Graduate School

Robert George Alexander

We, the dissertation committee for the above candidate for the

Doctor of Philosophy degree, hereby recommend

acceptance of this dissertation.

Gregory Zelinsky – Dissertation Advisor
Professor, Experimental Psychology

Nancy Franklin - Chairperson of Defense
Associate Professor, Experimental Psychology

Hoi-Chung Leung – Committee Member
Associate Professor, Biopsychology

Carrick Williams - Outside Member
Associate Professor, Psychology – Mississippi State University

This dissertation is accepted by the Graduate School

Charles Taber
Dean of the Graduate School

Abstract of the Dissertation

Guiding Features in Visual Search

by

Robert George Alexander

Doctor of Philosophy

in

Psychology

Stony Brook University

2013

Visual search is the task of finding something with an uncertain location. Much research has been done on the features used in visual search, but most previous work has used only simple stimuli, like colored bars. Although people debate the features of even simple stimuli, most researchers agree that hue, shape, and orientation are important for guiding search. This study tested the role of hue, shape, and orientation in how we search for real-world objects, and estimates the importance of these features, both separately and in combination, and relates this importance to how reliably each feature defines the target. Nine experiments were conducted in order to describe mismatch effects, in which a decrease in performance is caused by differences between the way a target looks when participants are told to search for it and the way it looks when they actually find it (if you're looking for a red car, but the car is actually blue, this will make your search harder, but only if you are using color as a guiding feature). These mismatch effects provide an indicator of how important each feature is in guiding search (with performance decrements increasing with mismatch for features that are important to the task), a means for testing effects of feature reliability (by varying the probability of mismatch in a given feature dimension), and the combination of mismatch effects across different features (by manipulating more than one dimension at a time). Results demonstrated that search is largely driven by color information and that shape and orientation are only used in guiding search when color is not available to perform the task. Mismatch effects were more pronounced when participants were uncertain of what feature dimension would change, suggesting that you are less likely to rely on features which are known to change. Mismatch effects with combinations of feature dimensions also demonstrate that the different feature dimensions do not sum linearly, though most models of visual search predict a linear sum. A number of modifications to current models of search are suggested so as to adequately account for these mismatch effects.

Table of Contents

List of Figures

Chapter 1 - Introduction

1.1 Paper structure

This paper is organized as follows. Section 1 gives a general introduction to the relevant literatures, including background information on previous research characterizing guiding features, as well as existing work varying the mismatch between preview and search target, and a discussion of current models of visual search and their predictions regarding feature summation. Section 2 outlines the specific aims of the investigation. Section 3 outlines the general methods for the current study. Section 4 outlines the specific design and methods for each experiment, provides background information specific to each, and includes results and discussions for individual experiments. Finally, Section 5 provides a summary and a brief discussion.

1.2 Previous behavioral work characterizing search guidance

Visual search is the task of directing your attention to something when its location is uncertain. You do this countless times in daily life, whether looking for a friend in a crowd, your car in a parking lot, the computer icon to load a particular program, or simply looking for a pen on the table in front of you so that you can pick it up. Directing gaze to visual cues is important for driving, for playing sports effectively, and for a wide variety of other daily activities.

Search is an active process. When you are looking for a tomato, you probably don't look around randomly until you happen to be looking at one. Instead, you actively guide your gaze towards things that are red and round, and to places that tomatoes are likely to be. *Guidance*, the non-random direction of your attention and gaze to locations likely to contain targets (either because those locations contain target-relevant visual features, such as the color red in the above example, or due to other factors that allow search to proceed non-randomly), is an important aspect of visual search.

Understanding search is important for a variety of reasons. Most obviously, since we engage in search quite often, understanding search would help us better understand a significant part of how we navigate, interact with, and experience the world. Understanding search guidance would help us understand how attention is deployed, which has implications for identifying clinical deficiencies in attention. Understanding guidance may also lead to improved techniques or equipment for performing many important tasks which involve effectively guiding attention, such as finding tumors in x-rays, weapons in luggage, detecting combatants in military situations, and pursuing criminals through crowds.

It is generally accepted that various simple features can be detected and used in early stages of visual processing, in parallel across the display, prior to the deployment of attention or gaze to specific objects. This early processing presumably provides the information used to guide attention or eye movements to particular objects (Wolfe, 1994). However, while there is a consensus on a few guiding features (color, motion, orientation, and size), the feature dimensions that could be used in guiding search are largely unknown even in arrays composed of simple stimuli (Wolfe & Horowitz, 2004).

Researchers have relied on a variety of converging measures for identifying guiding features. If the slope of the reaction time by set size function is shallow for search for the target feature, and remains shallow even when there is distractor heterogeneity, if the feature could be "effortlessly" segmented from a background texture (preattentive texture segmentation), if the feature can contribute to illusory conjunctions (where the features are correctly identified but incorrectly combined across objects), and if search for the presence of the feature is easier than search for the absence of the feature, researchers then may have "some assurance" that the feature can guide visual search (Wolfe & Horowitz, 2004, p. 4). However, there are problems with using any of these criteria to determine what features guide search, with the strongest conclusions arising from converging evidence.

Early work on what features could be used to guide search relied primarily on *search slopes*, the degree to which reaction time (RT) increases with the addition of more items to the search display. If the slope was flat, indicating that additional items in the display did not increase RT, this was interpreted as evidence that the feature dimension that defined the target could be detected in parallel and used to guide search. These efforts, though informative, were met with a profound difficulty: RT slopes are not as easy to interpret as one might hope. Search tasks cannot simply be split into "parallel" and "serial" search based on slopes (Palmer, 1995; Townsend, 1990; Wolfe, 1998). In addition, there is no single search mechanism that completely determines reaction time. Manual RT measures, while useful for assessing task difficulty, do not tell us *why* a task is difficult. Long RTs could indicate that search was not guided to the target or that subjects simply required a long time to make their response after the target was fixated (Zelinsky & Sheinberg, 1997). RTs are also influenced by decision criteria: Longer RTs might simply reflect a preference for accuracy in a given task. Items added to the search display could increase RTs by decreasing search guidance, but they could also increase RTs by decreasing confidence (thus increasing decision times) or by making processing more difficult due to the addition of clutter.

Search asymmetries, where search *for* a feature is not affected by the number of distractors in the display but the number of distractors strongly affects search for the *absence* of the same feature (Treisman & Souther, 1985) has the same potential confounds as RT measures in general. In tasks where subjects are searching for the absence of the feature, distractors must (by definition) have that feature. If distractor objects become difficult to process due to the presence of that feature, which may make the distractors more complex or more self-cluttered, RT slopes might become steeper for reasons completely unrelated to guidance effects (but see Zelinsky & Sheinberg, 1997, for evidence that asymmetries could be reflected in eye movement measures). Stimulus factors, such as target-distractor similarity differences, can also result in search asymmetries (Rosenholtz, 2001).

Preattentive texture segmentation also does not necessarily indicate that a feature can guide search. Wolfe, Friedman-Hill, Stewart, and O'Connell (1992) demonstrated that some stimuli that could be "effortlessly" segmented from a background texture cannot be efficiently found in a search task, and some stimuli that result in efficient search cannot be “effortlessly” segmented from background textures.

Illusory conjunctions cannot be directly observed, and are instead inferred from error rates. The use of illusory conjunctions to test whether a feature can guide search relies on the assumption that illusory conjunctions only occur prior to attention being directed to the objects. However, Navon and Ehrlich (1995) found equivalent feature combination errors under conditions of attention and inattention, calling this assumption into question. Further, illusory conjunctions may not even exist, and may instead be incorrect interpretations of experimental artifacts (Donk, 1999, 2001). Earlier work may have overestimated the rate of illusory conjunctions or found evidence for illusory conjunctions when none actually occurred, for a variety of reasons: uncontrolled stimulus factors that allowed for sophisticated guessing strategies (Johnston & Pashler, 1990), feature misperceptions (Ashby, Prinzmetal, Ivry, & Maddox, 1996), postperceptual factors (Navon & Ehrlich, 1995), or failures to account for target-nontarget confusions (Donk, 1999).

Perhaps the most important limitation with these criteria is that they often lead to inconsistent or unconvincing conclusions. Past difficulty in identifying what features can guide search is likely due to the fact that movements of attention are not directly observable. Instead, it has been necessary to infer shifts of attention from manual RTs and other indirect measures, such as those discussed above. This inability to directly test guidance may explain why researchers are often hesitant to accept or reject features as guiding search. This lack of confidence may be best evident in Wolfe and Horowitz, 2004, in which all but five particularly well-studied features (color, motion, orientation, length, and spatial frequency) were classified as “probable”, “possible”, and “doubtful” sources of search guidance).

Recently researchers have begun to measure eye movements during search tasks. While it is possible for attention to be directed to a location in space without the execution of an eye movement (Klein & Farrell, 1989; Murthy, Thompson, & Schall, 2001; Posner, 1980), when an eye movement is made it is preceded by a shift of attention to that spatial location (e.g. Deubel & Schneider, 1996; Peterson, Kramer, & Irwin, 2004; see Findlay & Gilchrist, 2003, for additional discussion of the relationship between eye movements and attention, and Zelinsky, Rao, Hayhoe, and Ballard, 1997, for a discussion in the context of visual search). Eye movement recordings allow us to directly test the time it takes gaze to reach a target, whether gaze could be directed to a particular feature in the visual periphery, and how much of manual RT effects are actually the result of processes that occur after targets have been located and gaze directed to them. As a result of gaze being directly observable, eye movement measures typically provide more information than button press RT measures, and do not require inferences to be made regarding the relationship between attention and performance in a given task. In addition, eye movements have rich theoretical implications for models of visual search (discussed below, in 1.3). Many models dovetail nicely with the previous work discussed above, and can predict the reaction time patterns commonly involved in search, but eye movements provide a means of requiring that predictions be more precise and better distinguishing between alternative search models.

1.3 Models of visual search

The most common conception of how visual search operates posits that the search displays are represented by visual features, which are compared to the features of the search target in order to either detect or identify the target or to generate a guidance signal which directs attention to the target or target-like distractors. Search performance is often explained in terms of similarity along these feature dimensions. Since search is accomplished through the use of target features, search becomes more difficult when distractor objects are similar to the target (e.g. Duncan & Humphreys, 1989; Treisman, 1991). This relationship between target-distractor similarity and search performance is key to most search theories and models (e.g. Duncan & Humphreys, 1989; Pomplun, 2006; Treisman & Gelade, 1980; Treisman & Sato, 1990; Wolfe, 1994; Wolfe, Cave, & Franzel, 1989 ; Zelinsky, 2008). There are also models of search that do not include or emphasize the use of target features to facilitate the search task (e.g. Itti & Koch, 2000; Parkhurst, Law, & Niebur, 2002), but these models will be excluded from the following discussion as the current work is focused on the exploration of these top-down target features.

Attentional Engagement Theory (Duncan & Humphreys, 1989) posits that there are three components to search processes. In the first, a "perceptual description" is formed, consisting of a parallel, preattentive, representation of the features present in the search display. The display is broken down into "structural units" or "parts", each of which contains descriptors of its visual features and some nonvisual properties. Next, this input is matched to an internal top-down template of task-related features. Finally, structural units enter visual short-term memory through a limited-capacity process, allowing subjects to make manual responses to the stimuli and for the stimuli to reach awareness. Targets compete with distractor items that have features similar to targets (those that better match the internal template) in this last stage of selection.

Feature Integration Theory (FIT; Treisman & Gelade, 1980) posits that each feature dimension (such as color, orientation, shape, etc.) is processed in a separate neuronal map (called a "feature map"), and feature values within each feature map are processed and identified early, automatically, and in parallel across the search display. Once this parallel processing occurs, if the target is defined along a single basic feature and no distractors share that feature, the target will "pop-out" and could be identified immediately. Shifts of attention or additional processing are not needed, as the relevant feature value is already detected and identified, and subjects can immediately respond that the target is present. However, while features are processed in parallel, objects are processed in a later, serial stage. FIT predicts that in order to combine feature values across feature dimensions—a prerequisite for the creation of multi-feature objects—attention must be directed serially to the location of those feature values. If this does not occur, features could be incorrectly combined (illusory conjunctions). As a result, if multiple features need to be integrated within a single object in order to identify the target (as is the case with conjunction search, where targets are defined by multiple features that are also present on distractor items), search proceeds serially, with attention directed to one object at a time until the target is selected. The order of objects attended until the target is detected was essentially random, although subjects could use visual salience or strategies to direct attention to locations more likely to contain the target. A later revision of FIT (Treisman & Sato, 1990) supplements the moving "window" or "spotlight" of attention that has to be serially directed to each object with a process of feature-based inhibition. In this process, "attention" can control the level of activation on a "master map" on which the levels of activation from each feature map are combined. Locations that have a particular feature value could be inhibited through their connection with the relevant feature map. The signal-to-noise ratio in a search task could then be improved by decreasing the activation of locations that contain features that are only present on distractors. This process could also be performed on multiple locations at once, but when distractor features are similar to target features the signal-to-noise ratio would not be sufficient to locate the target and subjects would need to focus attention on one small region at a time.

Guided Search (Wolfe, 1994; Wolfe, Cave, & Franzel, 1989) is a two-stage model of visual search. The first stage is a preattentive, parallel coding of features to form feature maps (as in FIT), though the features in Guided Search are broadly-tuned channels (orientations, for example, are coded as "right-tilted" or "left-tilted", and "steep" or "shallow", rather than as specific values like 17°). Activation on these maps is determined both by local differences in features (bottom-up activation) and by a top-down task-related weighting of feature channels. While FIT does not specify how feature maps are then combined across dimensions into a master map (only that they are combined), Guided Search predicts that a master map (or "activation map") is then created through a weighted linear sum of each feature map. Attention is then directed or "guided" to locations with the highest overall activation. Recognition occurs only for attended locations, and if the location with the highest activation does not contain the target, attention is directed to the location with the second-highest activation, followed by the third, and so on.

A number of models are based largely on the same processes as the Guided Search model. Navalpakkam and Itti's (2005) model is in large part a computational implementation of Guided Search. Bottom-up and top-down features are combined through a weighted linear sum after spatial competition and non-linear interactions occur within the feature maps. The weights used for combining feature maps are learned from images containing the target, and recognition is modeled using the same features that are used to guide search to the target. Cave's (1999) FeatureGate model also follows a similar process as Guided Search, but is a neural network model that uses "attentional gates" rather than attention in order to limit capacity. Maps within FeatureGate are also combined through linear summation, and include both top-down and bottom-up activations.

Some models, designed to handle visually complex image-based stimuli, have implemented top-down target templates by determining the visual features of a target image (or category) and comparing those against the search display (e.g. Ehinger, Hidalgo-Sotelo, Torralba, & Oliva, 2009; Rao, Zelinsky, Hayhoe, & Ballard, 1996; Zelinsky, 2008). These models are designed to predict eye movements, not simply search RTs, although shifts of attention are expected to occur in combination with shifts in gaze. The Target Acquisition Model (Zelinsky, 2008), for example, performs a retinal transform of search displays (simulating visual acuity loss in the periphery), and then extracts feature values from the transformed displays and from target images with a variety of feature detectors. Target and scene representations are compared, and the visual similarity between them is used to create a "target map" that serves a role similar to the master or activation maps of FIT and Guided Search: Search progresses to points of activation on that map through a combination of parallel and serial processes.

Some models of visual search do not include a serial component. Instead, subjects are assumed to process features from across the visual field in parallel, and make their responses based on information accumulated peripherally. For example, Signal Detection Theory (Green & Swets, 1966) and Ideal Observer Theory (e.g., Geisler, 1989;) have been applied to search and have been successful in predicting many common findings without relying on a serial component (e.g. Eckstein, 1998; Palmer, Verghese, & Pavel, 2000). Eckstein, Thomas, Palmer, and Shimozaki (2000) presented a signal detection model of search in which each element in a search array gives rise to a noisy internal representation of the relevant feature dimensions for that element. The representations are then combined across dimensions and a decision is made, either in response to the maximum value of a weighted linear sum of the features across dimensions (detecting the target) or the maximum value of the feature that provides the *minimum* response (detecting the item that is *least* like a distractor, and thus most likely to be the target).

The extended generalized context model for visual search (EGCM-VS; Guest & Lamberts, 2011) also does not include a serial mechanism. EGCM-VS posits that every display item is processed in parallel, and evidence for target presence is a linear sum of the similarity in feature values between the items and the top-down representation of the search target over time. The similarity relationships change over time, and poor performance in search tasks may be due to relevant features needing more time to be processed.

Parallel models generally do not include serial shifts of attention or eye movements, and additional assumptions would need to be added to the models to explain why these shifts occur. One possible explanation is that subjects could collect information about display objects in parallel, but fixate objects in an effort to increase confidence about target presence (Kotowicz, Rutishauser, & Koch, 2010; Zelinsky, 2008). If this is the case, eye movements could be used as an online indicator of what display object is the current best guess for target location. Serial shifts in attention would then be expected in response to parallel feature processing, and might indicate what features had been processed (if subjects fixated objects that were similar in color to the target, for instance, this would indicate that color had been processed), even though these shifts would not contribute to task decisions.

1.4 Simple versus Realistic Stimuli

Most previous work on the features guiding search, both in behavioral testing and in modeling efforts, has used relatively simple stimuli. More recently, researchers have begun using more natural contexts, such as arrays of common objects (e.g., Castelhano, Pollatsek, & Cave, 2008; Schmidt & Zelinsky, 2009; Williams, 2010; Williams, Henderson, & Zacks, 2005;Yang, Chen, & Zelinsky, 2009) and computer-generated or fully realistic scenes (e.g., Brockmole & Henderson, 2006; Bravo & Farid, 2009; Eckstein, Drescher, & Shimozaki, 2006; Foulsham & Underwood, 2011; Neider & Zelinsky, 2006, 2008; Oliva, Wolfe, & Arsenio, 2004).

With this move to more realistic stimuli, it is important to ensure that core findings from earlier search experiments generalize beyond simple displays. The complexity involved in realistic stimuli, compared to simple stimuli, may impact how search is guided to these objects. If a complex target with thousands of features was presented with equally complex distractor objects, it is far less clear what features will be used to locate the target than when a green vertical bar target is presented among green horizontal distractor bars. When numerous feature dimensions are required to discriminate a target, search may operate in a qualitatively different fashion, possibly due to visual working memory constraints (Alvarez & Cavanagh, 2004; Luck & Vogel, 1997; Zelinsky & Loschky, 2005), or differences in how similarity relationships are determined (see Alexander & Zelinsky, 2011, for a discussion).

Some recent findings suggest that search with real-world objects uses the same kind of similarity relationships as simple stimuli. Alexander and Zelinsky (2011) found that visual similarity ratings provided by human subjects predicted categorical search performance with real-world objects. Subjects who were searching for butterflies took longer on trials composed of distractors that were ranked as more similar to butterflies, and were more likely to fixate more similar distractors. Similarly, subjects searching for teddy bears were more distracted by distractors that were ranked as similar to teddy bears. Computational modeling also demonstrated that the similarity effects found in these experiments were due to matches in visual features between the targets and distractors.

In another study, Alexander and Zelinsky (2012) manipulated the similarity between real-world objects (teddy bears) by transplanting parts between target and distractor objects. When distractors had more parts in common with the target, search became more difficult, replicating the standard effect of target-distractor similarity often found with simple stimuli. Another classic visual similarity effect, that search is easier when distractors are more similar to other distractors (Duncan & Humphreys, 1989) also replicated with these real-world images. Unfortunately, it is unclear what visual features were included in the similarity manipulation, since features were transferred in terms of parts rather than feature values.

With real-world images, the most straightforward method of assessing what features guide search is to correlate where subjects' eye movements land with what features are present in the search display. If subjects look more at objects which share the same color as the target than objects which share the same shape as the target, it could be inferred that subjects used the color of the target to guide their search. When a given dimension is highly predictive of eye movements during search, that dimension likely is guiding search in that task (see Hwang, Higgins, and Pomplun, 2009 and Pomplun, 2006, for examples of this method). While correlating features in the search display with eye movements can potentially inform us about what features were most important in guiding search to the target, the method relies on correlational inference rather than experimental control. If subjects' fixations are highly correlated with regions of a scene that share the same color as the target, this does not necessarily mean that color was used to guide search. The regions in the scene have the same color information as the target might also match the target on some other feature dimension, may be contextually important, or may be highly salient. Experimental approaches are needed to provide direct tests of these features.

Chapter 2 Specific Aims

A series of experiments were conducted exploring the contribution of several features to the guidance of visual search. Specific aims and predictions are detailed below.

2.1 Aim 1: Investigate the effects of mismatch between previews and search targets on orientation, hue, and shape feature dimensions.

Rather than relying on a correlational approach or forcing the search objects to be identical on all but one feature dimension (a near impossibility in the case of realistic objects), the current study examines search guidance using a method borrowed from Vickery, King, and Jiang (2005), in which the similarity between the search preview that precedes each trial and the search target that appears in the actual search array is manipulated. By varying the feature difference between the preview and the search target, rather than between the target and distractors, distractors could be free to vary along many feature dimensions, as is likely to occur in real life (see also Vickery et al., 2005). Since the target is uniquely defined on more than one feature dimension, how much each feature dimension is used when multiple dimensions are available to guide search could be explored. Unlike manipulating individual feature dimensions in isolation to test whether those features *could be* used to guide search, this methodology tests whether those features *are* used even when the task could be performed without their use.

In a series of experiments, Vickery et al. (2005) manipulated the difference in 2d orientation between the preview and the search target, with both simple geometric stimuli and with models of real-world objects. In addition, preview-target mismatch in the size dimension was tested with simple objects (Experiments 2 and 3) and mismatch in 3d orientation was tested with real-world objects (Experiments 4 and 5). Vickery et al. used this method to demonstrate that cueing subjects with a preview that exactly matches the search target results in faster search RTs than previews that differ in orientation from the target, or previews that only contained the semantic identity of the target. The degree of orientation mismatch was found to affect performance only for real-world objects (both for 2d and 3d orientation) up to 90° of change, with greater mismatch resulting in longer RTs (the effects plateaued after 90°). For simple objects, no reliable effect of the degree of orientation mismatch was found: Any mismatch resulted in longer RTs, but the *amount* of mismatch was irrelevant to performance. The authors speculated that their simple stimuli were too ambiguous in orientation to result in effective mental rotation, but subjects may have been efficiently rotating the real-world objects during search (with larger degrees of mental rotation translating into longer RTs).

Bravo and Farid (2009) also manipulated the distance between previews and search targets in terms of orientation, and extended Vickery et al.'s conclusions to changes in size and reflection across the vertical axis with real-world objects. Whenever the preview was altered, performance was worse than when the preview was an exact match to the target. Whether the degree of mismatch mattered was not evaluated.

Because only RT and accuracy measures were used in the Bravo and Farid (2009) and Vickery et al. (2005) experiments, it is unclear whether these mismatch effects were caused by decreases in guidance or an increase in the amount of time it took to reject distractors or confirm that an object was the target. By using this task but including eye movement measures, the current study teases these factors apart. In addition, both Bravo and Farid (2009) and Vickery et al. (2005) used tasks that may have overestimated the impact of mismatch by preventing or limiting the use of color to guide search: Vickery et al. (2005) used only grayscale stimuli, and Bravo and Farid (2009) used very colorful background images (coral reef) that likely created some target-background similarity, especially as their target images were all of sea animals, which may have evolved their colors to match those of coral. The current study examines the effects of mismatch both when color information is unavailable to guide search (by using grayscale stimuli in "grayscale" experiments) and when color could be used to guide search (in "full-color" experiments using color stimuli).

Further support for the existence of preview-target mismatch effects can be found in experiments comparing search performance between trials where the preview was an exact match to the search target to trials where the preview was the category name of the target (Bravo & Farid, 2009; Castelhano, Pollatsek, & Cave, 2008; Malcolm & Henderson, 2009; Maxfield & Zelinsky, 2012; Schmidt & Zelinsky, 2009; Vickery et al., 2005; Wolfe, Horowitz, Kenner, Hyle, & Vasan, 2004; Yang & Zelinsky, 2009). All six papers reported that search RTs were longer with a text cue (category name) rather than an exact picture cue. Castelhano et al. (2008), Malcolm and Henderson (2009), Schmidt and Zelinsky (2009), and Yang and Zelinsky (2009) measured eye movements, and found that saccades were directed to the target more quickly with picture cues than with text cues. Castelhano et al. and Malcolm and Henderson also found that the time spent after fixating the target was also faster with picture cues, demonstrating that cue type affected both guidance and verification processes. When given the text cue, subjects presumably formed a mental representation of the target that did not match the search target as precisely as the exact picture cue, and the difference in RTs and eye movements in these experiments could be due to a preview-target mismatch effect.

Aim 1 replicates and extends findings describing a preview-target mismatch effect (e.g. Bravo & Farid, 2009; Vickery et al., 2005) to new feature dimensions (color and aspect ratio) and specifies what aspects of search these mismatches affect (determining, for instance, whether these effects are due to changes in guidance or other processes). Previous work has demonstrated that RTs increase when there is a mismatch between preview and target. The current work identifies whether these manual RT effects are due to changes in search guidance (how long it takes subjects to fixate the target) or due to other processes (i.e. distractor rejection or target verification). Because effects of feature mismatch demonstrate that subjects are using information from that feature dimension, it is of obvious importance for modeling efforts to determine what aspects of the search task actually show mismatch effects.

The current study also explores a wide range of mismatch differences within each feature dimension, in order to better characterize the effects of each. Previous work either did not test whether the *amount* of mismatch affected performance (Bravo & Farid, 2009; Wolfe et al., 2004), or tested only a few large (30°) increments of orientation (Vickery et al., 2005). The current experiments also explore the full range—from 0-360°—of two of the feature dimensions: Hue angle (described in section 3.3.2) and orientation. A fuller exploration of this range allows maximum and average effects of mismatch on guidance to be directly compared across feature dimensions, indicating relative judgments of feature importance to search.

The question of which feature dimensions are most important in guiding search has been a topic of interest for decades (e.g. Williams, 1966). In tasks where multiple feature dimensions are available to guide search, correlating where subjects fixate with what features are present at those locations has been the most common means of testing which dimension is most useful (e.g. Hwang *et al.*, 2009; Pomplun, 2006; Rutishauser & Koch, 2007). The current experiments allow a different test of feature importance: More important features should result in larger decrements to performance when those features are changed. A dimension that is more important for guidance should have a larger average and/or maximum mismatch effect, though this requires that the full range of the feature is tested, such as in the current experiments. If only part of the range of values on a dimension is tested, that leaves open the possibility that the maximum mismatch effect occurs in an untested part of the range.

It is hypothesized that, consistent with previous work, search performance will decrease when the preview does not exactly match the target when color information is unavailable (in the grayscale experiments) or when the mismatch is along the hue dimension. This will likely occur not only for the previously-examined orientation dimension, but for shape (aspect ratio) and color as well. Although this has not been previously explored, I predict that mismatch will affect search performance in very global ways, affecting search guidance as well as distractor rejection and target verification. It is also hypothesized that when color stimuli are used (in the "full-color experiments"), the use of hue information to guide search may limit the usefulness of other dimensions, and only limited mismatch effects was found, if at all. This would be consistent with findings that color dominates search (e.g. Hwang *et al.*, 2009).

2.2 Aim 2: Investigate the role and extent of dimension weighting using the preview-target mismatch paradigm.

In odd-one-out search tasks (where no preview is given, and subjects search for "the odd item"), RTs are longer when the feature dimension that differentiates a target from the distractors is unknown (Treisman, 1988). To explain this finding, Müller and colleagues (e.g. Müller, Heller, & Ziegler, 1995; Found & Müller, 1996; Müller, Reimann, & Krummenacher, 2003) proposed a dimension-weighting account in which feature values in saliency maps are amplified by weights that are set based on target knowledge (either from task instruction or from prior trials). When subjects are told in advance what feature dimensions are relevant, weights could be preferentially assigned to those relevant dimensions. When subjects have no knowledge of the relevant dimensions, each dimension is given equal weight. Then, after a trial begins, a process occurs which determines appropriate weighting and assigns that weighting. This process takes some time, resulting in RT increases. These weights are carried over onto following trials, such that the weights are correct if the relevant dimensions remain consistent and a re-assignment of weights is unnecessary. If the relevant dimensions change, RT increases as the weights are reassigned.

Unlike in tasks where the target is defined by a single dimension, when searching for real-world images—which are defined on many dimensions— subjects presumably utilize all relevant dimensions they are able to encode. Unlike singleton feature tasks, subjects may be unable to identify the target using a single dimension: The color brown cannot be used to identify a brown chair target if some of the distractors are also brown. Though dimension weighting has been shown to occur in singleton conjunction search tasks (Weidner & Müller, 2009) in which weight needs to be assigned to two dimensions, and the logic of the theory readily extends to multiple dimensions, this process has not been tested in the context of realistic stimuli and it is possible that weighting may not operate in the same fashion in that context. For example, Weidner & Müller (2009) suggest that dimension-change costs in their tasks may have been caused by a large amount of weight being given to a primary dimension, with a reduced amount of processing of the secondary dimension. It may be the case that this weighting process can only add weight to a single dimension, which would limit the usefulness of the process in the context of realistic objects (although search tasks in which one dimension best discriminates the target from distractors would still benefit).

In the experiments in Aim 1, a single feature dimension was manipulated throughout the course of each experiment. As a result, that feature dimension is unreliable and should be deweighted, minimizing the effect of mismatch on that dimension and maximizing performance. Note that Guided Search (Wolfe 1994) also predicts that feature dimensions are weighted when being combined (though these weights are assigned based on the target definition rather than being reassigned based on each individual search display).

Aim 2 tests, in the context of real-world objects, whether color, orientation, and shape dimensions are deweighted when those dimensions are known to change. Unlike Aim 1, the experiment in Aim 2 manipulated all three dimensions within the same set of experimental trials. This stands in contrast to Vickery et al. (2005), who manipulated only a single feature dimension (orientation). While Bravo and Farid (2009) did vary multiple dimensions, they did not vary whether one or more feature dimensions changed across trials.

If target feature values on a given feature dimension are unknown, top-down biasing cannot occur for that dimension, and for that reason it seems likely that subjects would attempt to either not encode feature values from target dimensions that are likely to be incorrect, or to deweight that feature dimension. In the experiments in Aim 1, if less weight is given to dimensions with uncertain or unreliable feature values, then the feature dimension that is manipulated should be given less weight than it would otherwise receive. For example, if hue is manipulated, hue should be given less weight and orientation and shape should receive relatively greater weight.

In Aim 2, mismatch was created along the same three dimensions manipulated in Aim 1. The experiment in Aim 2 manipulates all three dimensions of mismatch on interleaved trials instead of in separate experiments. Doing so creates uncertainty in which dimension will change on a given trial. If participants can adjust their target templates to give less weight to features that are expected to change, preview-target mismatch effects should be significantly greater when trials with mismatches on different dimensions are interleaved (as in Aim 2) compared to when only a single dimension occurs throughout the course of the experiment (as in Aim 1). As all three feature dimensions were equally likely to mismatch, weightings should be distributed more equally across dimensions rather than a single dimension potentially being largely deweighted. Aim 2 also tests whether these weights are reassigned on every trial—as predicted by Müller et al. (1995)—by comparing trials where the same dimension changed as on the previous trial to trials where a different dimension changed than on the previous trial. If weights are reassigned when different dimensions are relevant, this should result in an RT cost on dimension-change trials. Note that if hue sufficiently dominates search guidance, hue may be largely immune to deweighting and these switching costs might not be observed when color information is available to guide search, even though the presence of hue-mismatch trials might limit the usefulness of that information. To explore switching costs in the absence of hue information, a grayscale experiment was also run, in which only aspect ratio and orientation were manipulated. Again—unlike in Aim 1—both dimensions were interleaved rather than blocked to manipulate uncertainty and allow switching costs to be explored.

2.3 Aim 3: Characterize the summation of mismatches across feature dimensions and compare the empirical summation to the linear summation expected by most models of visual search.

Aim 3 tests an assumption underlying most models of visual search: The assumption that feature dimensions sum linearly (see section 1.3). The top-down comparison process between a target representation and a search display often requires combining features from multiple dimensions, particularly in the case of realistic objects. Many models of search simply assume that these dimensions are combined linearly (e.g. Cave, 1999; Eckstein, Thomas, Palmer, & Shimozaki, 2000; Navalpakkam & Itti, 2005; Palmer et al., 2000; Wolfe, 1994; Wolfe, Cave, & Franzel, 1989), because there is no firm reason within the search literature for modeling this combination differently. Models of search, after all, are not necessarily models of how information is combined across dimensions, and the focus of these models is on explaining other aspects of task performance.

There is some evidence suggesting that feature map combination does not occur through linear summation. Nothdurft (2000) found that subjective saliency appeared to be combined non-linearly across feature dimensions. The feature combinations affected saliency less than expected by linear summation, especially for color and orientation combinations, which Nothdurft suggested may be due to shared mechanisms between the dimensions (resulting in saturation in neurons coding for several dimensions). It is unclear whether these findings extend to how features are combined to create top-down guidance signals, as the saliency computation may involve different mechanisms than those which guide search—saliency could be biased towards making those judgments based on whichever single dimension has the largest difference, for instance, while guidance may not. Huang and Pashler (2005) extended these findings to an objective measure: the distraction (in terms of increases in RT) of a salient distractor in a visual search task. It is unclear from RT, however, whether this distraction affect made locating the target more difficult, or distracted subjects when making their discrimination response. In addition, saliency tends to be overridden by top-down signals in real-world search (Chen & Zelinsky, 2006), so it unclear whether this discounted additivity of features extends to more realistic stimuli.

In addition to simply confirming or disconfirming whether features are combined through linear summation, the specific patterns that characterize feature combination can provide information about the processes involved in search. Changes along multiple dimensions may have less of an impact than single-dimension changes if several dimensions are coded by overlapping sets of neurons. If, for instance, color and orientation are processed in part by the same neurons, those neurons may become saturated, resulting in a lower total activation than was expected by the linear sum of activations from the individual features (Nothdurft, 2000). Note that this may be true even if orientation does not significantly affect search when color information is available: So long as orientation is processed by the same neurons as color, some saturation may occur.

Conversely, changes along multiple dimensions may have *more* of an impact than changes on a single dimension. Neider and Zelinsky (2008) demonstrated that objects which are sufficiently dissimilar to the target are not considered as potential distractors. Objects which are target-dissimilar on multiple dimensions may be more likely to be excluded from the set of distractors than objects which are dissimilar on only a similar dimension. Similarly, participants may be quicker to (or more likely to) reject objects that are target-dissimilar on multiple dimensions. This would result in greater mismatch effects on trials with multiple dimensions changed. Note that this may be true whether or not color is the feature which primarily drives search guidance: Orientation and shape may contribute to objects being included in the set of considered distractors even if effects on search guidance and verification are expressed only in color.

In Aim 3, "combined mismatch" trials were included in which pairs of the dimensions—or all three dimensions—were manipulated on the same trial. These combined mismatch trials were compared to trials in which only a single dimension is altered to test how the feature maps are combined. The same increments of feature mismatch in terms of each individual feature were used for single-dimension change trials and combined change trials, to allow for direct comparisons. The mismatch effect on RT (or other measures) on combined mismatch trials should be the combination of the individual RT effects on single-dimension mismatch trials. If, for instance, changing hue angle 15° increases RT by 100 ms and 180° of orientation change increases RT by 100 ms, when both 15° of hue angle and 180° of orientation are changed on a trial, RT should increase by 200 ms if the effects on guidance sum linearly. This is true even though different feature dimensions are probably extracted or encoded using different mechanisms which do not use the same dynamic range and different scales. Aim 3 includes two experiments: A grayscale experiment in which only aspect ratio and orientation were manipulated, and a full-color experiment in which all three dimensions were manipulated so as to see whether orientation and aspect ratio combine differently in the absence of color.

Chapter 3 - General Research Design and Methods

3.1 Participant recruitment and restrictions

Twenty-five different subjects participated in each experiment for Aim 1 and 2 (experiments 1, 2, 3, and 4). Twenty-two subjects participated in Experiment 5. No subject participated in more than one experiment. For Experiments 1, 3, 4, and 5, the same numbers of subjects were also run in grayscale versions of the same experiments. The number of subjects was determined based on counterbalance: Each search display image was counterbalanced across every preview-target mismatch condition (further discussed in section 3.4). As a result, the grayscale and full-color versions of Experiment 5 had fewer subjects than Experiments 1-4, due to a reduced number of within-subjects conditions (discussed in section 4.3.1.1). Subjects were recruited from the undergraduate Psychology subject pool at Stony Brook University. Subjects all reported normal or corrected-to-normal vision and normal color vision.

3.2 Equipment/Apparatus

Gaze position was recorded using an SR Research EyeLink 1000 eye tracking system with default saccade detection settings. This video-based eye-tracker has a sampling rate of 1000 Hz and is typically accurate within 0.25-0.5°. Initial calibrations were not accepted unless the average spatial error was less than 0.45° and the maximum error was less than 0.90°. Head position and viewing distance were fixed at 70 cm from the screen with a chin rest. The experiment was displayed on a flat-screen CRT monitor at a resolution of 1280 x 960 pixels (subtending 28.7° x 22.3°) using a refresh rate of 60 Hz. Manual responses were made using a GamePad controller connected to a USB port.

3.3 Stimuli generation

Target and distractor images were hand-selected from the Hemera Photo-Objects collection (volumes 1, 2, and 3), and were from a wide a variety of categories. Several constraints were placed on image selection. Images that contain little to no hue information (images that are mostly black, white, or gray, such as steel kitchen utensils or soccer balls) prevent a meaningful exploration of changes in hue and were therefore not be used. Specifically, images were required to have 80% or more of their pixels differing by at least 3.9% of the possible range of RGB values (a difference of at least 10) on either the red, blue, or green channel. For example, a given pixel having RGB values of 249, 250, 255 would not be acceptable but a pixel having RGB values of 244, 250, 255 would be accepted. Excessively tall or wide images—those having a width that is more than 150% of the height or a height that is more than 150% of the width—were excluded to prevent extreme or inaccurate changes when aspect ratio is altered and to prevent stimuli from extending off the screen or occluding other objects in the display. Also to prevent compromising the aspect ratio manipulation, images that appeared oriented diagonally in the image canvas were excluded. Round images (such as balls, pizzas, and plates viewed from above) were also not included, as those images do not have a global orientation. Note that these constraints were applied to both target and distractor images. Finally, no distractor was allowed to share the target category; if the target on a given trial was a pig, none of the distractors on that trial were pigs. As the task required subjects to search for objects that were not identical to the preview, including same-category distractors might conceivably result in subjects becoming confused about the task and misidentifying those objects as targets.

A script, programmed in MATLAB (version 7.0.4), was used to implement the various feature changes to these images that comprised the experimental conditions. The amount of feature change performed on each image was counterbalanced across subjects. After feature changes were made, images were rescaled so that all of the images had 6000 non-white pixels (or as close as possible) in order to normalize for size. This rescaling was done after the feature change to accommodate the cropping required by the orientation feature manipulation (section 3.3.1) and so that changing the aspect ratio could be accomplished without upsampling (section 3.3.3). On average, the resulting images were ~1.77° of visual angle in size. The specifics of each feature manipulation are discussed below, in the following three sections.

3.3.1 Orientation

Orientation was manipulated using MATLAB to rotate the matrix of pixel values, without interpolation, by the required number of degrees. The script then cropped any rows or columns of pixels containing no non-white values from the edges of the image. Angles of rotation are indicated in the counter-clockwise direction, with 0° as the starting value. See Figure 1 for example orientation-changed stimuli.

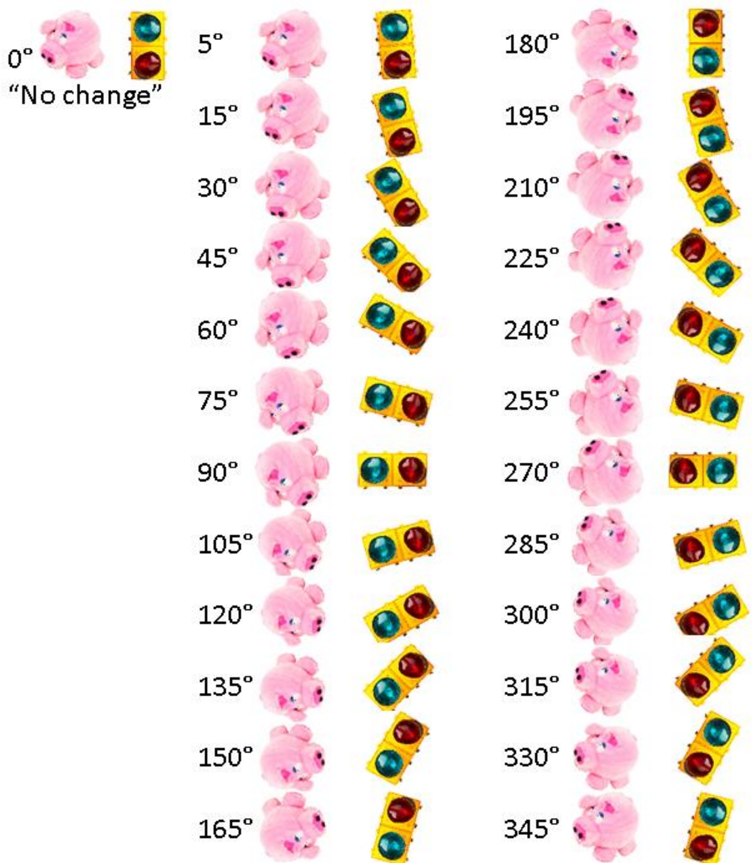

Figure 1. Example orientation-changed stimuli at each increment of change. Aim 1 used all 25 levels of rotation depicted here. Aim 2 selectively used 0°, 30°, 60°, 90°, 120°, 150°, 180°, 210°, and 240° (see section 4.2.1.1). Aim 3 used 0°, 60°, 120°, and 180° (see section 4.3.1.1).

3.3.2 Hue

The hue of images was altered by converting each image from RGB color space into HSV color space, adding a pre-defined value to the hue angle of each pixel, and then converting the image back to RGB color space. HSV is a cylindrical-coordinate model of color, treating three aspects of color (Hue, Saturation, and Value) as separate, independent dimensions. The hue, saturation, and value dimensions are confounded in RGB space, and the hue manipulation was performed in HSV space in order for us to test hue independently of saturation and value. In HSV space, saturation is represented as the radius of the cylinder (the distance from the central vertical axis), value as the height of the cylinder (the distance along the central vertical axis), and hue as the angle around the central vertical axis (see Figure 2).

Changing the hue angle changes the color of an image (red to blue, for instance) but does not change the saturation or the value of the image. By breaking color down into these three components, these individual features comprising what we normally think of as color can be explored independently, with one caveat: There are some instances in which several perceptual features vary with change in a single dimension in HSV space. Though HSV space separates “hue” from “saturation” and “value,” distances in HSV space do not perfectly capture distances in the *perception* of hue, saturation, or brightness. For example, distance in hue angle in HSV space does not perfectly correspond to perceptual similarity distances based on dominant wavelength of light leaving the image. Blue and yellow are not the same distance from each other in HSV space as red and green, though both are pairs of perceptual opposites. The most obvious difference between HSV space and perceptual space, however, is in terms of value: The white pixels in the center of Figure 2 are treated as having the same value as the more saturated pixels on the circle’s circumference, which clearly does not map onto human perception. As a result, changes in saturation can also change perceived brightness. While changes in hue angle can also change perceived brightness (e.g. dark blue is considered to have the same value as yellow), this is much less pronounced than for changes in saturation. Note that unlike changes in saturation, changes in hue angle do not vary *systematically* with perceived brightness: A change in 20° of hue angle could make the image appear either brighter or darker. Additionally, this relationship between hue angle and perceived brightness is less true for images that are low in saturation (e.g. all off-white colors appear bright) or value (e.g. images with a value of 0 all appear dark). For these reasons—and because most prior research on the effects of color on visual search (see 4.1.2.1) have focused on hue—the color manipulation in the current experiments alters hue angle rather than saturation or value.

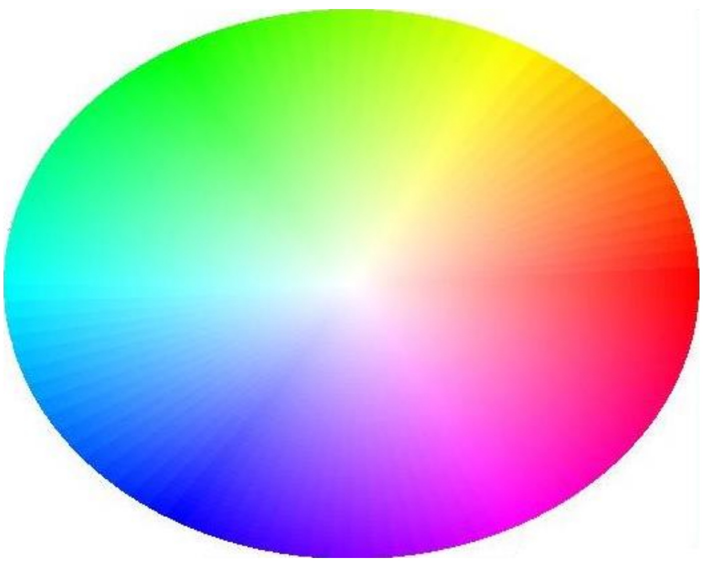

Figure 2. A visualization of the hue and saturation dimensions of HSV color space. Hue is represented as the angle around the circle and saturation is the distance from the center. Value (white-to-black) would extend in the z-axis; the value of this cross section is maximally bright (a value of 1).

Hue was manipulated by using MATLAB to change the hue angle in HSV space for each pixel that has hue information. Black, white, and gray pixels that carry no hue information (i.e. that have exactly the same amount of blue, red, and green) were not altered. Hue angles in this paper are given in the counter-clockwise direction, with 0° as the starting value. See Figure 3 for an example of a hue-changed stimulus.

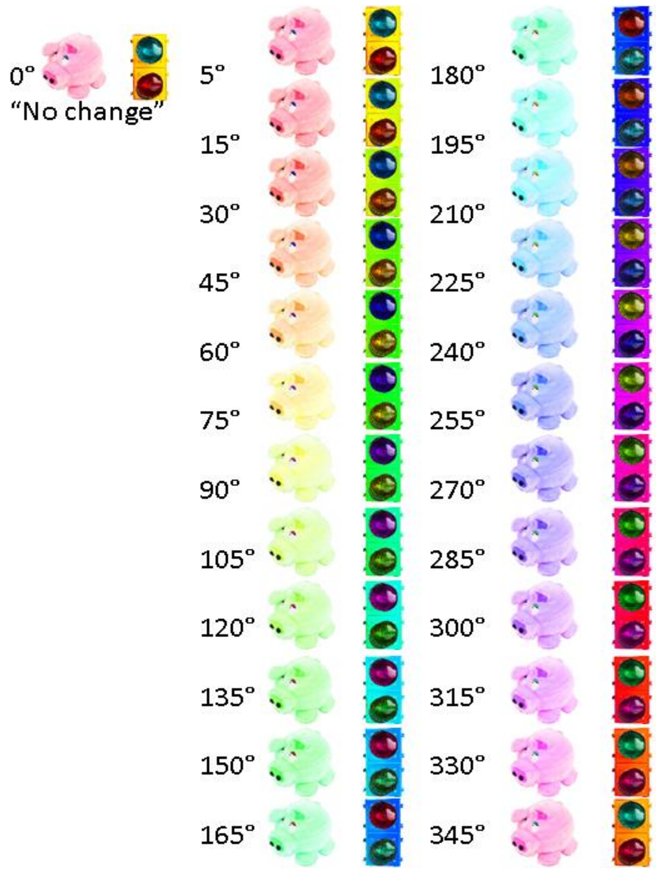

Figure 3. Example of a hue-changed stimulus at each increment of change. Aim 1 used all 25 levels of hue angles depicted here. Aim 2 selectively used 0°, 30°, 60°, 90°, 120°, 150°, 180°, 210°, and 240°. Aim 3 selectively used 0°, 60°, 120°, and 180°. Note that black, white, and grey grayscale pixels are unaffected by the manipulation.

3.3.3 Shape

The term “shape” includes many aspects of images: curvature, length, width, global shape (square, triangle, etc.), the number of corners or edges, any geometrical description of an object, and so on. In Experiment 3, aspect ratio was manipulated as a means of altering shape. The reasons for selecting aspect ratio, rather than a different shape feature, are straightforward. First, the process of altering the aspect ratio of real-world objects is very simple, straightforward, and thus easily replicable. Most image editing software allows aspect ratio to be changed, while altering the number of edges in an image in a systematic way is not so simple (aside from occluding some edges, or deleting parts of the object, which is arguably not a shape manipulation to the object). Second, there is some suggestion that aspect ratio may be preattentively available to guide search (Treisman and Gormican, 1988), thereby elevating the importance of this feature. Third, altering the aspect ratio of images alters a number of simpler shape features (length, width, curvature, and the angles of any corners). One goal of Aim 1 is to examine whether shape information is used to guide search to real-world objects. Encompassing a number of shape features in the manipulation is desirable as it allows the inclusion of many objects that would need to be excluded if these shape features were used individually (objects without corners can obviously not be used if a "number of corners" shape feature was to be manipulated). Altering aspect ratio, which is a more global shape feature, also allows for the shape of the entire object to be changed rather than a small part of the object (which might not even be visible in the periphery).

Aspect ratio was manipulated by using MATLAB to rescale (through downsampling) the height of each image. This manipulation was performed before images were resized to normalize the visual angle between images (see section 3.3). The end result of the process is that the *width* of the image was increased and the height decreased. Because all of the images began larger than the desired visual angle, the resizing process downsampled the images to the required pixel area. In this way, no upsampling was required on the images. As upsampling requires the interpolation of new pixels, spurious features could be created. While it is true that downsampling also could create spurious features (such as creating new edges on "blocky pixels" where curves could not be captured with fewer pixels), these images need to be further downsampled regardless to ensure that they share the same visual angle. Downsampling at both stages, rather than both upsampling and then downsampling, minimizes image distortion. See Figure 4 for example shape-changed stimuli.

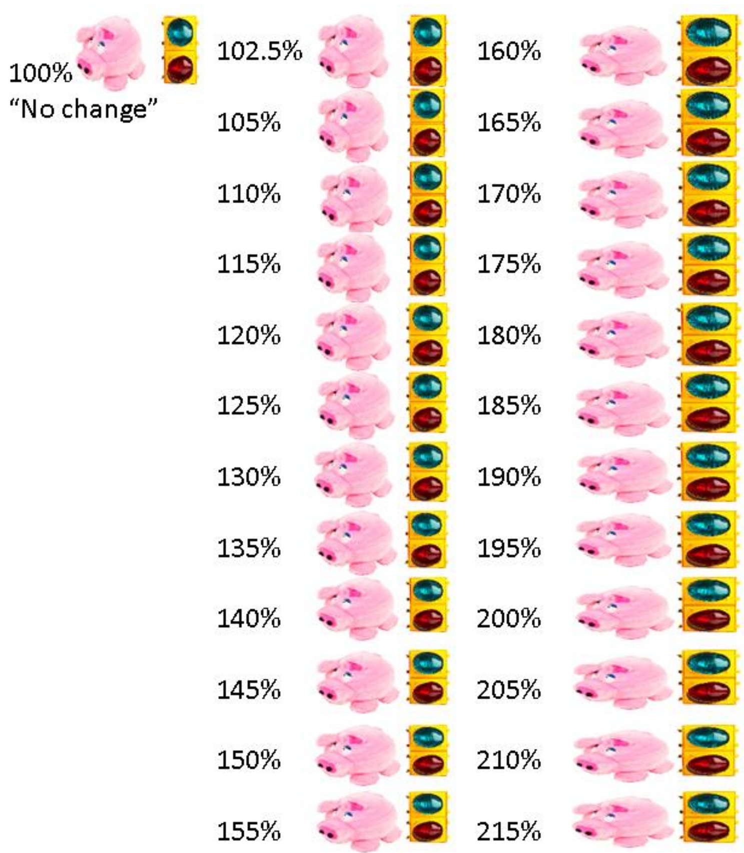

Figure 4. Example of a shape-changed stimulus at each increment of change. Aim 1 used all 25 levels of aspect ratio changes depicted here. Aim 2 selectively used 100%, 110%, 120%, 130%, 140%, 150%, 160%, 170%, and 180%. Aim 3 selectively used 100%, 120%, 140%, and 160%.

3.4 General procedure

Figure 5 shows a graphic depiction of the procedure. The subject fixated a central dot and pressed a button on the controller to initiate each trial. A preview of the target then appeared in the center of the screen and remained visible for 200 ms. The search display appeared, consisting of eight evenly spaced objects (the target and seven distractor objects) arranged on an imaginary circle with a radius of 275 pixels (~6.4º) relative to the center of the screen. The target was always one of the objects in the search display. Subjects were instructed to fixate the target and—while looking at the target—press a trigger button on the controller with the index finger of their dominant hand. Subjects were told that the target would always be present, but might look different from how it appeared during preview. Note that feature-manipulated targets were still considered to be targets.

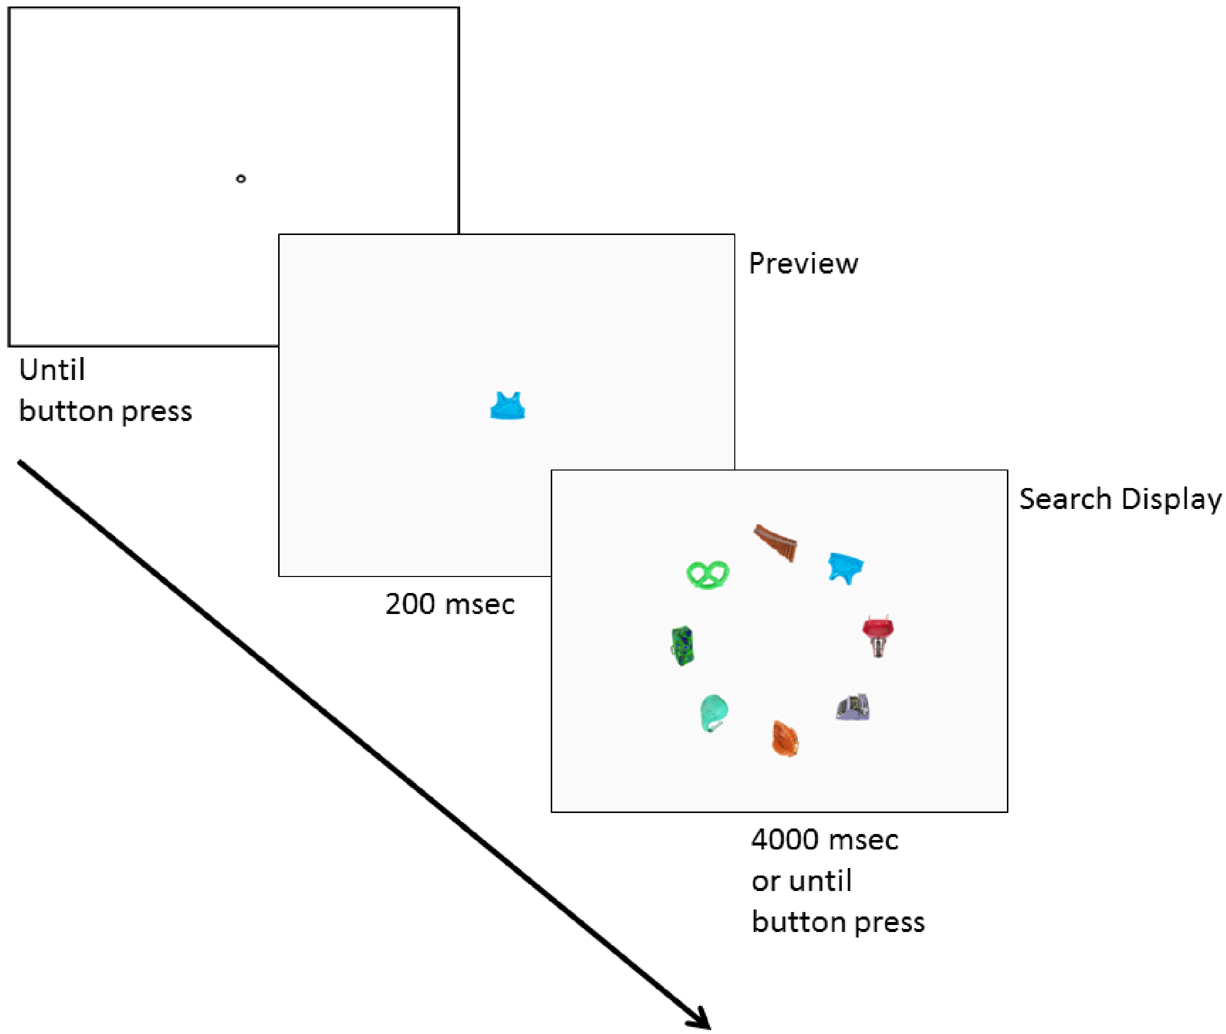

Figure 5. Example trial sequence for an orientation-change trial (Experiment 1).

To control for any potential artifacts caused by feature manipulation that could bias search behaviors, feature manipulations were performed on each of the seven distractor objects in the search displays. Specifically, one object had its hue changed, the second had its shape changed, the third had its orientation changed, the fourth had both hue and shape changed ("hue*shape"), the fifth had both hue and orientation changed ("hue*orientation"), the sixth had both shape and orientation changed ("shape*orientation"), and the seventh had all three features changed ("hue*shape*orientation"). This means that when the target was manipulated on a given dimension, more than half of the distractors also had been altered on the same dimension. This manipulation was meant to discourage subjects from creating unusual search strategies such as looking for objects with abnormal color. The particular amount of change that was applied to each distractor was randomly determined on every trial. For hue and orientation, the random value of the change ranged between 1-359°, and was restricted to whole numbers/degrees. For aspect ratio, the value ranged between 101%-215%, and was restricted to whole numbers.

A single set of trials was created, with trial-unique targets and distractors. Targets and distractors were selected randomly from the stimulus set. Distractors that shared the same category as the target on a given trial were replaced with another random distractor to prevent target-distractor confusion (see 3.3). Stimuli repeated five times throughout the list of trials, but in different contexts each time (different target/distractor and in different mismatch conditions). The same search displays were used for each experiment. Manipulations performed on the target on each trial, however, depended on the particular experiment. In other words, subject 1 in Experiment 2 had the same image pairings as subject 1 in Experiment 1, but target hue was altered in Experiment 2 and target orientation was altered in Experiment 1. In this way, all factors are identical between experiments except the feature manipulations performed on the target images. The preview-target mismatch conditions were counterbalanced across subjects, such that each trial appeared in a different mismatch condition for each subject. For example, the first trial on the list in Experiment 1 was a no change trial for subject 1, a 5° change for subject 2, a 15° change for subject 3, and so on. The change applied to distractor objects was randomized for every subject.

Interest areas—invisible circles with a radius of 100 pixels (2.3º) around the center point of objects in the search displays—were created in Experiment Builder. Fixations that fell within these interest areas were considered fixations on that object. If subjects pressed the response button while fixating within 100 pixels of the center of the target, the trial was counted as correct and terminated without feedback. If a subject pressed the response button while fixating a distractor object, an error noise sounded and "INCORRECT" appeared on the screen for 500 ms. Trials timed out after four seconds, at which point the error noise sounded and "TIMED OUT" appeared on the screen for 500 ms. If subjects pressed the response button while not fixating any of the items in the search display, an error noise sounded but the trial did not terminate. Subjects were instructed that, should this occur, they should look at the target and press the button again, and then inform the experimenter. If this occurred on two trials in a row, the experimenter interrupted the task and recalibrated the eyetracker, in case the subject was fixating the target but the calibration had become poor and the eyetracker was not correctly recording the eye position as being on the target object. A break was given halfway through the experimental trials and preview-target mismatch conditions were interleaved throughout the experiment. Upon completion of the experiment, subjects were debriefed as to the purpose of the experiment.

Chapter 4 - Details and results for each experiment

4.1 Experiments for Aim 1

In Aim 1, the experiments investigated the effects of mismatch between previews and search targets on three specific feature dimensions: Orientation, hue, and shape. Five experiments were conducted. Two versions of Experiment 1 (one grayscale version and one full-color version) manipulated mismatch on the orientation dimension. Experiment 2 manipulated mismatch in hue. No grayscale version of Experiment 2 was included, for obvious reasons. Two versions of Experiment 3 (one grayscale version and one full-color version) manipulated aspect ratio. These experiments were identical in all other respects.

4.1.1 Experiment 1: Orientation Mismatch

4.1.1.1 Background information regarding orientation as a guiding feature

Orientation has been shown to meet the criteria previously used to identify guiding features. Search can be efficient to orientation-defined targets, and oriented targets can “effortlessly” pop-out (Foster & Ward, 1991). Orientation can “effortlessly” be segmented from a texture (Nothdurft, 1991). Search asymmetries have been found for orientation, with search being easier for tilted line segment targets among vertical (0°-oriented) distractors than for the reverse (Foster & Ward, 1991; Poirier & Gurnsey, 1998; Treisman & Souther, 1985). Orientation asymmetries in singleton search tasks with simple line stimuli have also been found in the direction of initial saccades (Foster, Savage, Mannan, & Ruddock, 2000). Orientation information was even found to correlate with search performance in realistic scenes (Pomplun, 2006).

There are some reasons to believe that orientation may operate differently than other feature dimensions. Most previous work on orientation as a possible guiding feature have relied on stimuli that cannot be rotated a full 360° (stimuli that do have a full range have often been used, but typically not in studies where the effects of orientation were parameterized). For example, line segments rotated 180° are identical to unrotated line segments. Wolfe, Klempen, and Shulman (1999) used a variety of relatively simple stimuli (such as letters), including stimuli that had obvious axes of orientation, stimuli which did not have obvious axes but were polar objects (objects that, when rotated 180°, appear to be upside-down), and objects that had obvious axes and were polar. Search for objects rotated 180° was inefficient even for objects with strong polarity, and search for objects rotated 90° from distractors was more efficient than search for objects differing in 180° from distractors, suggesting that search is guided based on 180° representations, and that guidance does not use a full 360° range. As a result, an upside-down object is treated by guidance mechanisms the same way as a right-side-up object, and an object rotated 45° as equivalent to one rotated 315°. The authors noted, however, that this may not be true for real-world objects, for which features such as lighting direction may allow their directionality to be better processed preattentively. If this previously reported pattern holds in the current experiments, the data should show an increasing mismatch effect up to 90°, followed by a decrease as mismatch increases from 90° to 180°. This prediction is also consistent with the data of Vickery et al. (2005), who found an effect of mismatch on RTs that increased up to 90° in their task. However, since they did not use eye movement measures, it is unclear whether their data pattern was due to guidance effects or due to small rotations causing difficulty in target verification or other processes (for instance, subjects may have mentally rotated the images for the purpose of recognition in trials where orientation change was below 90° and may have used other strategies with higher amounts of change).

In Experiment 1, the orientation of the target was varied relative to the preview, and was designed to systematically evaluate the effect of preview-target orientation mismatch on search guidance and verification, with and without the availability of hue information. This experiment replicates and extends the work of Vickery et al. (2005) and Bravo and Farid (2009) to include more increments of mismatch and a full-color/grayscale manipulation. Experiment 1 included twenty-five conditions of change, where Vickery et al. (2005) used five (Experiment 3) and seven (Experiment 5). The additional conditions of the current experiment allow for a fine-grained analysis of small orientation changes across the full range of the feature dimension, and allows for testing whether search relies on a 180° or a 360° frame of reference (further discussed below in section 4.1.2.3).

Methods

4.1.1.2 Stimuli/procedure specifications

Twenty-five within-subject orientation conditions were created by manipulating the targets in the search displays (using the methods described in 3.3.1) in 15° increments from 0° to 360°. In addition, a 5° mismatch condition was included, to explore a finer-grained amount of mismatch. The conditions were: 0° ("no change"), 5°, 15°, 30°, 45°, 60°, 75°, 90°, 105°, 120°, 135°, 150°, 165°, 180°, 195°, 210°, 225°, 240°, 255°, 270°, 285°, 300°, 315°, 330°, and 345° of change. Each search display was counterbalanced across all 25 conditions across subjects. Subjects performed 25 practice trials (one for each mismatch condition). Twenty-five experimental trials were included for each condition, for a total of 650 experimental trials. In all other respects the procedure was the same as the general procedure (section 3.4). The grayscale version was identical to the full-color version in all respects save that the stimuli were grayscale and hue-change conditions were omitted. As such, there were only 17 practice trials and 425 experimental trials.

4.1.1.3 Results and discussion for Experiment 1

Orientation mismatch had no reliable effect on any measure (RTs, accuracy, time-to-target, target fixated first, and verification time) in the full-color version of Experiment 1 (all *p*s > .20). When hue was available to guide search, participants did not use orientation either to guide their search or to identify whether the currently fixated object was the target or a distractor. Note that although previous work did find mismatch effects with orientation, the use of hue information in those tasks was either likely minimized by the use of highly colorful backgrounds (Bravo & Farid, 2009) or was completely unavailable through the use of grayscale stimuli

(Vickery et al., 2005). In addition, because the mismatch conditions in Bravo and Farid (2009) confounded orientation change with changes in size and reflection across the vertical axis, it is unclear which changes resulted in the mismatch effect reported in that paper. As such, the grayscale version of the current experiment (detailed below) is more comparable to previous work. That no significant effects of orientation mismatch were found in the color version of this experiment suggests, however, that orientation is either not used—or is only minimally used—to guide search when hue information is available. Since most real-world objects have hue information, this would suggest that orientation does not normally contribute significantly to search in the real-world (but is available to the search process when necessary, as when searching for objects without hue information, or on very colorful backgrounds). This finding is consistent with prior work suggesting that search is predominantly driven by color (e.g. Williams, 1966; Pomplun, 2006).

The results below all refer only to the grayscale version of the experiment. All ANOVAs in this section are one-way within-subjects ANOVAs comparing the specified dependent measure to the amount of orientation mismatch, except where otherwise specified.

Manual reaction times and accuracy

To the extent that subjects rely on orientation in order to perform the task, RT was expected to differ across orientation mismatch conditions. Unlike in the full-color version of Experiment 1, this is precisely what was found: RTs significantly increased as orientation mismatch increased, $F(24,576) = 5.18$, $p < .001$. Note that this function generally followed the expected u-shaped pattern, with RTs increasing as orientation change increased to around 180°, followed by decreasing RTs as orientation changed in the direction of 360° (decreasing mismatch, as a result of the circular nature of orientation)—Figure 6A. Orientation mismatch did not reliably affect accuracy—$F(24,576) = .67$, $p = .76$—indicating that the effect in RTs was not due to a speed-accuracy trade-off. Error trials (including timed-out trials) were removed from all other analyses (including the RT analysis).

As discussed in section 2.1, effects of preview-target mismatch on RTs have previously been found, and so are not unexpected. However, the inclusion of numerous smaller increments of mismatch in the present experiment allowed mismatch effects to be tested for linearity. To account for the circularity of the orientation feature space when testing for linearity, trials were averaged across clockwise and counterclockwise rotation (e.g. 15° and 345° were combined)—see Figure 6B. No change trials, 5°, and 180° were omitted from the analysis, as they had no clockwise/counterclockwise matching conditions). Both linear and quadratic contrasts were significant, $F(1,24)>24$, $p < .001$, indicating that RT generally increases with increasing mismatch in orientation, demonstrating that orientation was used in performing this task, and that this increase levels off at higher amounts of mismatch. This quadratic component may reflect an ability to reject an object as a potential target when it is sufficiently dissimilar to the target preview: A 45° orientation mismatch may still be treated as a target-similar orientation, but once the orientation mismatch is greater than ~60° subjects may be considering the orientation to simply be "incorrect." At large amounts of mismatch, distractors are more likely to be more target-like in orientation, which may play a role in this pattern (sufficiently target-dissimilar orientations may be rejected as too dissimilar to the target simply because another object in the display has become more target-like in orientation than the target itself).

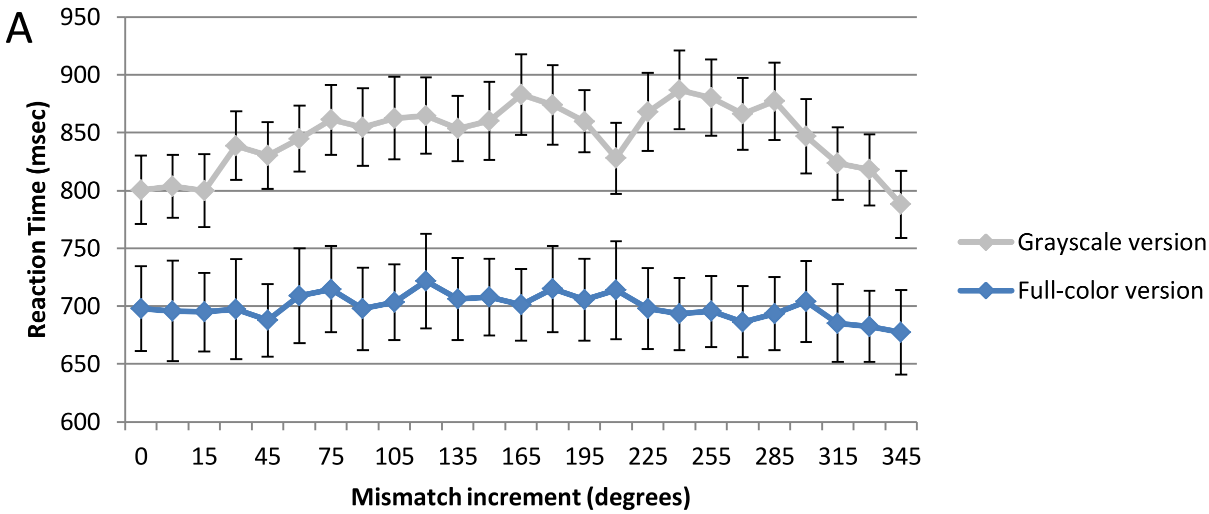

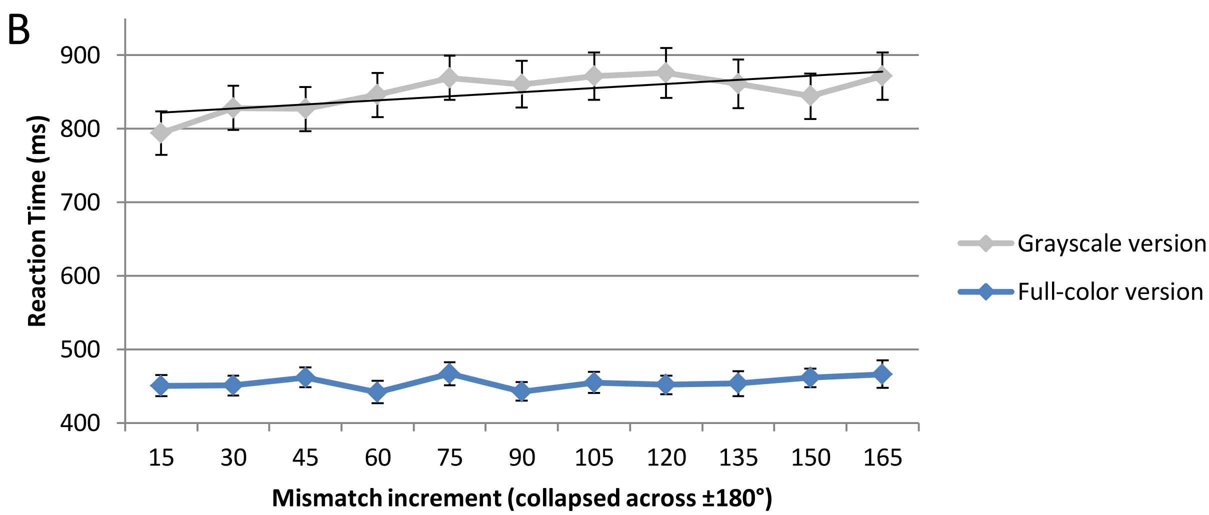

Figure 6. Change in RTs across orientation mismatch conditions in Experiment 1, for each mismatch condition (A) and collapsing across clockwise and counterclockwise change in orientation, with a best-fit line indicating the linear component of the increase of RTs with increasing mismatch in the grayscale version of the experiment (B). Error bars indicate one standard error of the mean.

Because manual RTs could be affected by a variety of factors other than search guidance, eye movement measures were used to break this coarse analysis down into two finer-grained measures: Time-to-target and verification times (e.g. Castelhano, Pollatsek, & Cave, 2008; Maxfield & Zelinsky, 2012; Schmidt & Zelinsky, 2009). Time-to-target (the time from display onset to when subjects first fixate on the target) is primarily driven by guidance processes. Verification times (the amount of time it takes subjects to make their response after fixating the target) are primarily driven by recognition processes, and should be relatively unaffected by guidance to the target (as the target has already been fixated by this point).

Guidance

Time-to-target changed with orientation mismatch—$F(24,576) = 2.64$, $p < .01$—suggesting that—in the absence of hue information, and consistent with previous work—orientation is used in guiding search to the target. Once again, both linear and quadratic contrasts were significant—$F(1,24) > 9$, $p < .05$. These data suggest that as orientation mismatch increased, gaze was guided less efficiently to the target, suggesting that target orientation was being coded at preview and used in search—see Figure 7. Moreover, the quadratic component suggests that small amounts of mismatch had relatively large effects but that the size of the mismatch effects leveled off at larger amounts of mismatch, potentially due to subjects treating large amounts of mismatch as "incorrect" and fully target-dissimilar regardless of the exact amount of mismatch. The presence of this quadratic component not only in RTs but in time-to-target suggests that this pattern begins early in the task and affects the guidance system

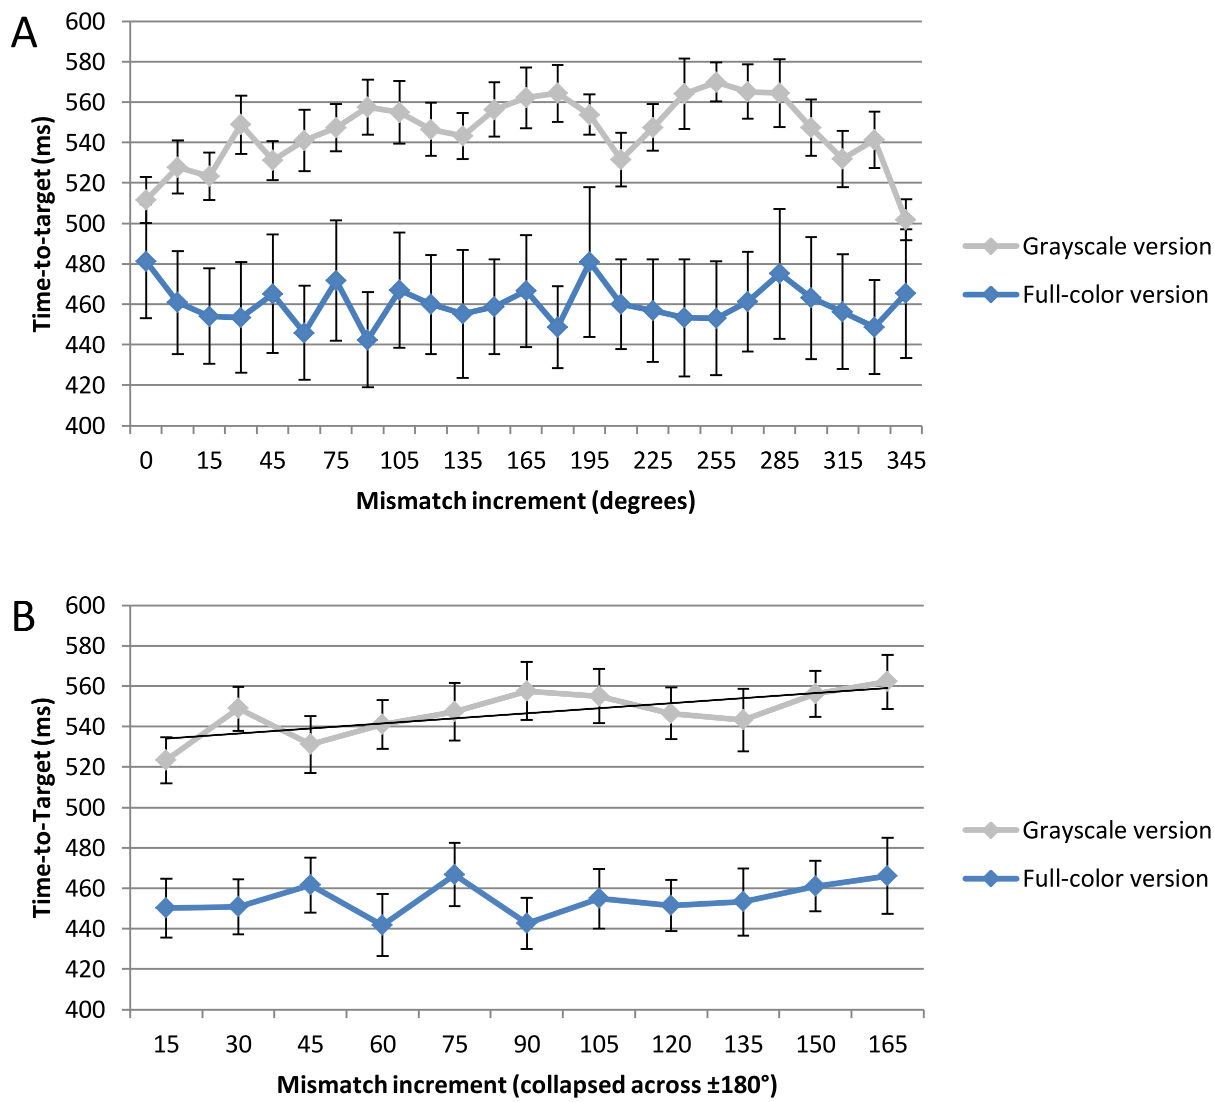

Figure 7. Change in time-to-target across orientation mismatch conditions in the grayscale version of Experiment 1, for each mismatch condition (A) and collapsing across clockwise and counterclockwise, with a best-fit line indicating the linear component of increases in time-to-target with increasing mismatch in the grayscale version of the experiment (B). Error bars indicate one standard error of the mean.

Because time-to-target could be affected by factors unrelated to guidance (such as the time taken to reject distractors before target fixation), a more conservative measure of early search guidance was used: The percentage of trials where the target was the first object fixated. This measure was above chance levels for all mismatch conditions—all $t(24)>9.47, p < .001$—demonstrating that eye movements were guided early in the task. To test whether orientation mismatch affected this early guidance, the proportion of target first fixations was then compared across mismatch conditions. If observers first fixate the target less frequently as the amount of orientation mismatch increases, then this is evidence that early search guidance processes rely at least partly on orientation. Note that the inclusion of this measure provides a means of examining the time course of guidance effects in these experiments: An effect in time-to-target but not in target first fixated could indicate a feature mismatch effect that does not affect early guidance, but begins to affect performance only after the first eye movements are made. Consistent with this, orientation mismatch did *not* reliably affect the proportion of trials where the target was first fixated—$F(24,576) = 1.34$, $p = .20$, see Figure 8—suggesting that any effect of orientation on guidance takes place *after* the first eye movements are made.

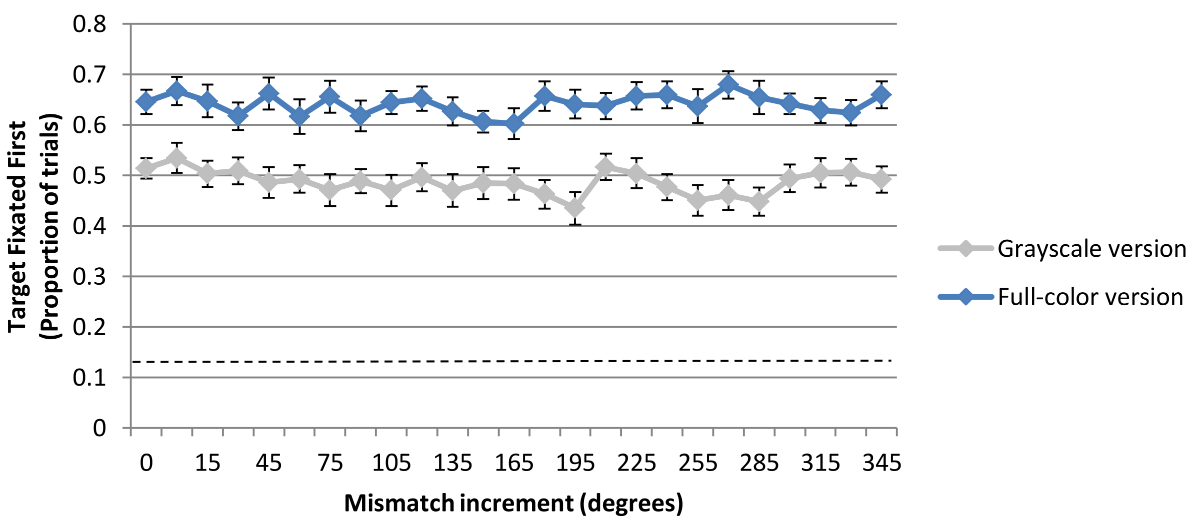

Figure 8. Change in target first fixated across orientation mismatch conditions in Experiment 1. The dashed line indicates chance levels of performance (0.125). Error bars indicate one standard error of the mean.

Verification times

Verification times were significantly affected by orientation mismatch, $F(24,576) = 3.31$, $p < .001$, see Figure 9, indicating that orientation was used in responding to the target in this task. As in RTs, verification times increased up to 180° of mismatch, and then decreased with increasing mismatch, consistent with the circularity of the orientation feature space.

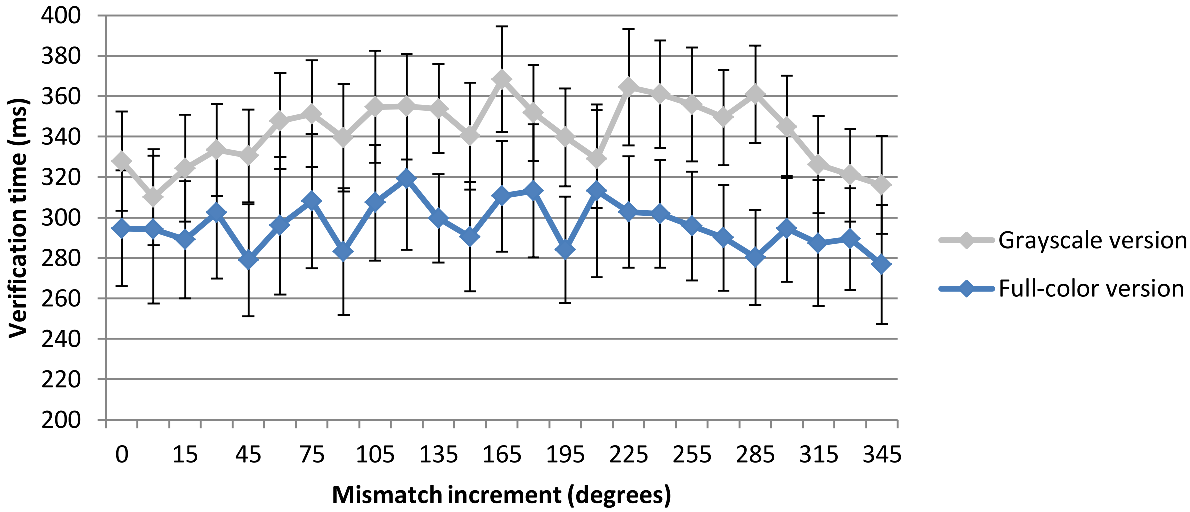

Figure 9. Change in verification time across orientation mismatch conditions in Experiment 1. Error bars indicate one standard error of the mean.

4.1.1.4 Analyses specific to Experiment 1

In addition to testing the general predictions for Aim 1 (section 2.1), several more specific analyses were conducted for the orientation manipulation in Experiment 1. In the real world, when presented with an object viewed from a perspective that is so unfamiliar as to render it difficult to recognize, one effective strategy for identifying the object is to pick it up and rotate it until it can be recognized. Similarly, when presented with an image of an object, subjects might "mentally rotate" the object to facilitate recognition. This is particularly likely when comparisons are required, such as the comparison between the preview and the object in the search displays of the current experiment. If an image looks like the preview, but the orientation is clearly wrong, subjects may mentally rotate the target image to the previewed orientation in order to confirm that it is, in fact, the target. Previous work has shown that this kind of mental rotation may be easier at certain special degrees of rotation. Specifically, Jolicoeur (1985) found that RTs are relatively short at 180° of rotation. Murray (1997) has suggested that this specific speeding of 180° of rotation may be due to an ability of subjects to "flip" objects across the vertical axis, rather than mentally rotating the objects as subjects do at other rotations.

Regardless of the reason for the reported speeded 180° RTs, shorter RTs might be expected here at 180° of mismatch. However, no difference was found between the 165° condition and the 180° condition (p = .70), indicating that 180° did not allow for a special process to speed search by flipping the images across the vertical axis in the present study. While it is possible that this null finding is due to a lack of power, the 9ms difference in RT here are far short of the >100ms differences reported elsewhere (e.g. Murray, 1997). It is possible that speeded 180° RTs are task-dependent and do not emerge in the context of visual search (instead emerging in picture naming or recognition tasks, such as those in Murray, 1997). It is also possible that the complexity of the real-world objects used in this study limits either the benefit of this process (resulting in too small an improvement to be detected) or prevents the use of this strategy.

Previous work (Wolfe, Klempen, & Shulman, 1998) demonstrated that, at least for simple stimuli, the guidance system relies on a 180° frame of reference, rather than a 360° frame of reference. As a result, right-side up people and upside-down people would both be treated as "vertical"—their orientation identical—by the guidance system. The authors suggested that this finding might not extend to real-world objects (noting, for example, that unlike the stimuli they used, real-world objects typically have an unambiguously-defined orientation in a 360° reference frame) and, in fact, evidence for a 180° reference frame was not found in the present experiment: Targets in the 180° mismatch condition in the current experiments had worse (rather than equal) guidance to the no change condition, by ~52.7ms. If subjects did use a 180° frame of reference, however, we should not only find equal guidance between those two conditions, but also when comparing 240° and 60° (a 23.2ms difference in the present experiment, indicating again that subjects did not use a 180° reference frame). The current results suggest instead that a 180° frame of reference is not used for guiding search for real-world objects. This is consistent with findings from previous work (Vickery et al., 2005), in which 180° had a mismatch effect relative to 0°, though their measures did not distinguish between guidance and verification processes. The current results provide converging evidence for the use of a 360° reference frame in visual search for real-world objects, as well as extending previous work to demonstrate that this 360° reference frame is used both in guiding search and in verifying the target once it is fixated.

4.1.2 Experiment 2: Hue Mismatch

4.1.2.1 Background information regarding color as a guiding feature

Color was identified early on as a feature that guides search (Williams, 1966), and color meets all of the criteria previously used for identifying guiding features (Wolfe & Horowitz, 2004). Numerous experiments have demonstrated that search for a color-defined target can be efficient (Carter, 1982; Duncan, 1989; D'Zmura, 1991; Nagy, Sanchez & Hughes, 1990), and color can effortlessly pop-out of displays (Bauer, Jolicoeur & Cowan, 1996; D'Zmura, 1991; Treisman & Gelade, 1980). Although color search is sometimes not efficient when distractors are heterogeneous, efficient search can be found even with heterogeneous distractors (Duncan, 1989; Smallman & Boynton, 1990; Wolfe, Yu, Stewart, Shorter, Friedman-Hill, & Cave, 1990). The particular conditions in which search for a color target among heterogeneous distractors is efficient appear to be when the colors are far apart in color space (Duncan, 1989; Smallman & Boynton, 1990; Wolfe et al., 1990), or when they are linearly separable in color space (Bauer, Jolicoeur & Cowan, 1996; D'Zmura, 1991). Search asymmetries involving color have also been found (Treisman & Souther, 1985; Treisman and Gormican, 1988). While some have suggested that color affects visual search in terms of color categories (e.g. Daoutis, Pilling, & Davies, 2006; Wolfe, 1994; Yokoi & Uchikawa, 2005), more recent evidence suggests that it is the perceived similarity between color values, not the category of colors per se, that matters (Reijnen, Wallach, Stöcklin, Kassuba, & Opwis, 2007; Vighneshvel & Arun, 2013).

Previous work on color as a guiding feature has primarily relied on stimuli with homogenous colors. It is unclear how much guidance relies on color for real-world stimuli, which contain many colors. When objects are viewed in the periphery (as in a search display), loss of resolution may cause more ambiguity in the color of real-world objects (blending colors together) than in simple homogenous-color stimuli. Color does, however, appear to still be useful in guiding search in such complex displays: Guidance correlates more strongly with color features than with other features in scenes (Hwang, Higgins, & Pomplun, 2009), and search is faster in color displays than grayscale displays, at least for some classes of targets (such as flowers; Hayakawa, Kawai, & Masataka, 2011).

Experiment 2 directly tested whether the preview-mismatch effects found by Vickery et al. (2005) and Bravo and Farid (2009) extend to the hue dimension in real-world stimuli. Importantly, when considered in combination with the fact that orientation mismatch effects were only observed in the grayscale version of Experiment 1, the existence of strong hue mismatch effects in Experiment 2 would suggest that hue dominates search guidance and, when present, overrode the orientation guidance signal in Experiment 1.

4.1.2.2 Stimuli/procedure specifications

Twenty-five different increments of hue change were created using the methods described in section 3.3.2 and tested in 15° increments from 0° to 360°. In addition, a 5° mismatch condition was included, to explore a finer-grained amount of mismatch. These conditions matched the equivalent angle changes in Experiment 1, and are: 0° ("no change"), 5°, 15°, 30°, 45°, 60°, 75°, 90°, 105°, 120°, 135°, 150°, 165°, 180°, 195°, 210°, 225°, 240°, 255°, 270°, 285°, 300°, 315°, 330°, and 345° of clockwise change (where degrees of hue change are measured in terms of hue angle in HSV space). Roughly speaking, a clockwise change of 30° results in red changing to orange or orange changing to yellow. Note that because HSV color space is not a color-opponency space it is not true for all colors that there is a constant amount of hue angle difference between a given color and its perceptual opposite or complement. For example, the distance between red and green is approximately 120°, while the distance between blue and yellow is approximately a 180° change (see Figure 2). Subjects performed twenty-five practice trials (one for each increment of change). Twenty-five experimental trials were included for each increment of change, for a total of 650 experimental trials. In all other respects the procedure was the same as the general procedure (section 3.4).

4.1.2.3 Results and discussion for Experiment 2

All ANOVAs in this section are one-way within-subjects ANOVAs comparing the specified dependent measure to the 25 within-subject mismatch conditions in this experiment, except where otherwise specified.

Manual reaction times and accuracy

Reaction times increased with hue mismatch—$F(24,576) = 26.60$, $p < .001$—indicating that hue was used to perform this task. There was also a marginally significant accuracy decrease with increasing hue match—$F(24,576) = 1.96$, $p = .05$. Note that this does not indicate a speed-accuracy tradeoff. Increasing RTs with decreasing accuracy instead indicates a general effect of task difficulty due to increasing hue mismatch. Error trials were removed from all other analyses.

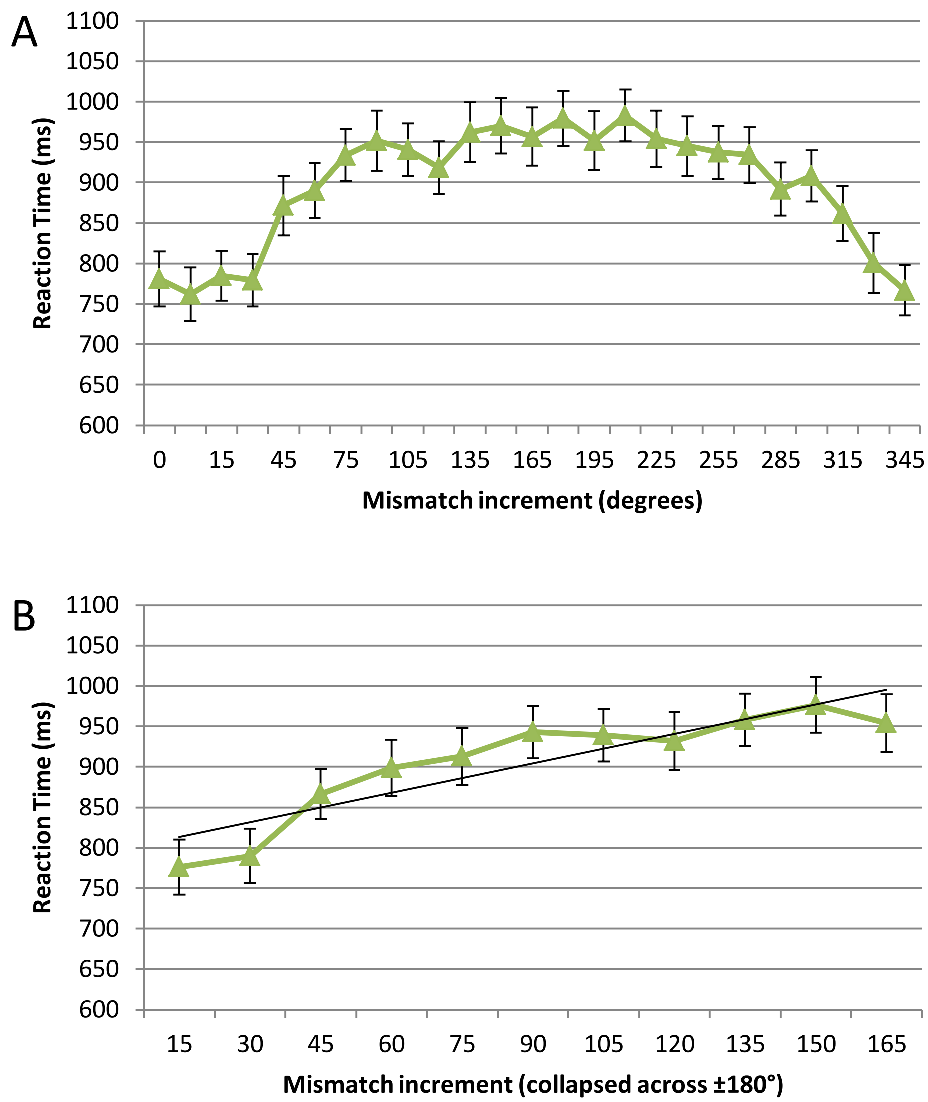

Figure 10. Change in RT across hue mismatch conditions in Experiment 2, for each mismatch condition (A) and collapsing across clockwise and counterclockwise, with a best-fit line indicating the linear component of the mismatch effect in the grayscale version of the experiment (B). Error bars indicate one standard error of the mean.

Reaction times increased (with accompanying decreases to accuracy) as mismatch increased towards 180° of change, and decreased between 180-360° of change, consistent with the circular nature of hue angle; see Figure 10A. It is important to note, however, that this effect (and similar effects in the measures discussed below) plateaus around ±75° of change, suggesting that hue changes beyond that distance in HSV space always result in the "wrong" hue regardless of the amount of change. This is apparent in Figure 3, where the predominately-red object becomes yellow, green, or blue within that range, all of which are no longer red.

As discussed in section 2.1, effects of preview-target mismatch on RTs have previously been found, and so were expected here. However, the inclusion of numerous smaller increments of mismatch in the present experiment allowed for linearity of the testing mismatch effects. In addition, while previous overall RT effects have been found, the current work better characterizes these effects by breaking down RTs into eye movement-based measures, as discussed below. As in Experiment 1, linear and quadratic contrasts were significant, $F(1,24)>48.60$, $p < .001$, with mismatch effects generally increasing linearly, but leveling off at higher amounts of mismatch--see Figure 10B. This extends the finding that mismatch effects includes both linear and quadratic components. In the case of hue, however, there is an additional reason to expect this particular pattern: Because HSV space is not a perceptually-defined color space, distances in hue angle do not accurately reflect a perceptual similarity metric. Importantly, once hue angle has changed by a sufficient amount, it is likely that the hue of the target has changed to a different color category (i.e. red becoming blue) and it may no longer be relevant what the actual distance in hue angle is. Instead, all values beyond that range are simply not validly matched to the preview.

Guidance

Time-to-target was significantly affected by hue mismatch—$F(24,576) = 17.94$, $p < .001$, Figure 11—indicating that hue was used in guiding search to the target. This is consistent with prior work suggesting that search is predominantly driven by color (e.g. Pomplun, 2006; Williams, 1966). Linear and quadratic contrasts were significant, $F(1,24)>4.89$, $p < .05$, consistent with the overall RTs and with Experiment 1: Time-to-target generally increased linearly, but leveled off at larger amounts of hue mismatch--see Figure 11B.

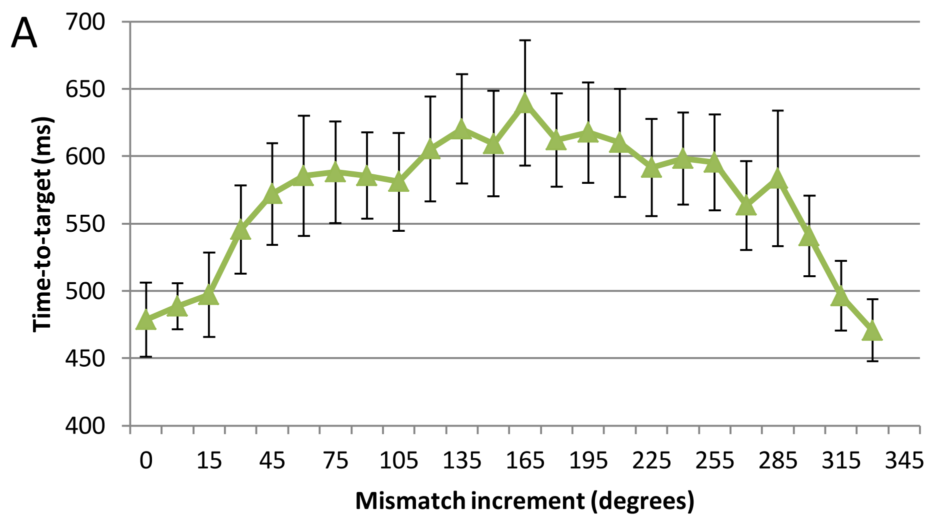

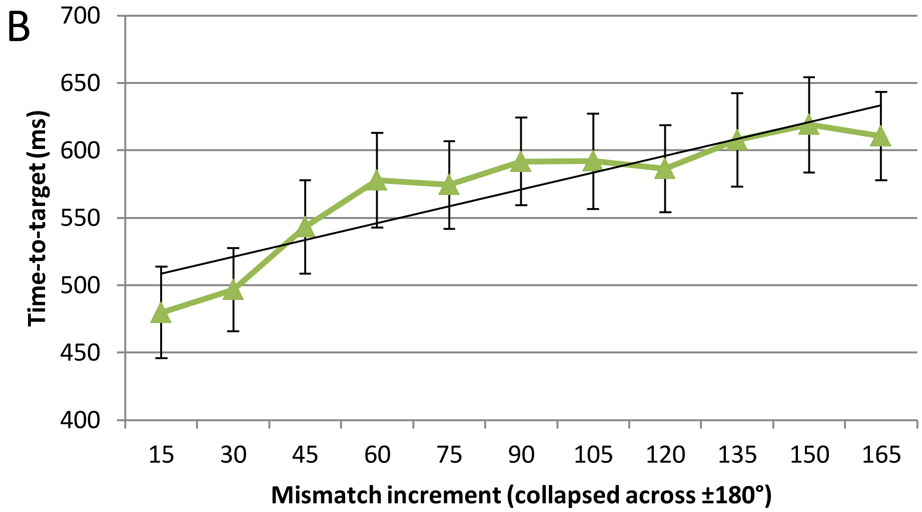

Figure 11. Change in time-to-target across hue mismatch conditions in Experiment 2, for each mismatch condition (A) and collapsing across clockwise and counterclockwise, with a best-fit line indicating the linear component of the mismatch effect in the grayscale version of the experiment (B). Error bars indicate one standard error of the mean.

The amount of hue mismatch significantly affected how often the target was the first object fixated—$F(24,576) = 17.97$, $p < .001$, Figure 12—demonstrating that early search guidance processes rely at least partly on hue information. If subjects made faster eye movements in conditions with poorer guidance, it would be possible that evidence of poor guidance in early saccades is actually the result of subjects making speeded eye movements, and not due to weaker guidance signals. However, the time to fixate any object first after display onset was *longer* with increasing hue mismatch—$F(24,576) = 5.10$, $p < .001$—demonstrating that the pattern in target fixated first was not due to a speed-accuracy tradeoff. Note also that more information than hue must be involved in guiding search to the target, as guidance never dropped to chance levels, $t(24)>6.75$, $p < .001$.

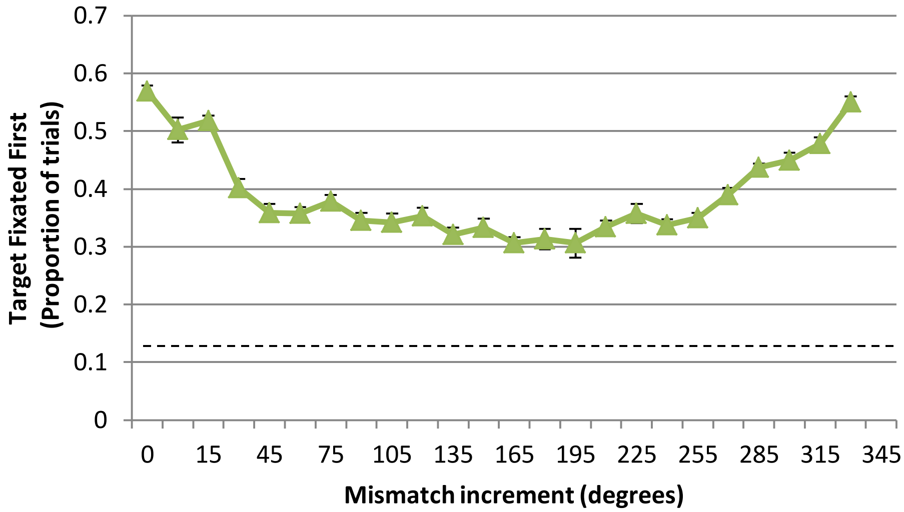

Figure 12. Change in the proportion of trials where the target was fixated first across hue mismatch conditions in Experiment 2. The dashed line indicates chance levels of guidance (0. 125). Error bars indicate one standard error of the mean.

Verification times

Verification times increased as hue mismatch increased up to 180° and then decreased as mismatch increased beyond 180°—$F(24,576) = 4.92$, $p < .001$, Figure 13. This demonstrates that subjects were using hue to verify the identity of the target, and is consistent with the circular nature the hue angle feature space. The similarities between the pattern of mismatch effects in verification times and in guidance measures suggest that hue may be important or useful in several aspects of the search task, or may suggest that a common process (utilizing the same features) may be involved in both guidance and target verification.

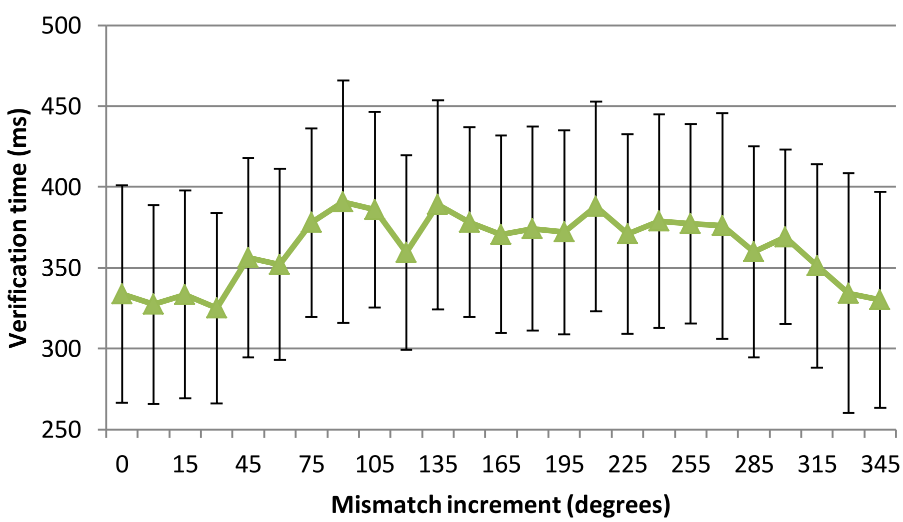

Figure 13. Change in verification time across hue mismatch conditions in Experiment 2. Error bars indicate one standard error of the mean.

The importance of individual feature dimensions for search

To test the relative impact of hue and orientation mismatches on search performance, the mean and maximum difference scores (observed value at each mismatch condition minus the observed value for the "no change" condition) for each dependent measure discussed above was computed for Experiments 1 and 2. For each dependent variable, independent samples *t*-tests were performed comparing the mean and maximum scores between each feature dimension. While it is true that the feature spaces for hue and orientation may be scaled differently, preventing direct comparisons of effects at each degree of change (45° of rotation is not necessarily the same perceptual distance from no change as 45° of hue angle), the full range of feature space for hue angle and orientation (0-360°) was sampled, providing an estimate of their effects across the entire range. This same logic cannot be applied to aspect ratio, as aspect ratio has no clearly defined maximum range: Effects of increased mismatch in aspect ratio could continue beyond the range the current experiments explored. Guidance (both in terms of maximum time-to-target and minimum proportion of trials where the target was fixated first) was significantly worse in Experiment 2 than in Experiment 1 (both $p < .001$), suggesting that altering hue has a greater impact on search guidance than orientation. It is interesting to note, however, that accuracy and verification times did not have significantly different extremes for Experiments 1 and 2 ($p \geq .14$), suggesting that this effect is specific to guidance, rather than other search processes. Previous work measuring correlations between eye movements and visual features of targets came to similar conclusions, that color drives guidance more than other features (e.g. Hwang et al., 2009; Pomplun, 2006; Rutishauser & Koch, 2007). The present finding confirms this earlier work but demonstrates experimentally the importance of hue information rather than relying on correlational measures.

Note too that although hue appears to be more important, other features must also be in use, as no condition in Experiment 2 reached chance levels of guidance. If hue was the *only* feature used to guide search, sufficiently change to a highly mismatching hue should have resulted in *below*-chance performance (as distractors become more likely to match the preview than the target itself)! This is also consistent with previous inferential tests of feature importance. For example, Hwang et al. (2009) demonstrated through the use of Pearson correlations that while eye movements are more likely to land on locations with target-similar hue than target-similar orientation, guidance was also at above-chance levels to locations with target-similar spatial frequencies and target-similar orientation.

4.1.3 Experiment 3: Shape Mismatch

4.1.3.1 Background information regarding shape as a guiding feature

Shape has been shown to meet the criteria previously used to identify guiding features (Wolfe & Horowitz, 2004). Search can be efficient to targets with at least some unique shape features (e.g. Levin, Takarae, Miner, & Keil, 2001), some shape features can easily be segmented from texture backgrounds (Simmons & Foster, 1992), and search asymmetries can be found for some shape features (Kristjánsson and Tse, 2001; Treisman & Gormican, 1988; Wolfe, Yee, & Friedman-Hill, 1992). Efficient search based on shape appears to rely on local shape features (line terminators, T-junctions, curvature, convexity, etc.), rather than on the global, holistic shape of the objects (Wolfe and Bennett, 1997). In other words, search cannot be effectively guided to a target based on the overall form of its outline, but could be guided to targets based on local attributes on that outline. There is some evidence that overall form can be processed preattentively, as demonstrated by the preattentive completion of shapes behind occluders (e.g. Rensink & Enns, 1998). To the extent that this is true, this would require determining at least some global form information, but these findings have recently been called into question (Wolfe, Reijnen, Horowitz, Pedersini, Pinto, & Hulleman, 2011). There is also limited evidence supporting preattentive processing of one global shape feature: Aspect ratio.

Treisman and Gormican (1988) varied aspect ratio in a task where ellipses with a height-to-width ratio of 1.42 were easier to find among circles (which have a height-to-width ratio of 1) than circle targets among ellipses. While they did not find that the ellipses popped-out (which would indicate preattentive coding of aspect ratio), the authors suggested that more extreme variations in aspect ratio might be coded preattentively. Their results suggest that aspect ratio could be processed in relatively early vision and could be used to discriminate search targets. This is not completely clear, however, as ellipses may be easier to find among circles for a variety of reasons: It may be easier to find an object with a less prototypical aspect ratio, as Treisman and Gormican suggest (features that deviate from a "default" prototypical value may be processed separately from the default value), or it may be that subjects use some other feature that varies between ellipses and circles. Distance between the sides differs between ellipses and circles, although Treisman and Gormican demonstrated that it is probably not the case that their subjects were simply using either the width or height of the ellipses to discriminate the targets, as there was no effect of having the orientation of the ellipses vary randomly. Also, differences in curvature also distinguish ellipses from circles, and curvature may effectively guide search (Levin, Takarae, Miner, & Keil, 2001; Treisman & Gormican, 1988; Wolfe, Yee, & Friedman-Hill, 1992; but see Kristjánsson & Tse, 2001). It is clear, though, that some shape feature played a role, and that that feature varied with aspect ratio.

Kristjánsson and Tse (2001) also tested for a search asymmetry using similar stimuli in their experiment 3, but instead of ellipses and circles they used half-ellipses and half-circles (curves that differed only in the rate of change of their curvature). Kristjánsson and Tse failed to find a search asymmetry in their task, and concluded that the rate of change of curves is not detected preattentively (although curvature discontinuities—*changes* in the rate of change—can). It is possible that the reason the ellipse and circle stimuli in Treisman and Gormican (1988) allowed preattentive processing of aspect ratio, while Kristjánsson and Tse's (2001) stimuli did not, is that Treisman and Gormican's stimuli were closed figures. The shape of an open figure may simply be more difficult to discriminate in the periphery.

With simple stimuli it is straightforward to have a target defined by a single shape feature. The target could be defined by length (a long line among short lines), size (a large square among smaller), curvature (curved lines among straight lines), aspect ratio, or any number of other dimensions. With real-world stimuli, however, modifying a single shape feature is likely to change other features as well. Elongating an object also changes the curvature of any curves present on the object and altering height-to-width ratio also obviously alters both height and width of the objects, not just the ratio between them. Experiment 3 therefore does not tease apart whether it is the aspect ratio per se that alters guidance to the targets. Instead, the goal is to determine whether shape (more generally) is used in searching for real-world objects. In other words, the concern is not the specific kinds of shape features, but whether or not shape in general is important.

Experiment 3, by varying the aspect ratio of the target relative to the preview, tested whether shape information (in terms of aspect ratio) is used in searching for real-world stimuli, even in the presence of uncertainty along that dimension. As in Experiments 1 and 2, knowledge that the aspect ratio information in the preview would almost always be inaccurate may have led to a de-weighting of this feature dimension. As a result, any effects of this manipulation would suggest that shape information is not only used in searching for realistic objects, but that such information is important enough to be used even when it is less valid than other accessible features. As in Experiment 1, grayscale and full-color versions were conducted to test whether shape was used with the presence and/or in the absence of hue information.

4.1.3.2 Stimuli/procedure specifications

Twenty-five aspect ratio conditions were created. The first condition had no change, and twenty-three other conditions stretched the images in 5% increments, with the most mismatched condition stretched 215%. As in Experiments 1 and 2, an additional finer-grained condition was included, stretching the image 2.5%. Stretching was always done in terms of increasing the object width (though the decision to increase width, rather than height, was arbitrary). The conditions are: 100% normal size ("no change"), 102.5%, 105%, 110%, 115%, 120%, 125%, 130%, 135%, 140%, 145%, 150%, 155%, 160%, 165%, 170%, 175%, 180%, 185%, 190%, 195%, 200%, 205%, 210%, and 215% increase in width (with a corresponding decrease in height).

Subjects performed twenty-five practice trials (one for each mismatch condition). Twenty-five experimental trials were included for each mismatch condition, for a total of 650 experimental trials. In all other respects the procedure was the same as the general procedure (section 3.4). The grayscale version was identical to the full-color version in all respects save that the stimuli were grayscale and hue-change conditions were omitted. As such, there were only 17 practice trials and 425 experimental trials.

4.1.3.3 Results and discussion for Experiment 3

All ANOVAs in this section are one-way within-subjects ANOVAs comparing the specified dependent measure to the within-subject mismatch conditions in each experiment, except where otherwise specified. All comparisons were $p > .13$ for the full-color version of the experiment, indicating that when hue is available to guide search, shape does not contribute significantly to the search process. This is consistent with Experiment 1 and with other work suggesting that color dominates search (e.g. Williams, 1966; Pomplun, 2006). However, significant effects in the grayscale version of Experiment 3 (discussed below) demonstrate that shape information *is available to be* used by the search process, but is simply not used (or does not contribute significantly) when hue is available. All further results reported in this section are for the grayscale version of the experiment.

Manual reaction times and accuracy

Accuracy was not reliably affected by aspect ratio mismatch—$F(24,576)=.72$, $p = .73$—indicating that any effects in RT are not due to a speed-accuracy tradeoff. Error trials were removed from all other analyses. Reaction times significantly increased with increased aspect ratio mismatch—$F(24,576)=1.98$, $p < .05$—indicating that subjects relied on shape in order to perform the task. Relative to the other experiments, the data in the grayscale version of Experiment 2 appear to be particularly noisy (Figure 14). Speculatively, this may reflect a particular vulnerability to distractors in the aspect ratio portion of the guidance signal: The aspect ratio of the distractors varied randomly, and the resulting values may have been more detrimental in some conditions than in others.

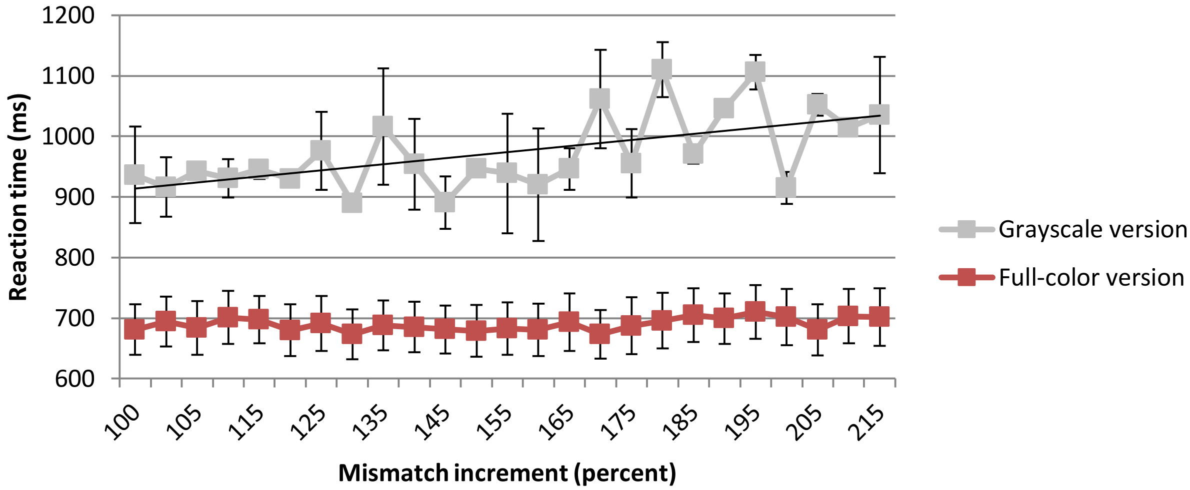

Figure 14. Change in manual RT across shape mismatch conditions in Experiment 3, with a best-fit line indicating the linear component of the mismatch effect in the grayscale version of the experiment. Error bars indicate one standard error of the mean.

Unlike the orientation and color spaces, aspect ratio is not circular and therefore mismatch conditions did not need to be collapsed across positive and negative rotation in order to test for linearity. Linear and quadratic contrasts were again significant, $F(1,24)>13.89$, $p < .001$, indicating that, as with the hue and orientation mismatch effects, the aspect ratio mismatch effect on RTs generally increases linearly but levels off at higher amounts of mismatch. With large amounts of preview-target mismatch, subjects may be treating any specific aspect ratio as simply "incorrect," or the aspect ratio of the target may be becoming more dissimilar to the preview than the aspect ratios of distractor items.

Guidance

Time-to-target significantly increased with increased aspect ratio mismatch—$F(24,576)=1.98$, $p < .05$—suggesting that shape is used in guiding search to the target when hue is not available (Figure 15). Because time-to-target could be affected by factors unrelated to guidance (such as the time it takes subjects to reject distractors they fixate before the target), a more conservative measure of early search guidance was again used: The percentage of trials where the target is the first object fixated. "Target first fixated" marginally decreased with increased aspect ratio mismatch—$F(24,576)=1.70$, $p = .07$ (Figure 16). This suggests that early search guidance processes might rely at least partly on shape, though other features were involved as well, indicated by guidance never dropping to chance levels when shape was mismatched ($p < .001$ for each condition). Unlike the pattern in overall RTs, in which linear and quadratic contrasts were significant, and unlike Experiments 1 and 2, only the linear contrast was significant for Experiment 3 ($F(1,24) = 33.07$, $p < .001$): Time-to-target increased linearly with increased aspect ratio mismatch, and did not level off at high amounts of change.

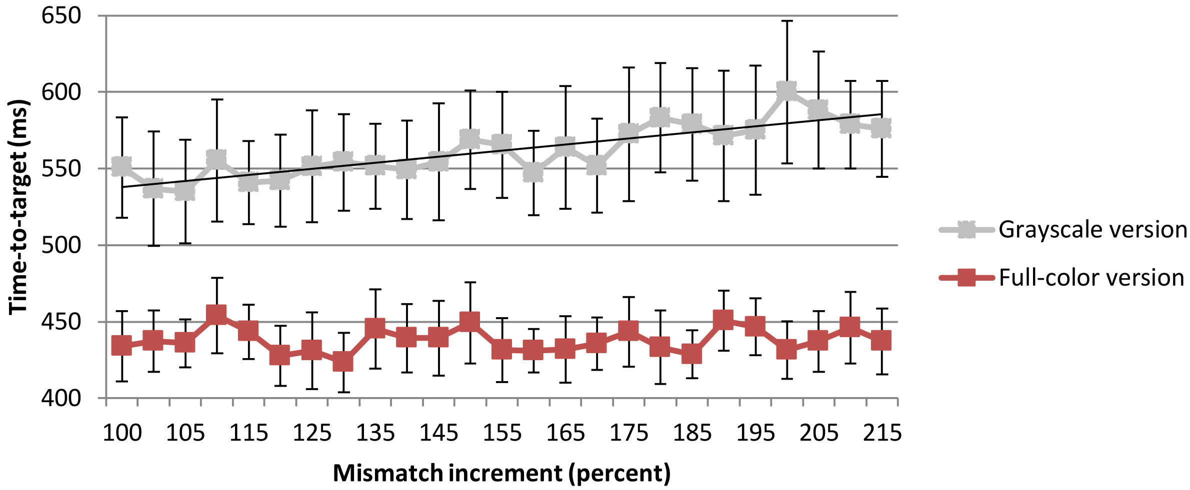

Figure 15. Change in time-to-target across shape mismatch conditions in Experiment 3, with a best-fit line indicating the linear component of the mismatch effect in the grayscale version of the experiment. Error bars indicate one standard error of the mean.

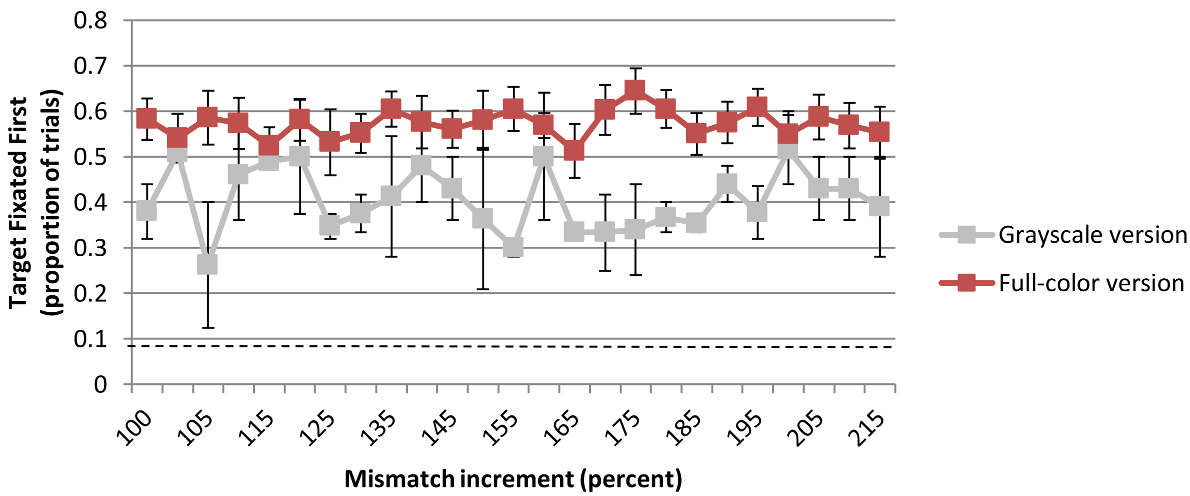

Figure 16. Change in the proportion of trials where the target was the first item fixated across shape mismatch conditions in Experiment 3. The dashed line indicates chance. Error bars indicate one standard error of the mean.

To test for speed-accuracy tradeoffs that might bias the early guidance measure (target first fixated), one-way ANOVAs were conducted to compare the effects of each increment of feature change on the time it took subjects to first fixate *any* object (target or distractor) after display onset. If fast eye movements result in fewer first target fixations, then this factor might explain the results, rather than indicating a weak guidance signal. However, as in Experiment 2 this was not the case: the time to fixate any object did not vary with mismatch increment ($F(8,392) = .29$, $p = .79$)

Verification times

Verification times significantly increased with increased aspect ratio mismatch—$F(24,576)=3.78$, $p < .001$—demonstrating that shape was used in responding to the target in this task (Figure 17).

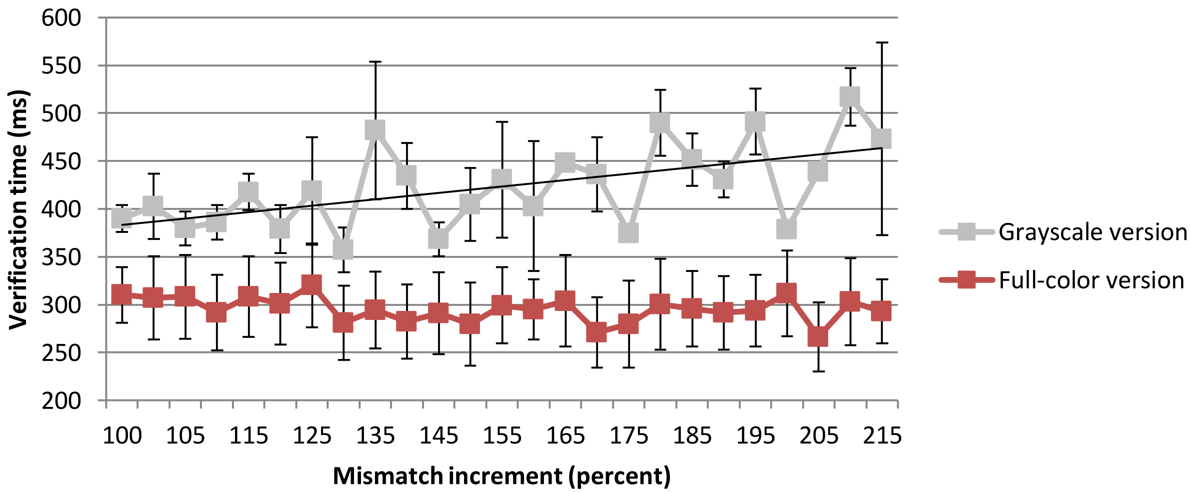

Figure 17. Change in verification time across shape mismatch conditions in Experiment 3, with a best-fit line indicating the linear component of the mismatch effect in the grayscale version of the experiment. Error bars indicate one standard error of the mean.

4.2 Experiments for Aim 2

Aim 2 investigated the effects of mismatch between previews and search targets when the dimension that changed on a given trial was uncertain. Two experiments were conducted, one full-color and one grayscale. The experiments were identical to those of Aim 1 save that in the full-color version all three of the feature dimensions manipulated in Aim 1 (hue, orientation, and aspect ratio) were manipulated, and the feature type that changes on a given trial was interleaved. Similarly, both orientation and aspect ratio were manipulated and interleaved in the grayscale version (hue-change trials were not included).

4.2.1 Experiment 4: Interleaved feature dimension mismatches

Aim 1 tested whether search guidance with real-world objects relies on color, orientation, or shape information, even when the values the target would have on that dimension are uncertain. Experiment 4 tested whether those dimensions are used when the values on those dimensions are perfectly valid on the majority of trials and when subjects are uncertain of *which* dimension will change between preview and the search display. This increased validity and decreased uncertainty was created by interleaving trials in which each of the three feature dimensions changed. Unlike the experiments in Aim 1, subjects did not know which feature dimension would change on a given trial and each feature dimension was substantially less likely to change than in Aim 1, in which the manipulated feature dimension changed in all trials except the no change condition ($1/25^{th}$ of trials in the experiment).

4.2.1.1 Stimuli/procedure specifications

Twenty-five within-subject conditions were created for Experiment 4. The first condition had no change on any dimension. The twenty-four other conditions consisted of eight conditions with hue changes, eight with orientation changes, and eight with aspect ratio changes. These conditions sampled those used in Experiments 1, 2, and 3, rather than including all 73 different conditions, so as to avoid a prohibitive number of trials. The eight hue-change conditions consisted of 30° increments, starting with 30° and ending with 240° of change (30°, 60°, 90°, 120°, 150°, 180°, 210°, and 240°). The eight orientation-change conditions consisted of the same 30° increments (30°, 60°, 90°, 120°, 150°, 180°, 210°, and 240°). The eight aspect ratio conditions consisted of increments of 10% change, from 110% original width to 180% original width (110%, 120%, 130%, 140%, 150%, 160%, 170%, and 180%). These conditions were interleaved, and as a result approximately $1/3^{rd}$ of the trials were followed by another trial in which the same dimension changed. Note that all of these conditions have an equivalent in either Experiment 1, 2, or 3, and only these overlapping conditions were included in comparisons made between Experiments 1-3 and Experiment 4 (see section 4.2.1.2). As in Experiments 1-3, the target previews were not manipulated.

Subjects performed twenty-five practice trials (one for each increment of change). Twenty-five experimental trials were included for each increment of change, for a total of 650 experimental trials. For the grayscale version of the experiment, the hue-change conditions were omitted. As a result, there were only 17 practice trials and 425 experimental trials, and there was an approximately 50% chance that the same feature changed on one trial changed again on the following trial. In all other respects the procedure for both experiments was the same as the general procedure (section 3.4).

4.2.1.2 Results and discussion for Experiment 4

4.2.1.2.1 Results and discussion for the full-color version of Experiment 4

Unlike Experiments 1 and 2, significant mismatch effects were found in both the full-color and grayscale versions of Experiment 4. As such, results from each experiment version will be considered in full. The results below all refer to the full-color version of the experiment. Results for the grayscale version are provided in section 4.2.1.2.2. Except where otherwise noted, all analyses in this section are 3(hue/shape/orientation) x 9(mismatch increments) ANOVAs, with Bonferroni-corrected post hoc *t*-tests.

Manual reaction times and accuracy

Accuracy was not reliably affected by mismatch increment ($F(8,192)=.48$, $p = .79$), feature dimension ($F(2,48)=2.34$, $p = .14$), or the interaction of mismatch increment and feature dimension ($F(16,384)=.40$, $p = .84$), indicating that any effects in RTs were not due to a speed-accuracy trade-off. Error trials (including timed-out trials) were removed from all other analyses (including the RT analysis). Consistent with the results of Experiment 2, RTs increased with increasing mismatch—$F(8,192)=33.52$, $p < .001$, and this was driven by differences in hue rather than aspect ratio or orientation, as confirmed by a main effect of feature dimension, $F(2,48)=195.47$, $p<.001$, and an interaction between feature dimension and the amount of mismatch, $F(16,384)=27.53$, $p<.001$—see Figure 18A. Post hoc tests revealed significant differences between hue and both orientation and aspect ratio (ps < .001). Aspect ratio did not reliably differ from orientation ($p=.23$). To further explore mismatch effect in the orientation and aspect ratio conditions, the same comparisons were performed with hue removed. Unlike in Experiments 1 and 3, a significant main effect of mismatch in orientation and aspect ratio was found, $F(8,192)=3.61$, $p<.01$, with mismatch effects increasing as the mismatch increment increased. Orientation and aspect ratio did not interact, $F(8,192)=.99$, $p=.43$, and orientation and aspect ratio only trended towards a difference, $F(1,24)=3.40$, $p=.08$. In the context of Experiment 4, where participants are uncertain of what feature dimension will change on a given trial, orientation and aspect ratio mismatch effects were found even in the full-color version of the experiment, suggesting that in Experiments 1 and 3, participants were able to deweight (at least to some degree) dimensions that were known to be invalid.

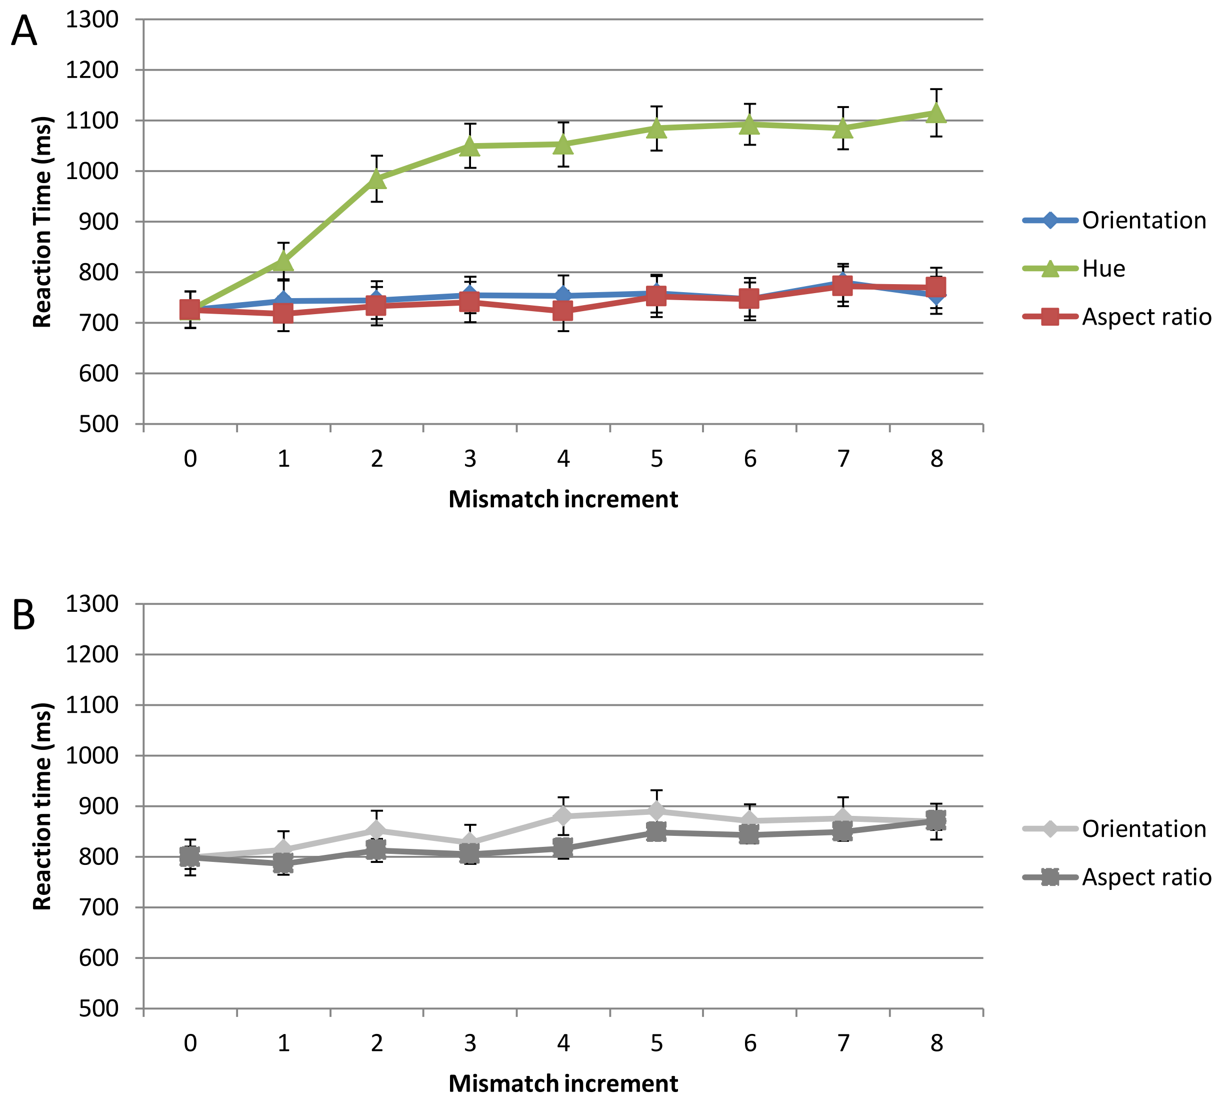

Figure 18. Change in RTs across feature dimension conditions in the full-color (A) and grayscale (B) versions of Experiment 4. Error bars indicate one standard error of the mean.

Guidance

Feature dimension and amount of mismatch significantly interacted to affect time-to-target—$F(16,384) = 19.152$, $p < .001$—with the largest effects at higher mismatches for the hue dimension—See Figure 19A. Main effects of mismatch—$F(8,192) = 22.01$, $p < .001$—and feature dimension—$F(2,48) = 112.97$, $p < .001$—were also found. As with RTs, these effects were driven by the mismatch effect in the hue condition, which was significantly different from orientation and aspect ratio mismatch conditions, both $p < .001$. Aspect ratio and orientation mismatch conditions did not differ, $p = 1.00$. To further explore the aspect ratio and orientation dimensions, the same analyses were again performed but with hue removed. Consistent with Experiments 1 and 3, no reliable main effect of mismatch in orientation and aspect ratio was found, $F(8,192)=.70$, $p < .64$. Orientation and aspect ratio did not interact, $F(8,192)=.68$, $p = .65$, and no main effect of feature dimension was found, $F(1,24)=.04$, $p = .85$. This pattern suggests that the difference in overall RTs was not driven by guidance effects.

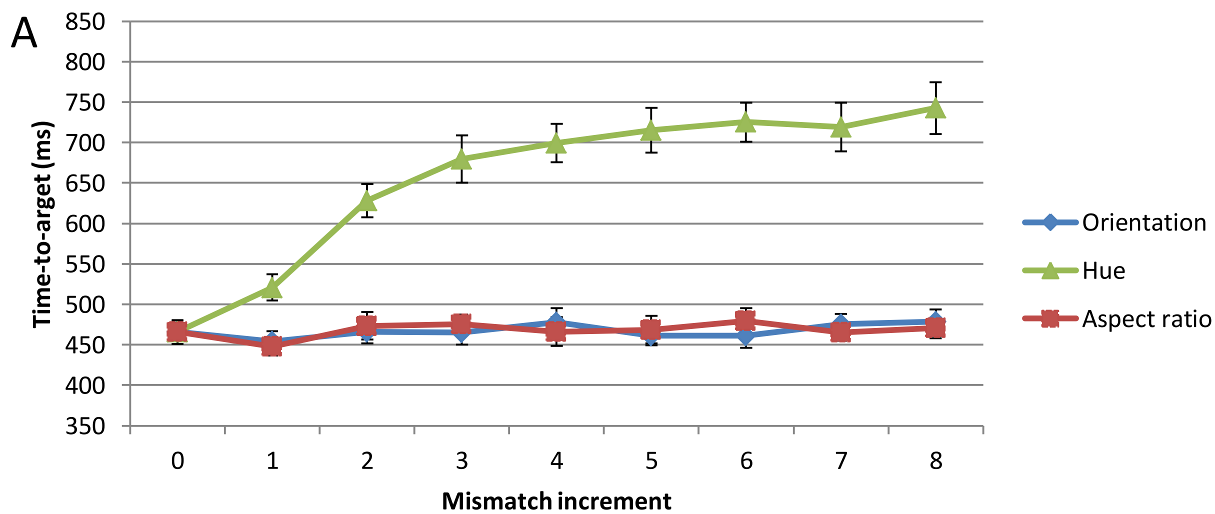

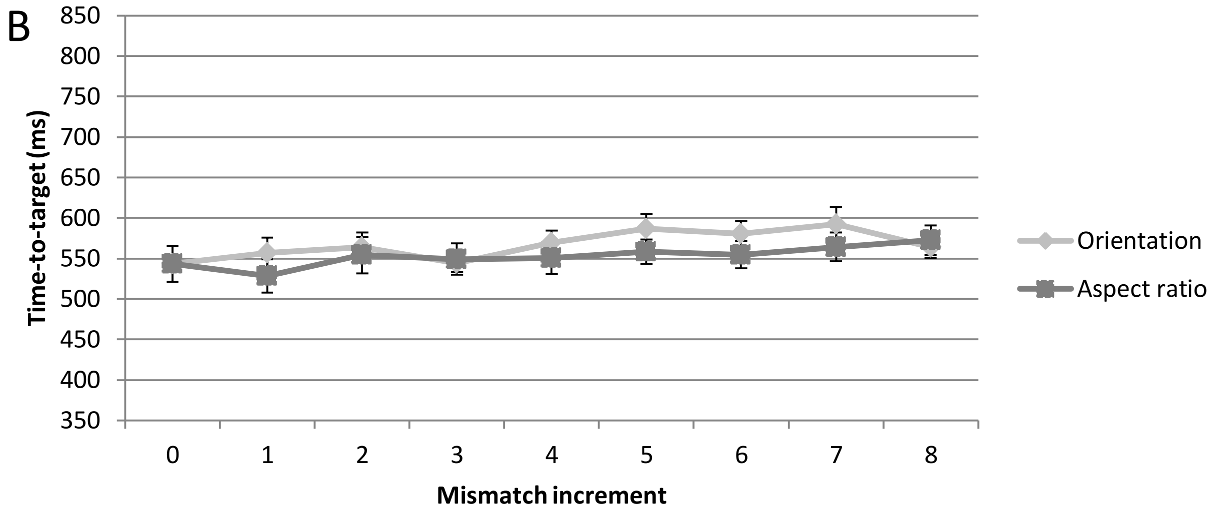

Figure 19. Change in time-to-target across mismatch conditions in the full-color (A) and grayscale (B) versions of Experiment 4. Error bars indicate one standard error of the mean.

A main effect of feature mismatch on whether the target was the first object fixated was found—$F(8,192) = 14.46$, $p < .001$—suggesting that the effect of mismatch on guidance took place early enough to influence the first fixations on objects (Figure 20A). Differences were also found across feature dimension—$F(2,48) = 118.99$, $p < .001$—and feature dimension interacted with mismatch increment—$F(16,384)=20.67$, $p < .001$—with an effect of hue mismatch driving the effects. Guidance was worst on hue mismatch trials ($p < .001$) and there was no difference between aspect ratio and orientation mismatch conditions ($p = .78$). To further explore the aspect ratio and orientation dimensions, the same analyses were again performed but with hue removed. Unlike in Experiments 1 and 3, no reliable main effect of mismatch in orientation and aspect ratio was found, $F(8,192)=1.48$, $p = .20$. Orientation and aspect ratio did not interact, $F(8,192)=.85$, $p = .53$, though a significant main effect of feature dimension was found, $F(1,24)=11.03$, $p < .01$, with lower guidance in orientation-change trials than aspect ratio-change trials. Consistent with time-to-target, this pattern suggests that the orientation and aspect ratio mismatch effects in overall RTs were not driven by differences in guidance. Guidance was above chance levels in all conditions, all $p < .001$, suggesting that there were other features remaining to be used in this task.

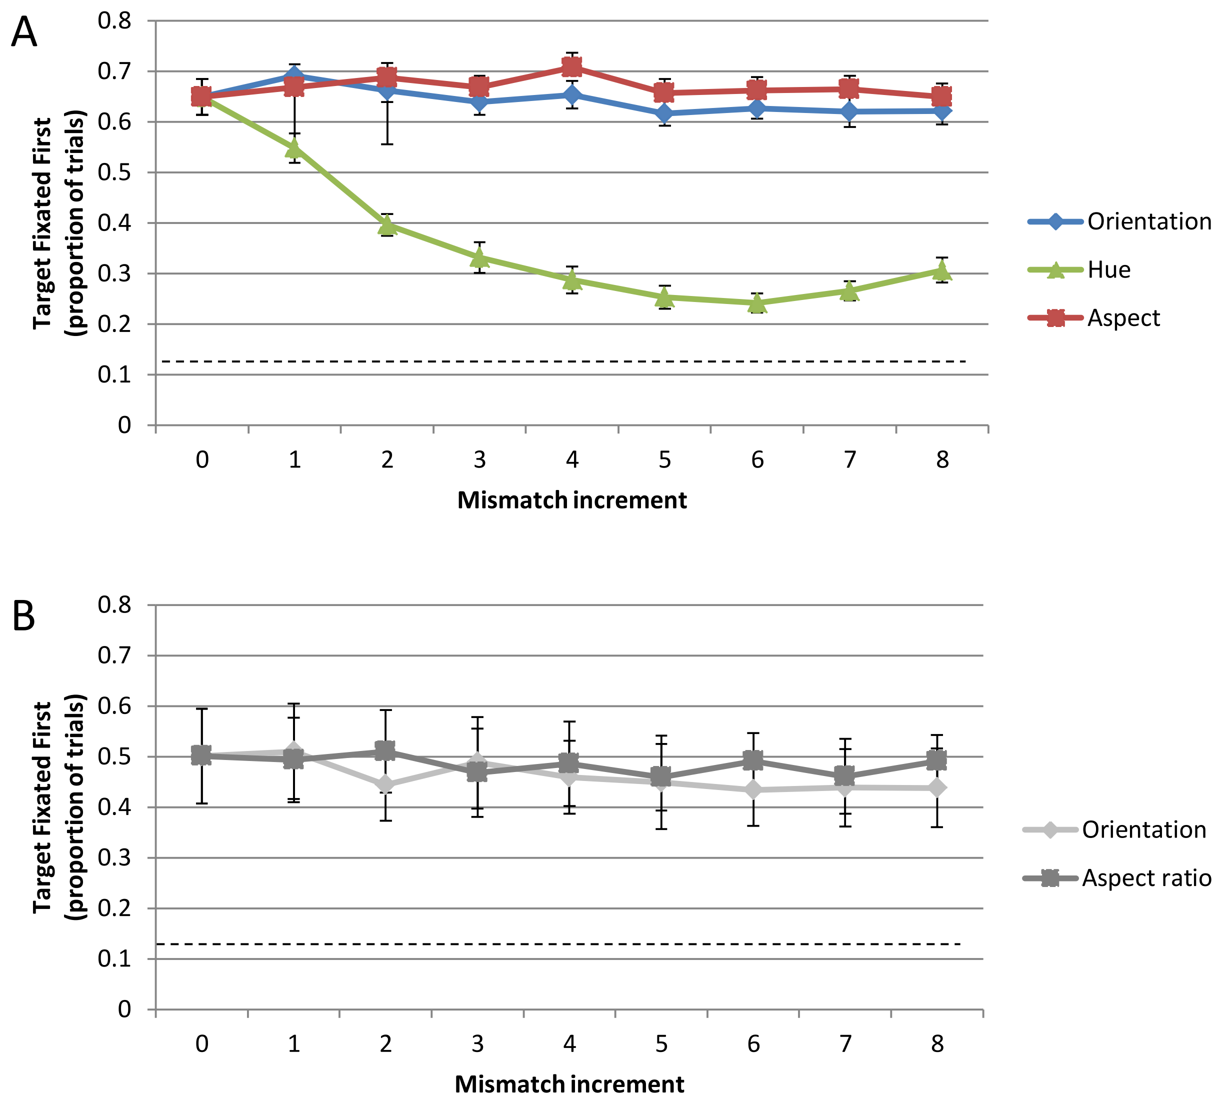

Figure 20. Change in target first fixated across mismatch conditions in the full-color (A) and grayscale (B) versions of Experiment 4. The dashed lines indicate chance levels of performance (0.125). Error bars indicate one standard error of the mean.

To test for speed-accuracy tradeoffs that might bias the early guidance measure (target first fixated), one-way ANOVAs were again conducted to compare the effects of each increment of feature change on the time it took subjects to first fixate *any* object (target or distractor) after display onset. As in Experiments 2 and 3, no evidence for such a speed-accuracy trade-off was found: The time to fixate any object did not vary with mismatch increment ($F(8,392) = .76$, $p = .53$), nor with feature dimension ($F(1,49)=2.11$, $p = .15$), nor did mismatch increment and feature dimension interact ($F(8,392)=.69$, $p = .48$).

Verification times

Verification times were again affected by feature mismatch (Figure 21A). Verification times were significantly affected by amount of mismatch—$F(8,192) = 6.99$, $p < .001$—which interacted with feature dimension—$F(16,384)=3.25$, $p < .01$. No main effect of feature dimension was found—$F(2,48)=1.93$, $p = .18$. Unlike in the guidance measures, no difference was found between hue and aspect ratio ($p = .46$). Consistent with guidance measures, however, hue mismatch affected performance significantly more than orientation mismatch ($p < .001$) and no difference was found between orientation and aspect ratio mismatch conditions ($p = 1.00$). To further explore the aspect ratio and orientation dimensions, the same analyses were again performed but with hue removed. Unlike in Experiments 1 and 3, a significant main effect of mismatch in orientation and aspect ratio was found, $F(8,192)=2.92$, $p < .05$. Orientation and aspect ratio did not interact, $F(8,192)= .86$, $p = .52$, and no main effect of feature dimension was found, $F(1,24)=.80$, $p = .38$. This pattern, taken together with the lack of orientation and aspect ratio mismatch effects in guidance measures (time-to-target and target fixated first), demonstrates that differences in overall RT were not driven by guidance effects. Verification processes were affected by orientation and aspect ratio mismatch in the context of the full-color version of Experiment 4, though not in the full-color versions of Experiments 1 and 3. This suggests that when participants knew that a given feature dimension was going to be invalid on most trials (Experiments 1 and 3), they were able to deweight those features in their decision-making process. When participants were uncertain which feature dimension would change on a given trial (Experiment 4), even though aspect ratio and orientation mismatches did not affect guidance, participants were unable to ignore the orientation and aspect mismatches when making their target present response.

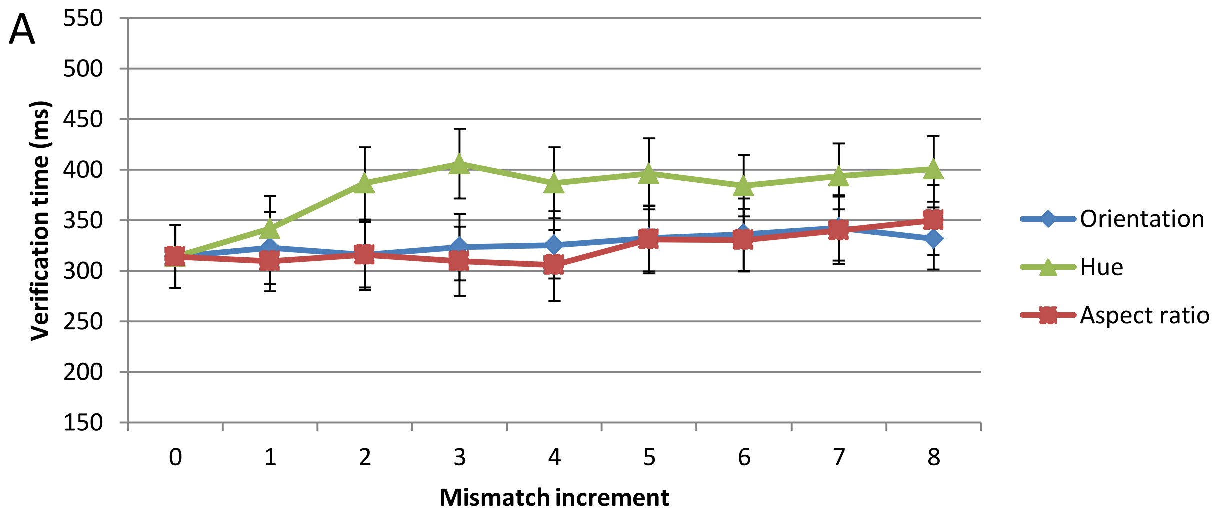

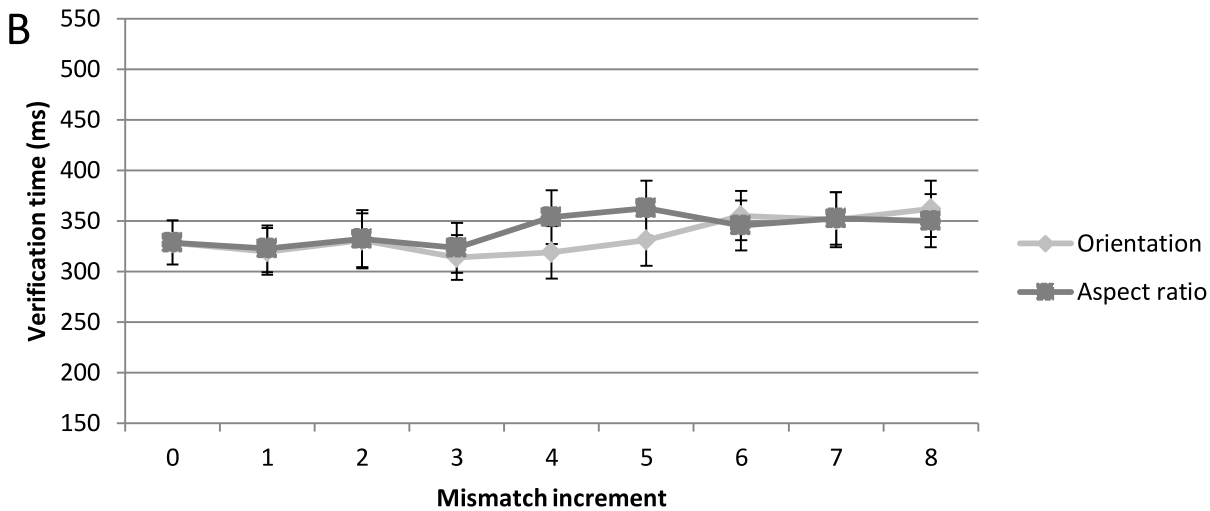

Figure 21. Change in verification time across mismatch conditions in the full-color (A) and grayscale (B) versions of Experiment 4. Error bars indicate one standard error of the mean.

The importance of individual feature dimensions for search

To test for relative levels of hue and orientation importance, rather than using independent sample *t* tests as in Aim 1, the mean and maximum hue-change and orientation-change effects were compared using paired samples *t*-tests for each DV. As in Aim 1, aspect ratio was excluded from the tests comparing mean and maximum effects, but was included in all other analyses. Consistent with the results of Aim 1 and with previous research (e.g. Pomplun, 2006; Williams, 1966), hue mismatch caused significantly worse guidance than orientation mismatch, both in terms of a higher maximum time-to-target ($t(24)=12.13$, $p < .001$) and a lower minimum proportion of trials where the target was the first object fixated ($t(24)=17.60$, $p < .001$). Unlike Aim 1, however, significant differences were also found in verification time and accuracy: Verification times were shorter ($t(24)=7.35$, $p < .001$) and accuracy was lower for hue-change trials ($t(24)=4.22$, $p < .001$), which may simply reflect a speed-accuracy tradeoff.

4.2.1.2.2 Results and discussion for the grayscale version of Experiment 4

The results below all refer to the grayscale version of the experiment. Except where otherwise noted, all analyses in this section are 2(shape/orientation) x 9(mismatch increments) ANOVAs, with Bonferroni-corrected post hoc *t*-tests.

Manual reaction times and accuracy

To the extent that subjects rely on orientation or aspect ratio to perform the task, RT was expected to differ across mismatch conditions. As in the grayscale versions of Experiments 1 and 3, RTs increased as mismatch increased—$F(8,192) = 5.82$, $p < .001$, Figure 18B. A main effect of feature dimension was also found, $F(8,192) = 29.96$, $p < .001$, with a larger effect in orientation than aspect ratio. Mismatch and dimension did not interact ($F(8,192)=1.29$, $p = .27$). Orientation and shape mismatch did not reliably affect accuracy—$F(8,192)=1.36$, $p = .25$—nor was there a main effect or interaction between orientation and shape (ps > .48), indicating that any effects in RTs was not due to a speed-accuracy trade-off. Error trials (including timed-out trials) were removed from all other analyses (including the RT analysis).

Guidance

Unlike the grayscale versions of Experiments 1 and 3, although time-to-target differed across feature dimension, $F(8,192)=7.64$, $p < .05$, with a larger effect in orientation-change conditions, feature dimension and amount of mismatch did not significantly interact to affect time-to-target—$F(8,192)=.94$, $p = .46$—nor was a main effect of mismatch found—$F(8,192)=1.67$, $p = .15$— see Figure 19B.

No reliable effect of feature mismatch on whether the target was the first object fixated was found—$F(8,192)=1.57$, $p = .17$, see Figure 20B—again suggesting that the effect of orientation mismatch on guidance did not take place early enough to influence the first eye movements made to objects, consistent with Experiment 1. The lack of an aspect ratio mismatch effect in this experiment may cast doubt on the marginal effect of aspect mismatch on whether the target was the first object fixated in Experiment 3: A larger mismatch effect was expected in Experiment 4 than in Experiment 3, and if aspect ratio has mismatch effects on early fixations, they should be found here. Guidance was worse on orientation-change trials than shape-change trials— $F(1,24)=9.01$, $p < .01$—though feature dimension did not interact with the amount of mismatch—$F(8,192)=1.38$, $p = .27$, and did not drop to chance levels in any condition (all $p < .001$).

To test for speed-accuracy tradeoffs that might bias the target first fixated measure, conditions were compared on the time it takes subjects to first fixate *any* object (target or distracto6r) after display onset. However, consistent with Experiments 2 and 3, and with the full-color version of this experiment, the time to fixate any object did not vary with mismatch or feature dimension, nor did the conditions interact (ps $\geq .15$).

Verification times

As shown in Figure 21B, verification times were significantly affected by amount of mismatch—$F(8,192)=3.10$, $p < .05$—indicating that orientation and shape were used in responding to the target in this task. As in the full-color version, feature dimension did not reliably affect verification time, $F(1,24)=1.92$, $p = .18$, nor did feature dimension interact with degree of mismatch, $F(8,192)=1.42$, $p = .22$.

Switching costs

Earlier work on guiding features (see section 1.2) primarily used tasks in which targets differed from distractors on a single dimension. If subjects can perform that task efficiently, one can conclude that people can guide search based on that feature dimension. However, evidence that a given dimension *can* result in efficient search does not necessarily mean that subjects *use* information from that dimension when other feature dimensions are available. When the target is defined by shape, for instance, it is necessary for shape information to be processed in order to perform the task. Would shape still be used when color, orientation, and other features are also available, as in the case of real world search tasks? This question can be addressed by comparing Aims 1 and 2. In Aim 1 subjects could assign little or no weight to the manipulated dimension and keep that weight low throughout the experiment. In Aim 2, however, subjects may have adopted very different strategies. It is possible that subjects assigned weights based on the degree of uncertainty and the usefulness of each feature dimension and left those weights unchanged throughout the experiment. Though it may take some time to learn the presence of uncertainty and the amounts of usefulness for particular feature dimensions, it is unclear how long learning these values would require. Switching costs might be present on early trials, but participants in Aim 1 likely knew by the end of the practice trials that only one feature was changing from trial to trial, and participants in Aim 2 similarly knew that the target was changing in multiple dimensions, which may have been sufficient to treat dimensions in Aim 1 as being less reliable than in Aim 2. Regardless, if weights were maintained through most of the experiment, no switching costs should be found when the feature that is manipulated on a trial is different from the feature dimension that was changed on the previous trial. If, instead, intentional or automatic processes adjust the weights given to these dimensions throughout the experiment, a RT cost should be found when there is a change in feature dimension on consecutive trials. When the trial began, the dimensions not altered on the previous trial would first be checked for target feature values. If the manipulated dimension had changed, and as a result there was no strong match to the target on the weighted dimensions, weights would be reassigned, resulting in a processing cost that would slow subjects for that trial.

To test these predictions, analyses were performed to check for dimension-switching costs. Trials where the manipulated feature dimension was different from the dimension manipulated on the previous trial were compared to trials where the same dimension was manipulated as on the previous trial. There were no reliable differences between hue-change trials that were preceded by other hue-change trials ("no switch") or no change trials ("switch")—$t(24)=1.13$, $p = .27$—nor when preceded by aspect-change trials—$t(24)=.20$, $p = .85$—or orientation-change trials—$t(24)=1.41$, $p = .17$. It is possible that the lack of differences here reflect the dominance of hue in performing this task: If hue is always highly weighted, weights may not be shifted sufficiently away from hue in order to create detectable performance differences. Similarly, if aspect and orientation changes do not impact search sufficiently to require a change to weights, then no differences between trials preceded by orientation or aspect ratio trials would be expected. However, orientation-change trials which were preceded by hue-change trials had significantly longer RTs than orientation-change trials preceded by other orientation-change trials, $t(24)=3.23$, $p < .05$, and aspect-change trials which were preceded by hue-change trials had significantly longer RTs than aspect-change trials preceded by other aspect-change trials, $t(24)=3.32$, $p < .01$. This pattern suggests that on (or after) hue-change trials, subjects deweight hue and/or assign more weight to aspect ratio, or possibly that the difficulty of the hue-change trials cause subjects to slow down on following trials in expectation of further difficulty. Contradictory with the former possibility (deweighting hue or adding weight to aspect ratio), and with the concept of reassigning weights during trials, no change trials that were preceded by hue-change trials did not reliably differ from no change trials preceded by other no change trials—$t(13)=1.39$, $p = .18$. Speculatively, this may instead suggest that more-or-less difficult trials alter subject expectancies, and that those expectancies affect RTs. In the case of the more difficult hue-change trials followed by no change trials, subjects are expecting the task to be very difficult but are instead confronted by a relatively easy no change trial, resulting in shorter RTs. In the case of aspect- or orientation-change trials preceded by hue-change trials, subjects are expecting the task to be difficult and, when confronted with a target that is not identical to the preview, may devote additional time to ensuring that the object is actually the target. Taken together, although there was a cost to switching away from hue-change trials, this pattern generally failed to support the presence of dimension-weighting on a trial-by-trial basis in this task.

Uncertainty

To test for changes in mismatch effects due to uncertainty in what dimension changes on a given trial, 2 (experiment) x 9 (mismatch condition) mixed ANOVAs were performed on each measure, for each feature dimension, comparing the results of Experiment 4 with those of Experiments 1, 2, and 3. Only the corresponding 8 conditions in each of the Aim 1 experiments were included in these analyses (the "no change" condition; 30°, 60°, 90°, 120°, 150°, 180°, 210°, and 240° of hue angle or orientation change, and 110%, 120%, 130%, 140%, 150%, 160%, 170%, 180% changes to aspect ratio); other conditions from Experiments 1-3 were excluded. Significant interactions between experiment and mismatch level would demonstrate an effect of uncertainty from trial-to-trial in what feature dimension changed. If subjects in Experiments 1, 2, and 3 were able to effectively de-weight the feature dimension that was being manipulated, the effect of mismatch would be significantly less in Experiments 1, 2, and 3 than in Experiment 4. Models of search that predict the use of valid cued features of the target—such as Guided Search (Wolfe, 1994; Wolfe, Cave, & Franzel, 1989)—predict this pattern to appear in guidance measures (time-to-target and target first fixated). Note that this difference might be substantially more or less pronounced in some feature dimensions than in others. If a dimension is very important or useful for guiding search, it is possible that it was given more weight than other dimensions even in the face of uncertainty on that dimension. A less important or less useful feature dimension might be used only when its value was reliable, and given little or no weight whenever there was uncertainty on that dimension. As a result, differences between feature dimensions could provide further insight into the relative importance of those dimensions and could shine some light on the relative importance of aspect ratio (which cannot properly be addressed through the comparison of mean and maximum effects of change).

If, instead, no difference was found between the mismatch effect in Aim 1 and Aim 2, this would suggest that top-down weighting is adapted preattentively to each individual search trial. Hwang, Higgins, and Pomplun (2009), for example, have suggested that this occurs in the context of search through real-world scenes. Rather than weighting according to which feature defined the target on a previous trial (or possibly in addition to such weighting), subjects are assumed to use a rough heuristic to determine, at scene onset, what features were most important for each particular target-scene pairing, and to weight feature dimensions accordingly. If this could be done in sparse real-world object arrays as well, no interaction between the mismatch effects in Aim 1 and Aim 2 would be found.

Analyses comparing mismatch effects in Aim 1 and Aim 2 were 2(experiment) x 9(mismatch increment) ANOVAs, and focused on the time-to-target measure because mismatch effects were consistently found using this measure. Hue mismatch effects on time-to-target differed between Experiment 2 and Experiment 4, $F(1,48)=50.96$, $p < .001$, and interacted with the amount of mismatch, $F(8,384)=7.13$, $p < .001$, with larger effects of hue mismatch in Experiment 4 than in Experiment 2—Figure 22A. This is consistent with participants weighting hue heavily in Experiment 4, since hue was valid on most trials, and deweighting hue information in Experiment 2, where hue was invalid on almost all trials. Orientation and aspect ratio were analyzed across the grayscale versions as well as the full-color versions of Experiment 1 (Experiment 3 for aspect ratio) and Experiment 4. For the grayscale versions of the orientation manipulation, there was no reliable difference in orientation mismatch effects between Experiment 1 and 4, $F(1,48)=1.65$, p = .21, nor was there an interaction with the amount of mismatch, $F(8,384)=1.52$, p = .16. Similarly, for the full-color versions, there was no reliable difference between Experiment 1 and 4, $F(1,48)=.52$, $p=.48$, nor was there an interaction, $F(8,384)=.92$, p = *.47*—Figure 22B. For the grayscale versions of the aspect ratio manipulation, there was no reliable difference between Experiments 3 and 4, $F(1,48)=.03$, p = .87, nor an interaction with the amount of mismatch, $F(8,384)=.59$, p = .74 (Figure 22C). For the full-color versions, however, there was again a significant main effect of aspect mismatch, $F(1,48)=6.48$, $p < .05$, and a marginal interaction with amount of mismatch, $F(8,384)=1.86$, p = .08, with a larger cost to guidance in Experiment 4. Taken together, this pattern suggests that subjects were able to change their feature weighting to account for uncertainty in what dimension would change, for hue and—at least in the context of the full-color experiments—for aspect ratio. It is unclear why similar effects were not found in orientation, although this may simply be due to the lower magnitude of mismatch effects in orientation than in aspect ratio. There was a cost to guidance in Experiment 4, presumably due to the greater uncertainty preventing a deweighting of the mismatching feature. The lack of aspect ratio mismatch effects in the grayscale version is similar to the lack of switching costs after orientation- or aspect-change trials, and may also suggest that changes in weighting only occur in the context of color. In Experiments 1 and 3, hue never mismatched, and participants relied on hue information, weighting it heavily. In the full-color version of Experiment 4, because hue sometimes changed, some weight was likely withdrawn from hue, and more weight could be given to aspect ratio or orientation as a result.

A

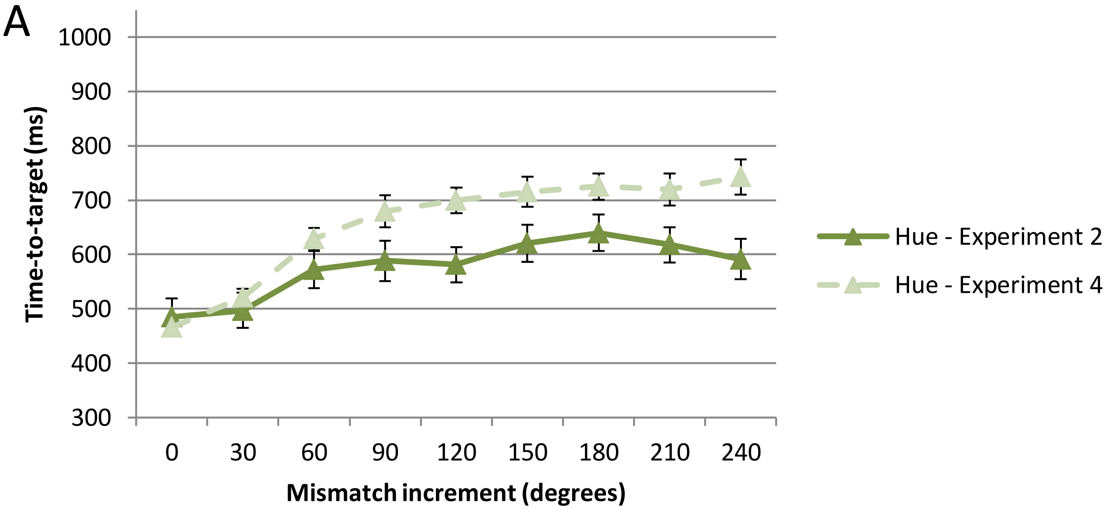

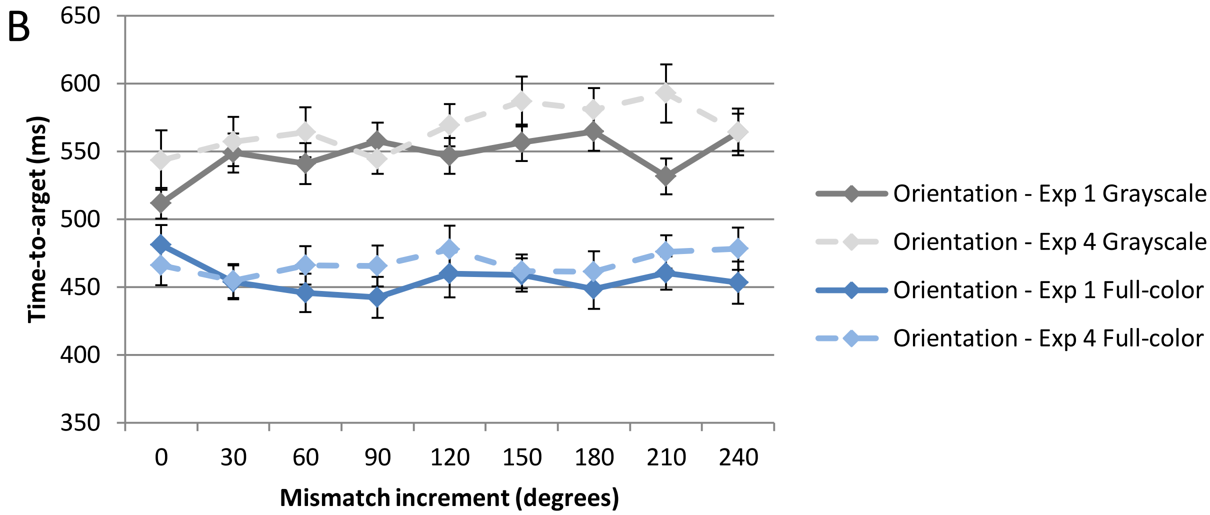

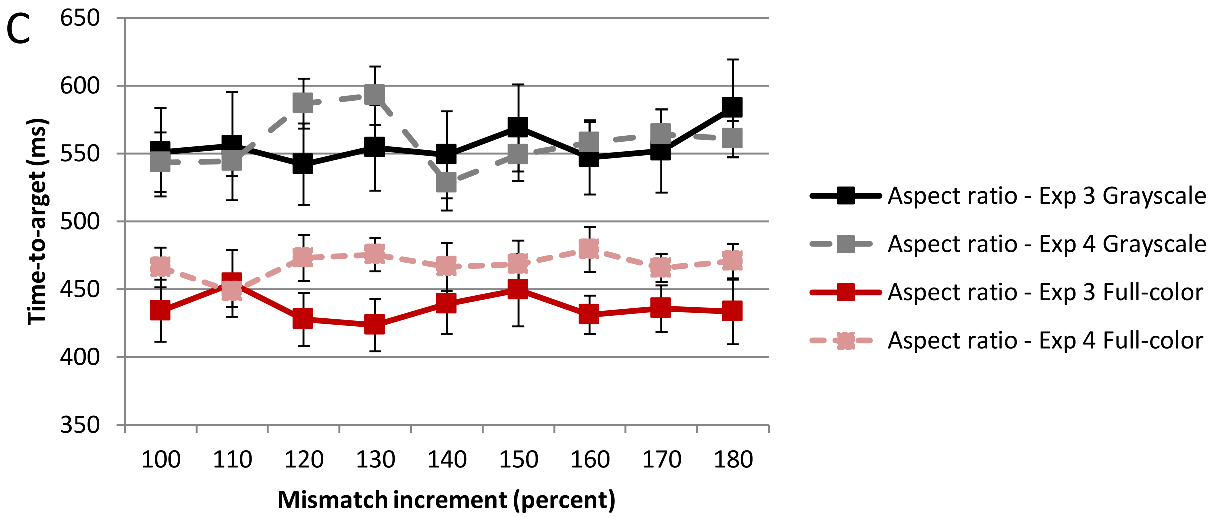

Figure 22. Change in time-to-target across hue mismatch conditions in Experiments 2 and 4 (A), and for both the full-color and grayscale versions of the orientation mismatch conditions in Experiments 1 and 4 (B), and aspect ratio mismatch conditions in Experiments 3 and 4 (C). Error bars indicate one standard error of the mean.

4.3 Experiments for Aim 3

Aim 3 investigated the effects of mismatch between previews and search targets when the mismatch existed on multiple dimensions simultaneously. Two experiments were conducted (one grayscale and one full-color). In the full-color version, hue, orientation, and shape were altered both independently (as baselines) and simultaneously. The grayscale version was identical save that hue-change trials—and trials where hue changed along with orientation or shape—were omitted.

4.3.1 Experiment 5: Simultaneous mismatch across multiple dimensions

The goal of this experiment was to test whether guiding features sum linearly (as predicted by most models of visual search) or not. In this experiment, feature dimensions were manipulated both independently and conjointly. By comparing the effects of feature change within a single dimension to the effects of feature change within several dimensions simultaneously, we can assess whether the combination of guiding features is linear. If the effects on guidance sum linearly, then the mismatch effect on RT (and other measures) in the combined mismatch trials should reflect the sum of the mismatch effects on the single-dimension mismatch trials. If, instead, the mismatch effect on combined mismatch trials is lower than the sum of mismatch effects on single-dimension trials, this would suggest that these feature dimensions are processed partly by the same neurons (Nothdurft, 2000). A higher mismatch effect on combined trials could be the result of objects that differ on multiple feature dimensions being excluded from the set of potential distractors (see Neider and Zelinsky, 2008, for evidence that dissimilar objects may be excluded from consideration).

4.3.1.1 Stimuli/procedure specifications

The full-color version of Experiment 5 included 22 mismatch conditions. The first had no change on any dimension. Three “orientation-change” conditions had only orientation manipulated, with increments of 60° (60°, 120°, and 180°). Three “hue-change” conditions had only hue manipulated, in increments of 60° of hue angle change (60°, 120°, and 180°). Three “shape-change” conditions had only aspect ratio manipulated, with increments of 80% change in width (180%, 260%, and 340% original width, with the accompanying decrease in height). Larger increments of aspect ratio mismatch were used in this experiment in order to magnify the aspect ratio effect, as there was no significant effect of aspect ratio in the full-color version of Experiment 3 and only a marginal effect on “target fixated first” in even the grayscale version of Experiment 3. Unlike hue angle and orientation, aspect ratio is not a circular feature dimension and the range of increments could be increased. Note that fewer levels of change for each individual feature dimension were used to accommodate additional conditions without creating a prohibitive number of trials. Three more conditions were included for each combination of hue, orientation, and shape manipulations: “Hue*aspect”, “orientation*hue”, “orientation*aspect”, and “orientation*hue*aspect*”. For these combined conditions, the same increments of change were used as for the hue-change, shape-change, and orientation-change conditions but were applied on both dimensions (or all three, in the case of orientation*hue*aspect). For example, the three orientation*aspect conditions were (1) stretched to 180% and rotated 60°, (2) stretched 260% and rotated 120°, and (3) stretched 340% and rotated 180°—see Figure 23. The grayscale version of the experiment had ten conditions: The no change condition, the three orientation conditions, the three shape conditions, and the three orientation*aspect conditions.

No change	Orientation	Hue	Shape	Orientation *Hue	Hue *Aspect	Orientation *Aspect	Orientation*Hue *Aspect
	60° 	60° 	180% 	60°x60° 	60°180% 	60°x180% 	60°x60°x180% 
	120° 	120° 	260% 	120°x120° 	120°x260% 	120°x260% 	120°x120°x260% 
	180° 	180° 	340% 	180°x180° 	180°x340% 	180°x340% 	180°x180°x340% 

Figure 23. Example stimuli, at each preview-target mismatch condition for Experiment 5.

Twenty-five experimental trials for each condition and 22 practice trials (one for each condition) were included, for a total of 572 trials. Because there were fewer trials in Experiment 5 than in the first four experiments, the full list of trials (images and x,y coordinates) from Experiments 1-4 could not be used in Experiment 5. Instead, the last 3 practice trials and the last 75 experimental trials were removed from the generated list used for the earlier experiments. The grayscale version was created by removing all hue-change trials and converting the stimuli to grayscale. As such, there were only 10 practice trials and 250 experimental trials in the grayscale version. In all other respects the procedure was the same as described in the general procedure (section 3.4).

4.3.1.2 Results and discussion for Experiment 5

4.3.1.2.1 Results and discussion for the full-color version of Experiment 5

The results below all refer to the full-color version of the experiment. Results for the grayscale version are provided below, in section 4.3.1.2.2. Except where otherwise noted, all analyses in this section are 3(hue/shape/orientation) x 4(mismatch increments) ANOVAs, with Bonferroni-corrected post hoc *t*-tests.

Manual reaction times and accuracy

Consistent with the results of Experiment 2, RTs increased with increasing mismatch—$F(3,63)=51.31$, $p < .001$, with the largest effects in hue rather than aspect ratio or orientation, as confirmed by a main effect of feature dimension, $F(2,42)=910.56$, $p <.001$, and an interaction between feature dimension and the amount of mismatch, $F(6,126)=346.93$, $p <.001$—see Figure 24. Post hoc tests revealed significant differences between all three feature dimensions (*p*s < .001).

If mismatch effects reflect the linear summation of feature dimensions, then the combined mismatch trials should have equivalent mismatch effects to the sum of each mismatch effect on the component single-mismatch trials. For example, the expected mismatch effect for the hue*aspect condition at one mismatch increment (60° for hue and 180% for aspect ratio) would be the sum of the difference between 60° hue mismatch condition and the "no change" condition and the difference between the 180% aspect ratio condition and the "no change" condition. These expected values were computed for each subject independently, so that variance between subjects could be included in statistical comparisons (and be shown as error bars in the accompanying figures). If the observed values for the combined mismatch trials differ from the expected values, this would reflect the involvement of a nonlinear process, such as the sharing of some neurons between different feature dimensions (which would result in a subadditive mismatch effect) or excluding objects that differ on multiple dimensions from set of potential distractors (which would result in a superadditive mismatch effect). To test this, RTs for the combined mismatch conditions were compared to the expected values predicted by linear summation in that condition.

Aside from the "orientation*aspect" condition, which did not significantly differ from the expected values predicted by linear summation in that condition (Figure 24B)—$F(3,63)=.80$, $p = .48$—all conditions combining mismatch on multiple feature dimensions ("hue*aspect", "orientation*hue", and "orientation*hue*aspect") had significantly larger RT mismatch effects than predicted by linear summation, all Fs≥ 4.51, $p < .01$, see Figure 24A, C, and D. This pattern is consistent with a process excluding objects with multiple dimensions of mismatch from the potential distractor set, though only when one of the dimensions is hue. This may result from a floor effect in the effects of aspect ratio and orientation, (see Figure 24B), or may reflect the relative dominance of hue in search: Even when both aspect ratio and orientation differ from the preview, a match on the hue dimension may override any attempt to reject this object from the set. Accuracy decreased with increasing mismatch increment—$F(3,63)=3.17$, $p < .05$—and differed across feature dimension, $F(2,42)=3.93$, $p < .05$), though accuracy never fell below 94%. Mismatch increment and feature dimension trended towards an interaction, $F(6,126)=2.11$, $p = .08$), with a larger decrease in accuracy in hue-change trials. This suggests that the RT effects were not due to a speed-accuracy trade-off, as accuracy decreased as RTs increased. Error trials (including timed-out trials) were removed from all other analyses (including the RT analysis).

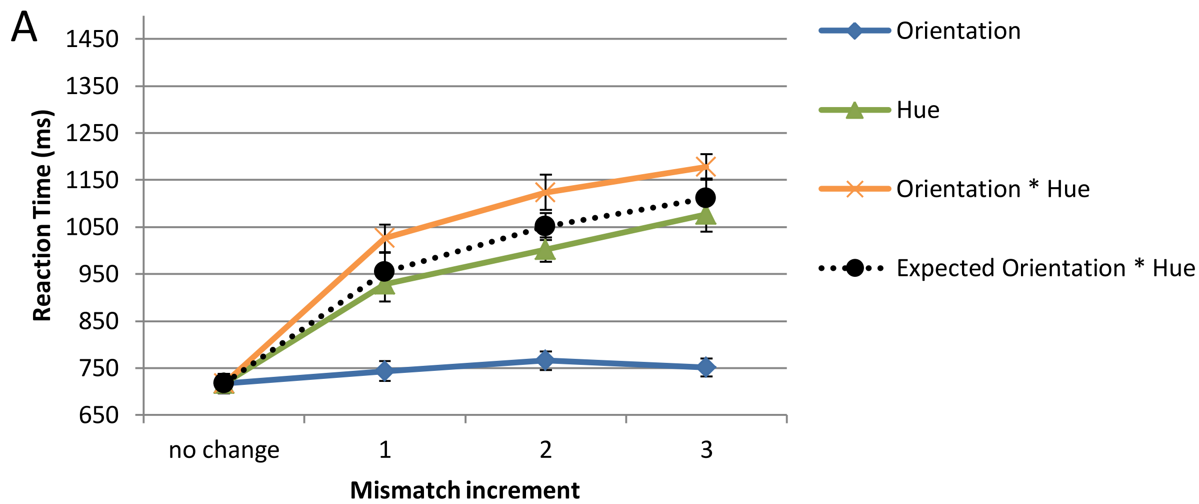
A
1450
1350
1250
1150
1050
950
850
750
650
Reaction Time (ms)
no change
1
2
3
Mismatch increment
Orientation
Hue
Orientation * Hue
Expected Orientation * Hue

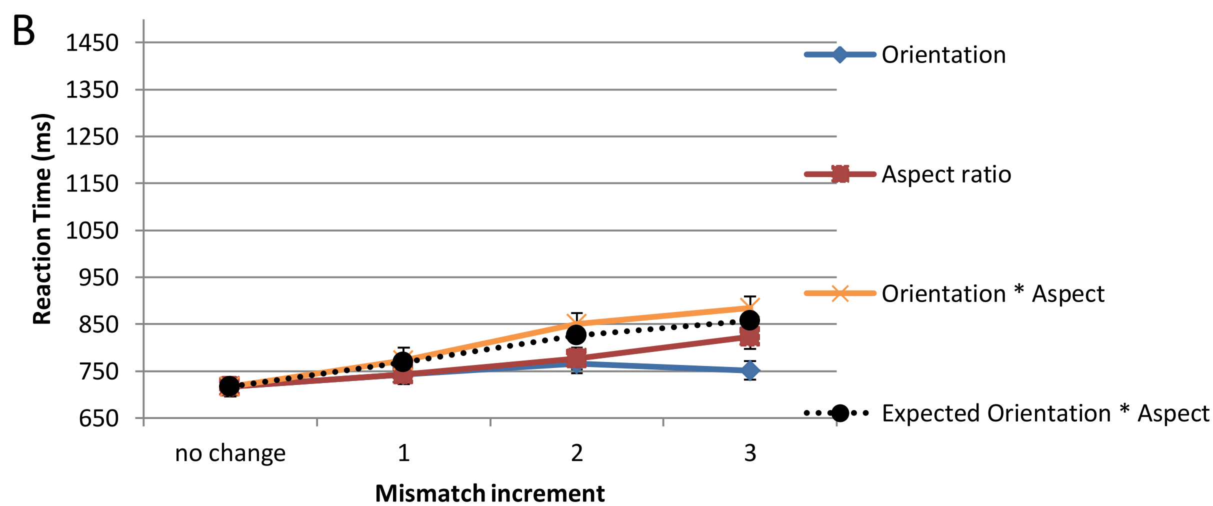
B
1450
1350
1250
1150
1050
950
850
750
650
Reaction Time (ms)
no change
1
2
3
Mismatch increment
Orientation
Aspect ratio
Orientation * Aspect
Expected Orientation * Aspect

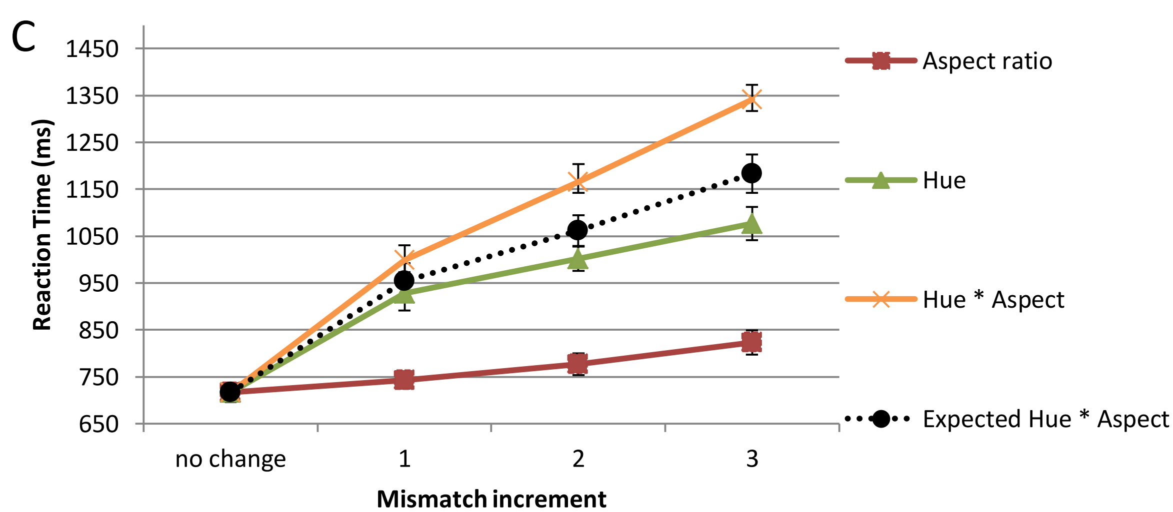
C
1450
1350
1250
1150
1050
950
850
750
650
Reaction Time (ms)
no change
1
2
3
Mismatch increment
Aspect ratio
Hue
Hue * Aspect
Expected Hue * Aspect

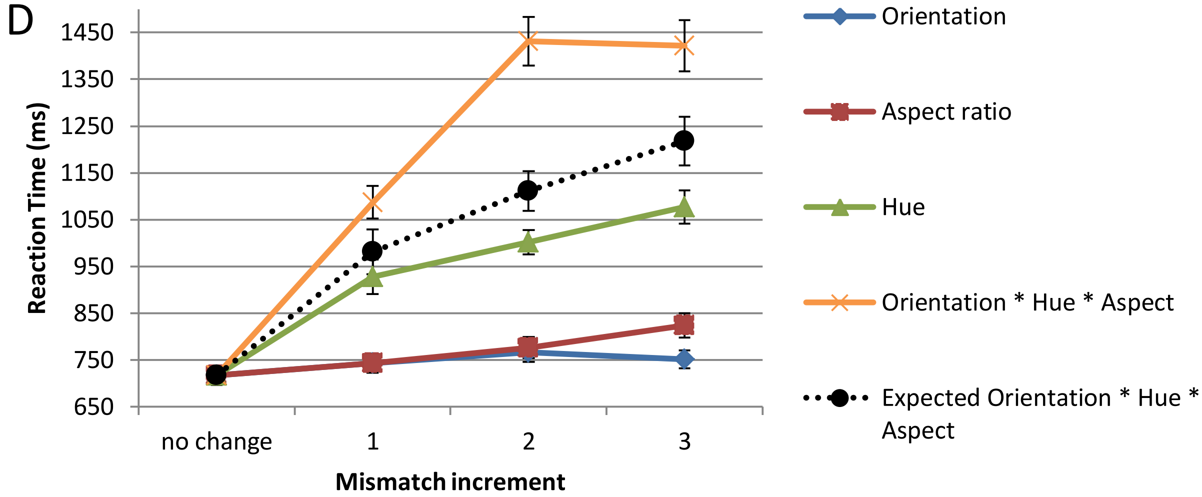

Figure 24. Change in RTs across mismatch conditions in the full-color version of Experiment 5. (A) Comparing orientation, hue, and orientation*hue conditions with the expected orientation*hue pattern that would be found if signals were summed linearly. (B) Comparing orientation, aspect ratio, and orientation*aspect conditions with the expected orientation*aspect pattern that would be found if signals were summed linearly. (C) Comparing hue, aspect ratio, and hue*aspect conditions with the expected hue*aspect pattern that would be found if signals were summed linearly. (D) Comparing the combination of all three feature dimensions (hue, aspect ratio, and orientation) with the expected pattern that would be found if signals summed linearly. Error bars indicate one standard error of the mean.

Guidance

Feature dimension and degree of mismatch significantly interacted to affect time-to-target—$F(6,132) = 22.59$, $p < .001$—with the largest effects at higher mismatches for the hue dimension—See Figure 25A. Main effects of mismatch—$F(3,66) = 33.73$, $p < .001$—and feature dimension—$F(2,44) = 83.36$, $p < .001$—were also found. If the feature maps used to guide search sum linearly, there should be no significant interaction between the combined feature manipulations and the sum of the mismatch effects when only the individual dimensions are manipulated. However, every condition combining mismatch on multiple feature dimensions (hue*aspect, orientation*hue, orientation*aspect, and orientation*hue*aspect) had larger mismatch effects than predicted by linear summation, as revealed by significant interactions between the conditions and their expected values, all Fs$\geq$9.46, $p < .001$, see all four panels of Figure 25. This pattern suggests that guidance discounts objects that differ on multiple dimensions.

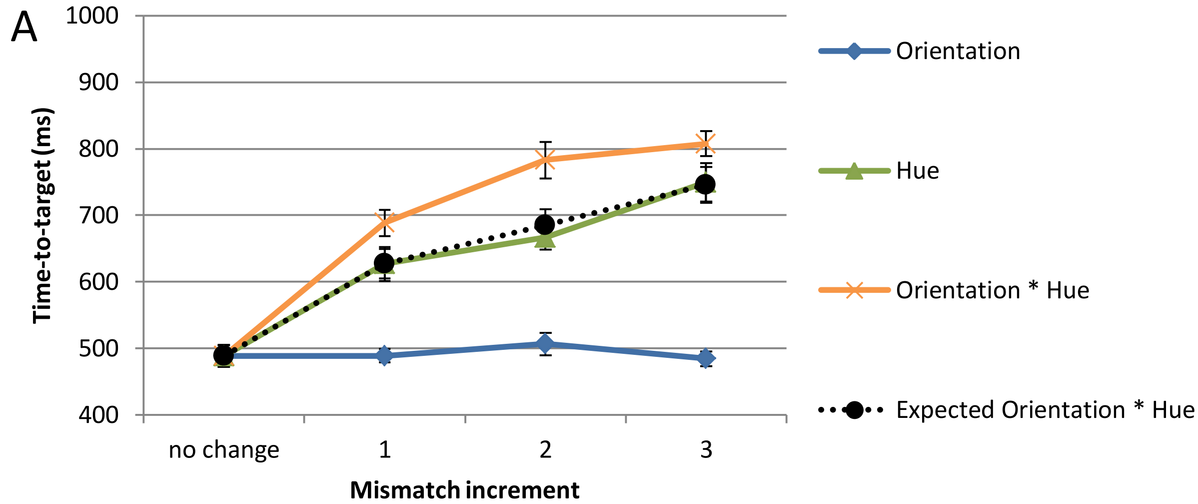
A
Time-to-target (ms)
1000
900
800
700
600
500
400
no change
1
2
3
Mismatch increment
Orientation
Hue
Orientation * Hue
Expected Orientation * Hue

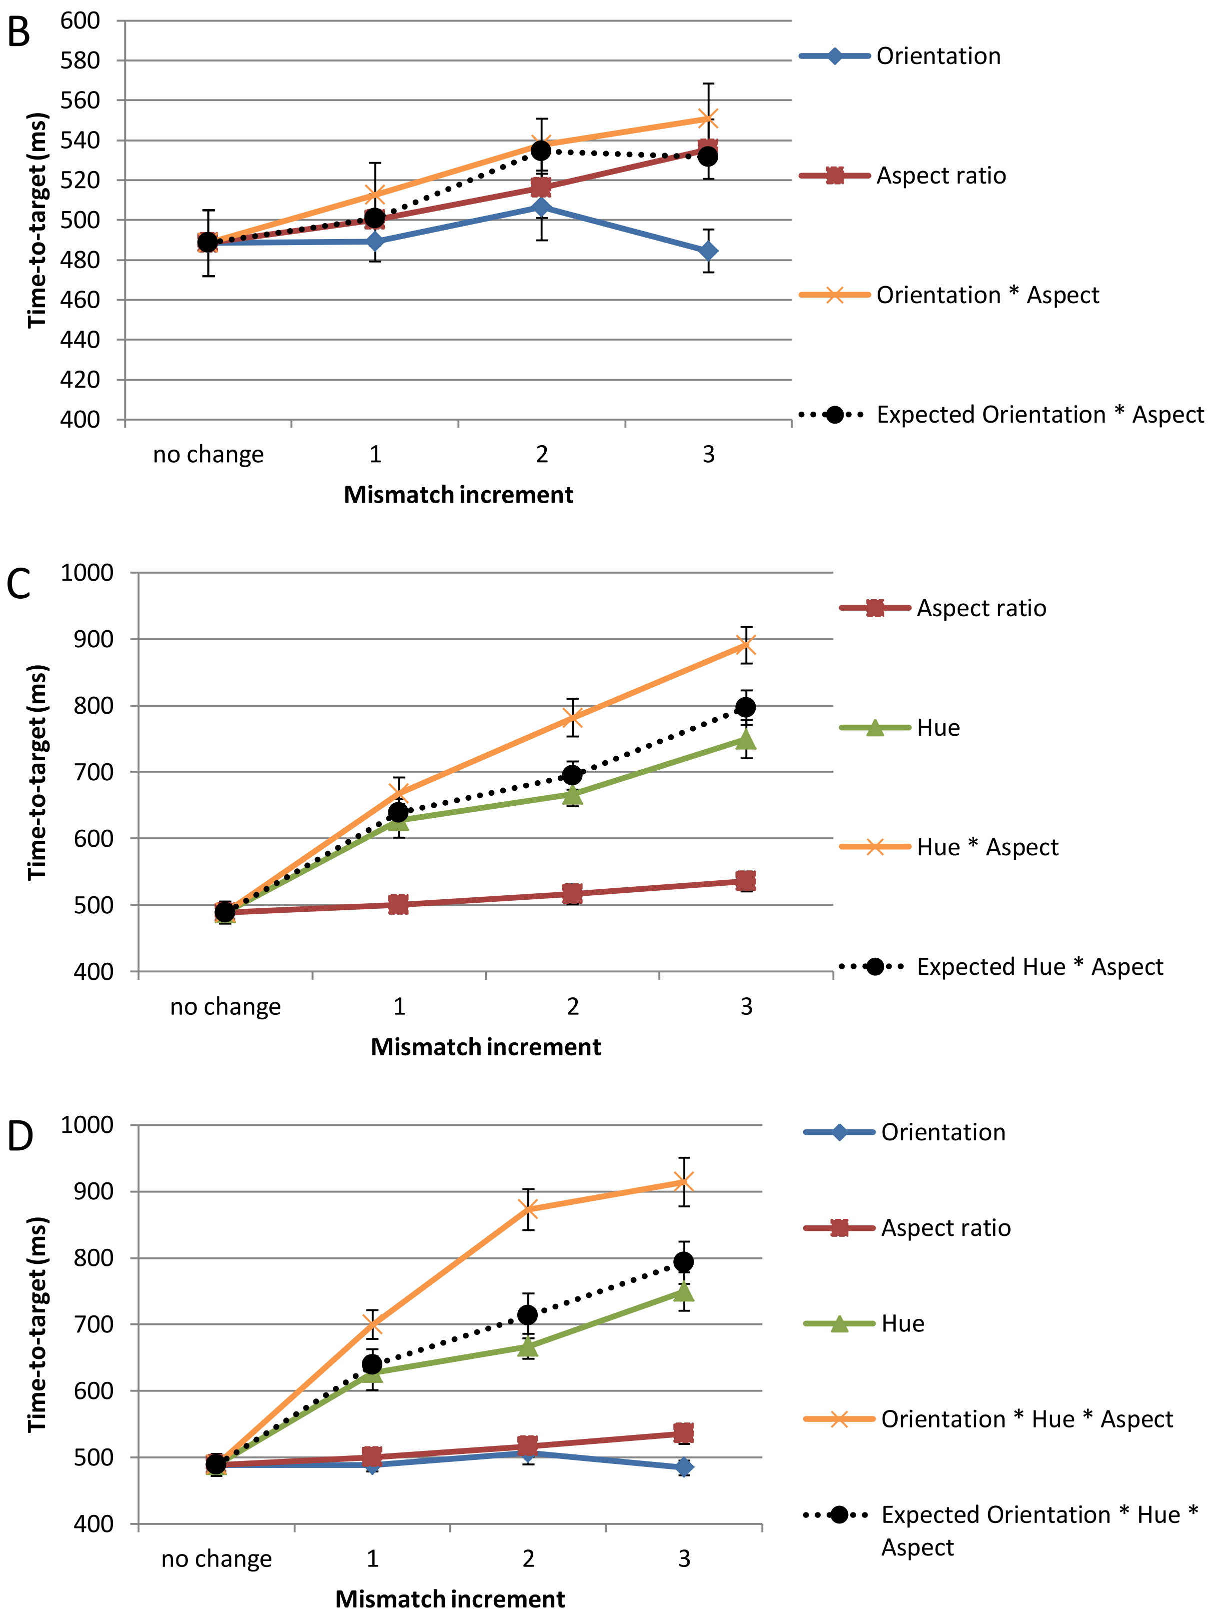
B
600
580
560
540
520
500
480
460
440
420
400
Time-to-target (ms)
no change
1
2
3
Mismatch increment
Orientation
Aspect ratio
Orientation * Aspect
Expected Orientation * Aspect
C
1000
900
800
700
600
500
400
Time-to-target (ms)
no change
1
2
3
Mismatch increment
Aspect ratio
Hue
Hue * Aspect
Expected Hue * Aspect
D
1000
900
800
700
600
500
400
Time-to-target (ms)
no change
1
2
3
Mismatch increment
Orientation
Aspect ratio
Hue
Orientation * Hue * Aspect
Expected Orientation * Hue * Aspect

Figure 25. Change in time-to-target across mismatch conditions in the full-color version of Experiment 5. (A) Comparing orientation, hue, and orientation*hue conditions with the expected orientation*hue pattern that would be found if signals were summed linearly. (B) Comparing orientation, aspect ratio, and orientation*aspect conditions with the expected orientation*aspect pattern that would be found if signals were summed linearly. Note that panel B uses a smaller scale than the other panels in this figure. (C) Comparing hue, aspect ratio, and hue*aspect conditions with the expected hue*aspect pattern that would be found if signals were summed linearly. (D) Comparing the combination of all three feature dimensions (hue, aspect ratio, and orientation) with the expected pattern that would be found if signals summed linearly. Error bars indicate one standard error of the mean.

In terms of whether the target was the first object fixated, all of the combined conditions did not reliably differ from the expected values if feature mismatch is linearly summed ($p \geq .14$; Figure 26A, B, C, and D). This pattern suggests that feature dimensions are summed linearly, possibly indicating that there is an early, linear process (affecting the first few fixations), followed by a process that results in non-linearities in later measures of guidance. For example, if participants exclude objects that differ on multiple dimensions from the set of potential distractors, this may not occur quickly enough to affect the first few eye movements. Caution should be taken in interpreting several of these findings, however, because the largest mismatch increment of hue*aspect dropped to chance levels of guidance ($p = .1$), as did the two largest mismatch increments in the hue*orientation*aspect condition (both $p > .4$). All other conditions were significantly above chance ($p < .001$).

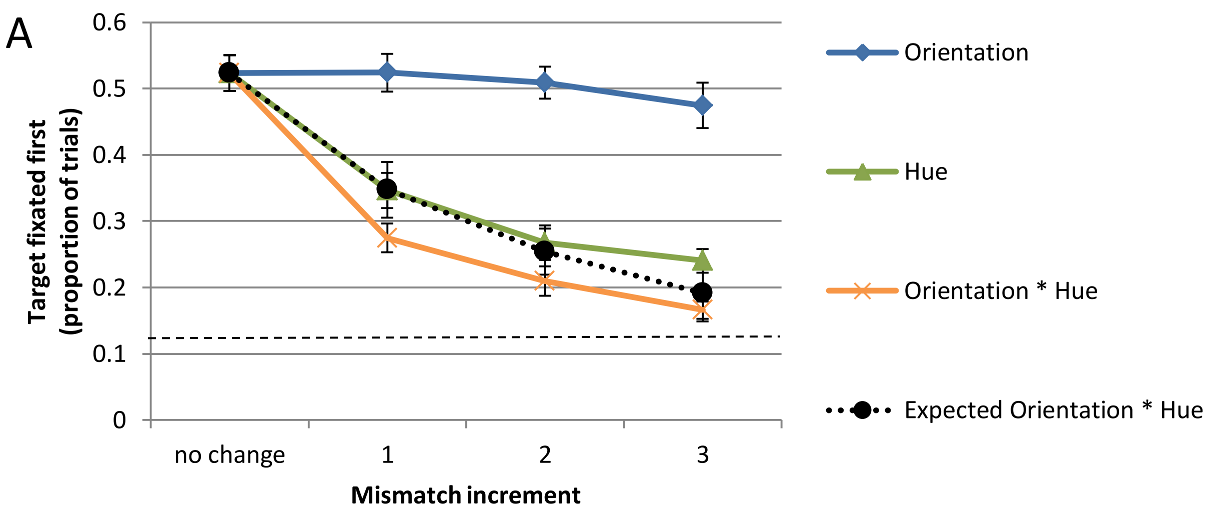

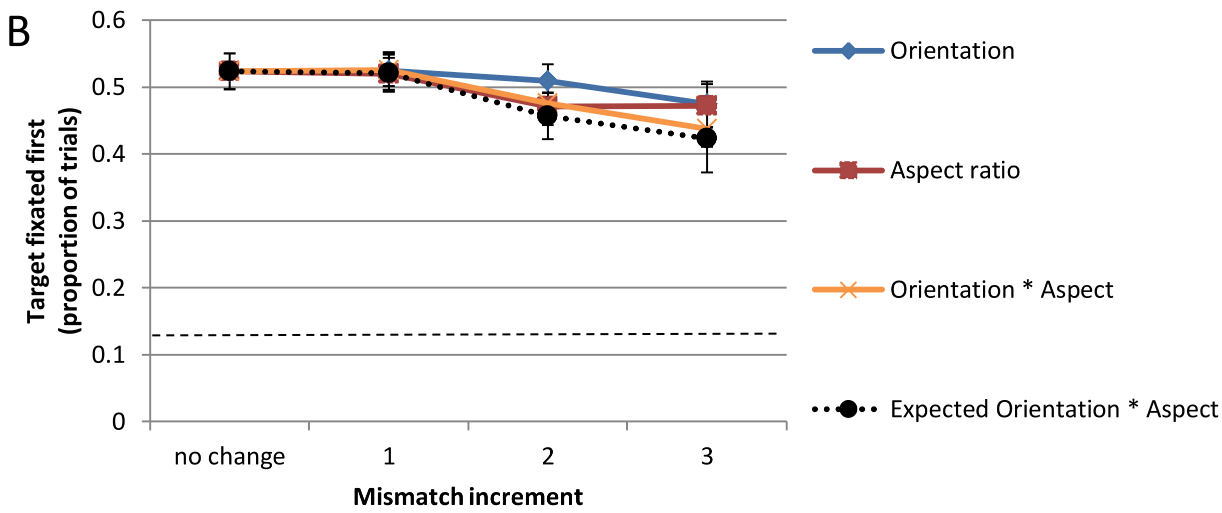
B
0.6
0.5
0.4
0.3
0.2
0.1
0
Target fixated first (proportion of trials)
no change
1
2
3
Mismatch increment
Orientation
Aspect ratio
Orientation * Aspect
Expected Orientation * Aspect

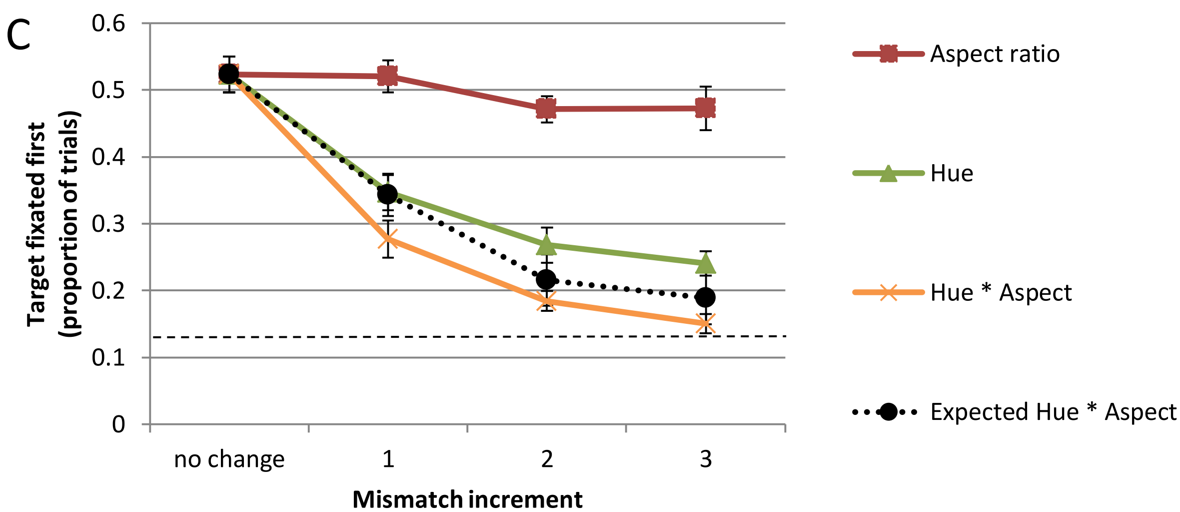
C
0.6
0.5
0.4
0.3
0.2
0.1
0
Target fixated first (proportion of trials)
no change
1
2
3
Mismatch increment
Aspect ratio
Hue
Hue * Aspect
Expected Hue * Aspect

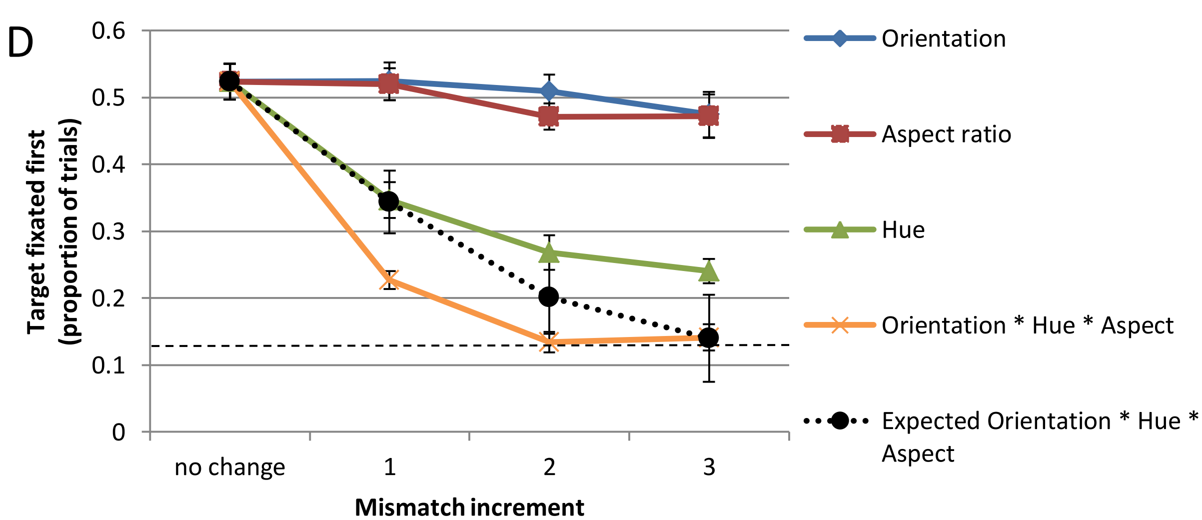
D
0.6
0.5
0.4
0.3
0.2
0.1
0
Target fixated first (proportion of trials)
no change
1
2
3
Mismatch increment
Orientation
Aspect ratio
Hue
Orientation * Hue * Aspect
Expected Orientation * Hue * Aspect

Figure 26. Change in target first fixated across mismatch conditions in the full-color version of Experiment 5. (A) Comparing orientation, hue, and orientation*hue conditions with the expected orientation*hue pattern that would be found if signals were summed linearly. (B) Comparing orientation, aspect ratio, and orientation*aspect conditions with the expected orientation*aspect pattern that would be found if signals were summed linearly. (C) Comparing hue, aspect ratio, and hue*aspect conditions with the expected hue*aspect pattern that would be found if signals were summed linearly. (D) Comparing the combination of all three feature dimensions (hue, aspect ratio, and orientation) with the expected pattern that would be found if signals summed linearly. The dashed lines indicate chance levels of performance (0.125). Error bars indicate one standard error of the mean.

To test for speed-accuracy tradeoffs that might bias the early guidance measure (target first fixated), one-way ANOVAs were conducted to compare the effects of each increment of feature change on the time it took subjects to first fixate *any* object (target or distractor) after display onset. However, this was not the case: the time to fixate any object *increased* with mismatch increment—$F(3,63) = 2.45$, $p = .10$—and interacted with feature dimension—$F(6,126)=6.62$, $p < .001$—such that initial saccade latencies were *longer* in conditions with poorer guidance, the opposite of what would be predicted from a speed-accuracy tradeoff.

Verification times

Verification times were again affected by feature mismatch. Though there was no reliable main effect of amount of mismatch—$F(3,63)=1.66$, $p = .21$—there was a main effect of feature dimension—$F(2,42)=26.25$, $p < .001$—and feature dimension interacted with amount of mismatch—$F(6,126)=4.15$, $p < .01$—with the largest effect in the hue condition. Hue mismatch affected performance significantly more than orientation mismatch or aspect ratio mismatch ($p < .001$) and no difference was found between orientation and aspect ratio mismatch conditions ($p = 0.30$). Verification times for each of the combined mismatch conditions significantly differed from the values expected by linear summation, all *p*s < .05, see Figure 27.

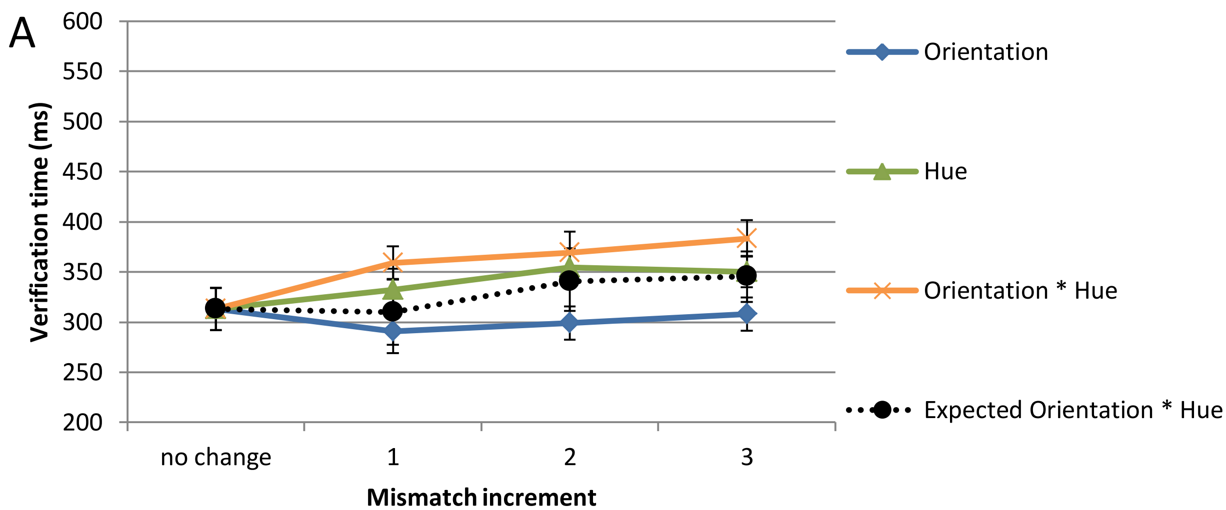

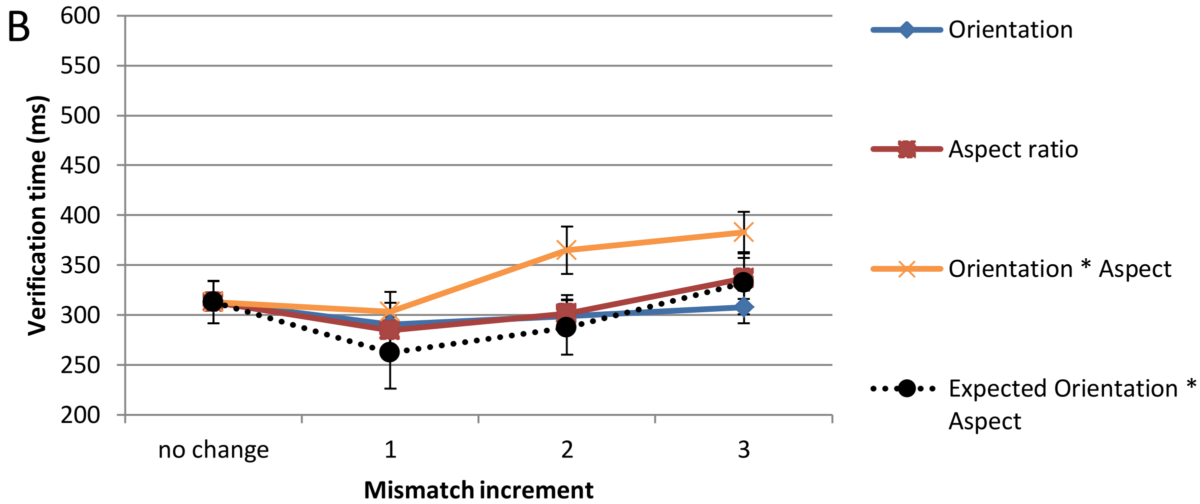

B
600
550
500
450
400
350
300
250
200
Verification time (ms)
no change
1
2
3
Mismatch increment
Orientation
Aspect ratio
Orientation * Aspect
Expected Orientation * Aspect

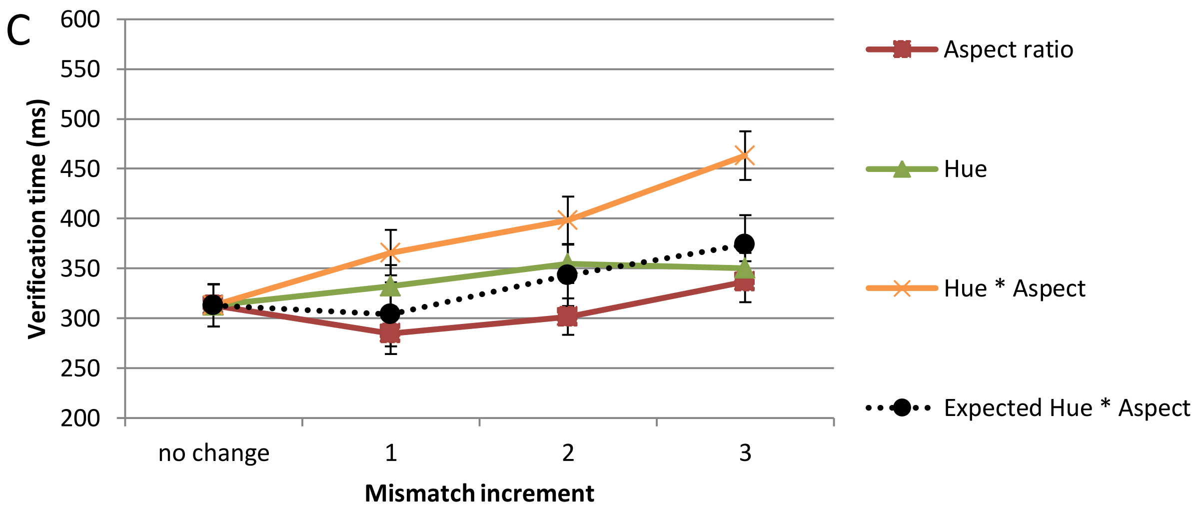

C
600
550
500
450
400
350
300
250
200
Verification time (ms)
no change
1
2
3
Mismatch increment
Aspect ratio
Hue
Hue * Aspect
Expected Hue * Aspect

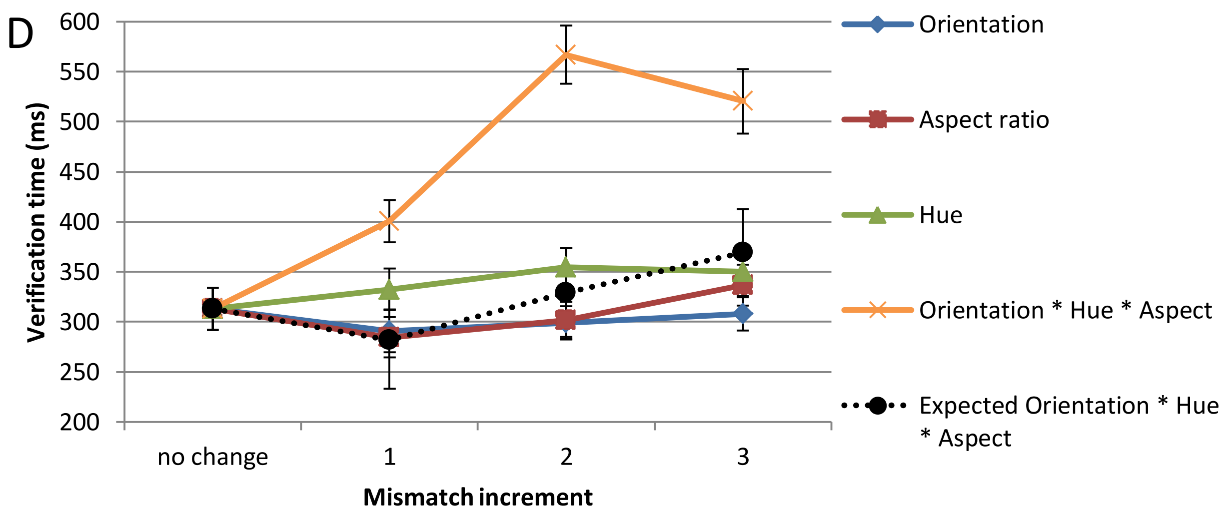

D
600
550
500
450
400
350
300
250
200
Verification time (ms)
no change
1
2
3
Mismatch increment
Orientation
Aspect ratio
Hue
Orientation * Hue * Aspect
Expected Orientation * Hue * Aspect

Figure 27. Change in verification time across mismatch conditions in the full-color version of Experiment 5. (A) Comparing orientation, hue, and orientation*hue conditions with the expected orientation*hue pattern that would be found if signals were summed linearly. (B) Comparing orientation, aspect ratio, and orientation*aspect conditions with the expected orientation*aspect pattern that would be found if signals were summed linearly. (C) Comparing hue, aspect ratio, and hue*aspect conditions with the expected hue*aspect pattern that would be found if signals were summed linearly. (D) Comparing the combination of all three feature dimensions (hue, aspect ratio, and orientation) with the expected pattern that would be found if signals summed linearly. Error bars indicate one standard error of the mean.

The importance of individual feature dimensions for search

To test for relative levels of importance for hue and orientation, paired samples *t*-tests comparing the mean and maximum effect of each DV for trials where hue changed were compared to trials where orientation changed. As in Aims 1 and 2, aspect ratio was excluded from the tests comparing mean and maximum effects, but was included in all other analyses. Consistent with the results of Aims 1 and 2, and with previous research (e.g. Pomplun, 2006; Williams, 1966), hue mismatch caused significantly worse guidance than orientation mismatch, both in terms of a higher maximum time-to-target ($t(21)=8.24$, $p < .001$) and a lower minimum proportion of trials where the target was the first object fixated ($t(21)=7.45$, $p < .001$). Unlike Aim 1, however, significant differences were also found in verification time and accuracy, though unlike Aim 2 these were not in the direction of a speed-accuracy tradeoff: Verification times were longer ($t(21)=3.34$, $p < .01$) and accuracy was lower for hue-change trials ($t(21)=4.92$, $p < .001$) than for orientation trials. This pattern suggests that, in the context of Aim 3 where multiple feature dimensions sometimes mismatch simultaneously, hue has an impact on verification processes. Participants may be less confident about target identity when multiple dimensions change, which may have allowed an effect in hue to emerge.

4.3.1.2.2 Results and discussion for the grayscale version of Experiment 5

The results below all refer to the grayscale version of the experiment. Except where otherwise noted, all analyses for single-dimension-change trials in this section are 2(shape/orientation) x 4(mismatch increments) ANOVAs, with Bonferroni-corrected post hoc *t*-tests. Two(orientation*aspect trials/expected orientation*aspect values) x 4(mismatch increments) ANOVAs were used to compare combined orientation*aspect-change trials to the expected orientation*aspect values predicted by a linear summation of mismatch effects from the aspect and orientation conditions.

Manual reaction times and accuracy

To the extent that subjects rely on orientation or aspect ratio in order to perform the task, RTs were expected to differ across mismatch conditions. Consistent with this prediction, as well as the grayscale versions of Experiments 1 and 3, RTs indeed increased as mismatch increased—$F(3,63) = 18.50$, $p < .001$. Post-hoc contrasts revealed that increasing mismatch had a linear effect on RTs, $F(1,21)=25.65$, $p < .001$. Although no main effect of feature dimension was found, $F(1,21) = 1.61$, $p = .22$, mismatch and dimension significantly interacted, $F(3,63)=3.33$, $p < .05$. RTs in the combined orientation*aspect condition did not reliably differ from the expected values for that condition, $F(1,21)=2.22$, $p = .15$, Figure 28, nor did they interact with mismatch increment, $F(2,42)=2.48$, $p = .10$, suggesting that these mismatch effects may combine linearly, consistent with the full-color version of the experiment in which all combined trials *except for* orientation*aspect showed deviations from linearity (Figure 24B).

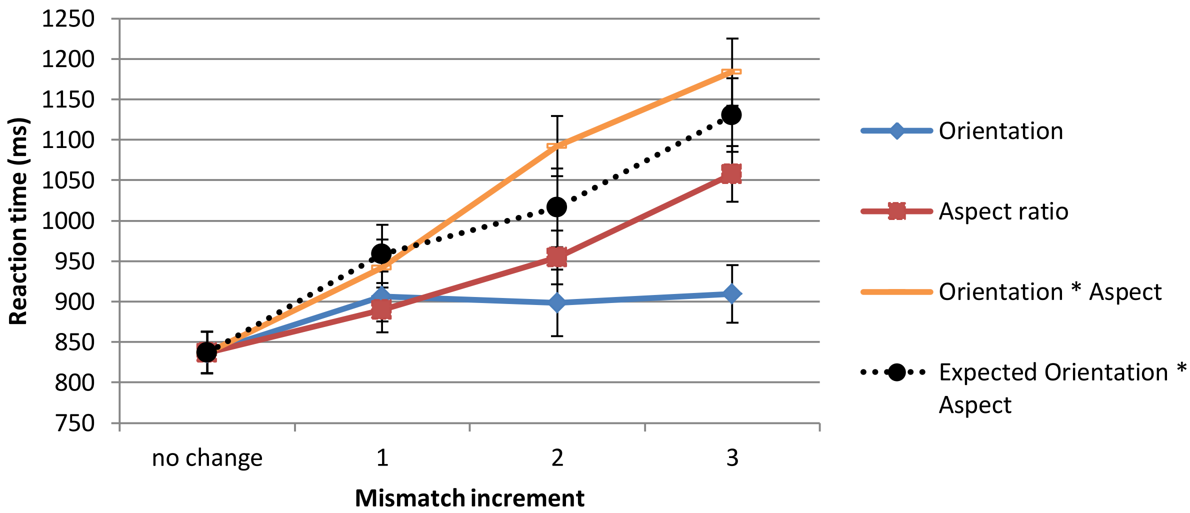

Figure 28. Change in RTs across mismatch conditions in the grayscale version of Experiment 5. Error bars indicate one standard error of the mean.

Increasing mismatch resulted in worse accuracy, $F(3,63)=5.99$, $p < .01$, though orientation and shape mismatch did not differentially affect accuracy, $F(1,21) = .52$, $p = .48$, nor did it interact with mismatch increment, F(3,63)=2.13, p = .11. This indicates that the difference between orientation and aspect ratio mismatch effects in RTs was not due to a speed-accuracy trade-off. The combined orientation*aspect trials, however, had lower accuracy than expected by the combination of mismatch effects when only one dimension was changed, $F(1,21)=11.87$, $p < .01$, suggesting that when targets change on multiple dimensions, it becomes more difficult to identify them as targets. Error trials (including timed-out trials) were removed from all other analyses (including the RT analysis).

Guidance

Consistent with the grayscale versions of Experiments 1, 3, and 4, time-to-target increased with increasing mismatch in orientation and aspect ratio, $F(3,63)=9.97$, $p < .001$. Orientation and aspect ratio had marginally different effects, $F(1,21)=3.91$, $p = .06$, and feature dimension trended towards an interaction with amount of mismatch—$F(3,63) = 2.45$, $p = .09$—with the largest effects at higher mismatches for the aspect ratio dimension, consistent with Experiment 4—See Figure 29. Consistent with overall RTs, the orientation*aspect trials did not reliably differ from the expected orientation*aspect values, $F(1,21)=.28$, $p = .60$, further suggesting that a linear summation of feature dimensions might account for mismatch effects in search guidance when hue mismatch is not present. However, time-to-target can be affected by processes other than guidance (for example, the rejection of distractors), and more conservative measures (such as whether the target was the first object fixated) are necessary to ensure that this pattern holds true for guidance *per se*.

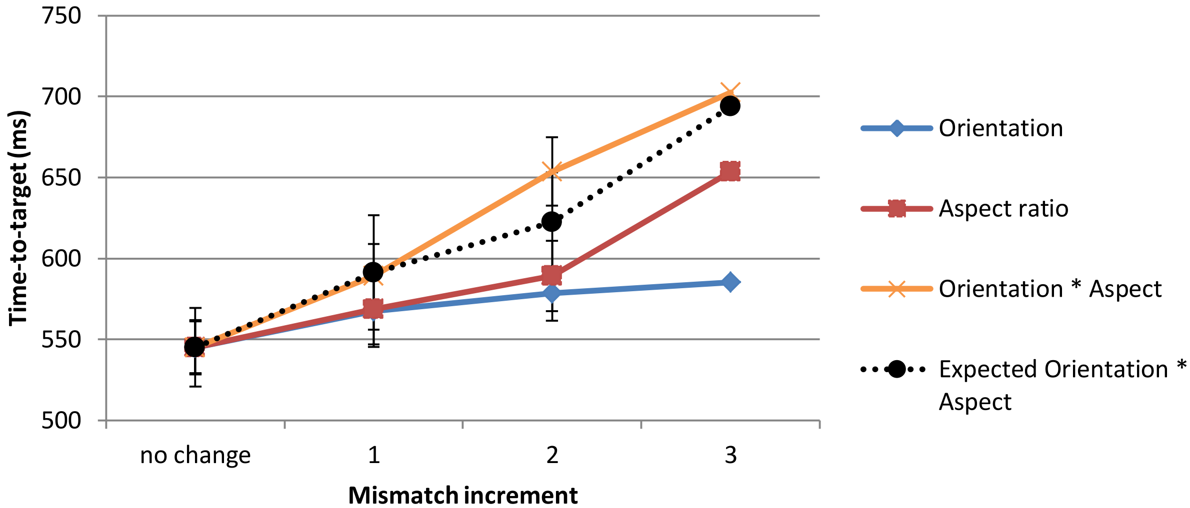

Figure 29. Change in time-to-target across mismatch conditions in the grayscale version of Experiment 5. Error bars indicate one standard error of the mean.

With increasing mismatch, the target was less often fixated first, $F(3,63)=8.51$, $p < .001$. No main effect of feature dimension was found, $F(1,21)=2.28$, $p = .15$, but feature dimension marginally interacted with amount of mismatch—$F(3,63) = 2.59$, $p = .07$—with the largest effects at higher mismatches for the aspect ratio dimension, consistent with time-to-target and with the results of Experiment 4—See Figure 30. The orientation*aspect trials did not reliably differ from the expected orientation*aspect values, $F(1,21)=.28$, $p = .60$, again suggesting that a linear summation of feature dimensions might account for mismatch effects in search guidance, as predicted by most models of visual search. No conditions dropped to chance levels of guidance, all *p*s < .001.

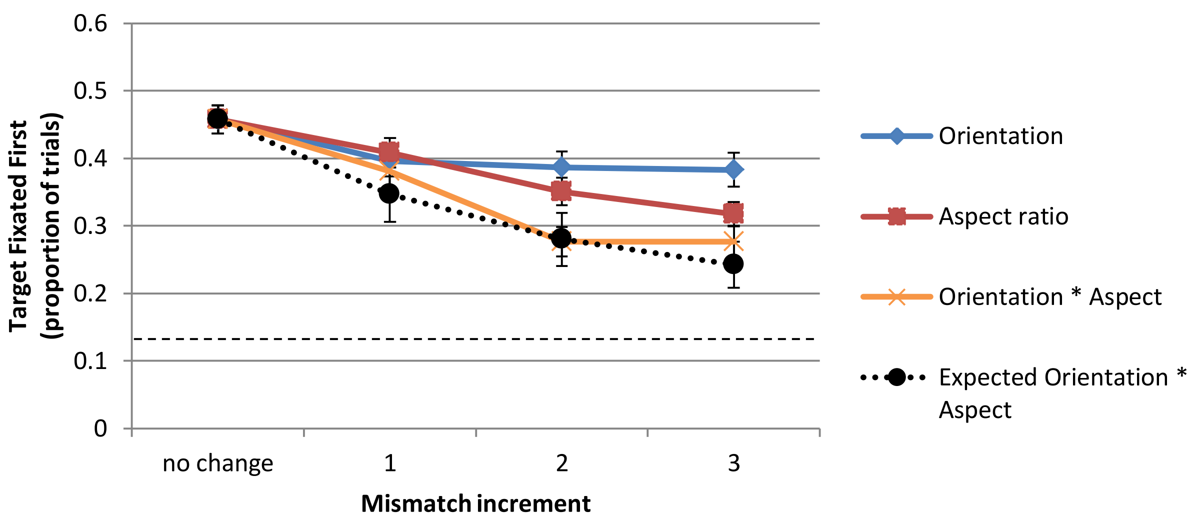

Figure 30. Change in the proportion of trials where the target was fixated first across mismatch conditions in the grayscale version of Experiment 5. The dashed line indicates chance. Error bars indicate one standard error of the mean.

To test for speed-accuracy tradeoffs that might bias the target first fixated measure, conditions were compared on the time it takes subjects to first fixate *any* object (target or distractor) after display onset. However, the time to fixate any object did not vary with mismatch or feature dimension, nor did the conditions interact ($ps \geq .21$). Similarly, the time to fixate any object did not differ between the orientation*aspect trials and expected orientation*aspect values, nor across different levels of mismatch, both $p \geq .39$.

Verification times

Verification times were significantly affected by amount of mismatch—$F(3,63) = 7.11$, $p < .01$, Figure 31—indicating that orientation and shape were used in responding to the target in this task. Feature dimension reliably affected verification time ($F(3,63) = 20.28$, $p < .001$), and feature dimension interacted with amount of mismatch, $F(3,63)=6.46$, $p < .01$. Verification times did not differ between the orientation*aspect trials and the expected orientation*aspect values, $F(1,21)=.08$, $p = .78$, nor did the mismatch effect in those conditions interact with the degree of mismatch, $F(3,63)=1.36$, $p = .27$, consisted with a linear summation of the orientation and aspect dimensions.

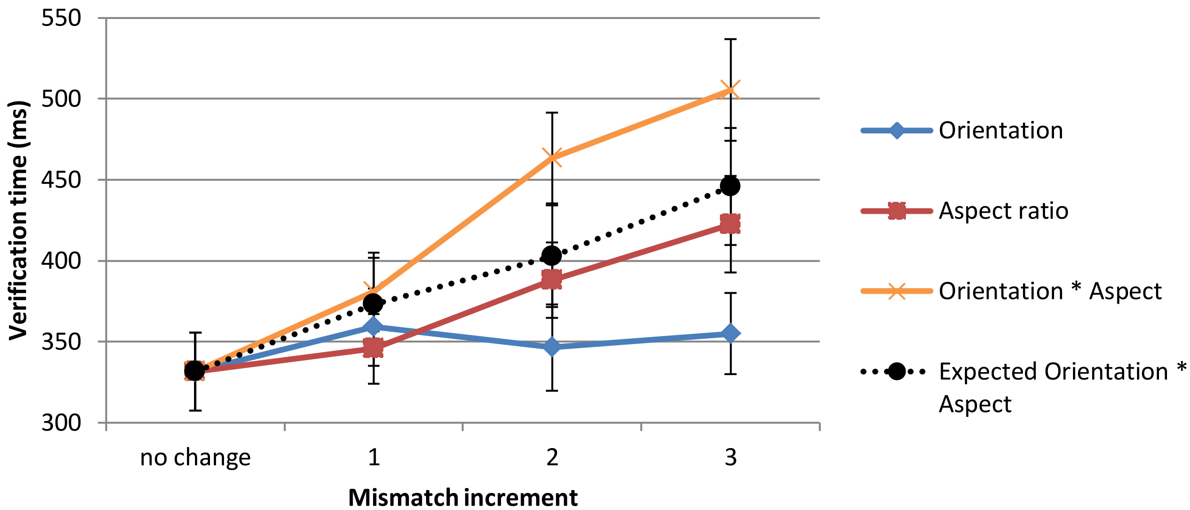

Figure 31. Change in verification time across mismatch conditions in the grayscale version of Experiment 5. Error bars indicate one standard error of the mean.

Chapter 5 - Conclusions, unresolved questions, and implications for future work

The goals of this project were motivated by the need for a clear understanding of what features are used to guide visual search with real-world stimuli. Most previous work exploring the target template has used simple stimuli and has focused on determining what features *can* be used to guide search, and what features are preattentively available to the visual system. It is likely, however, that different features might be used with real-world stimuli and—importantly—the features which *can* guide search are not necessarily those that *are* used to guide search. If, for example, participants were capable of guiding search efficiently to an angry face among neutral faces, this would suggest that emotional expressions are available (at least to some degree) preattentively, and can be used to guide search. However, it is extremely unlikely that facial expressions are a feature which is included in the target template for most search tasks (for example, facial expressions are useless for finding a pen on your desk or your car in a parking lot). The present experiments tested which features *are* used by allowing the targets to differ in many dimensions from distractor items, a more natural situation than tasks where targets are defined along a single dimension. The rise of image-based models of visual search (e.g. Itti & Koch, 2000; Navalpakkam & Itti, 2005; Oliva, Torralba, Castelhano, & Henderson, 2003; Parkhurst, Law, & Niebur, 2002; Pomplun, 2006; Rao, Zelinsky, Hayhoe, & Ballard, 2002; Torralba, Oliva, Castelhano, & Henderson, 2006; Zelinsky, 2008) has created a pressing need to formalize what features (and what feature weightings) are used to guide search to target images. The current experiments tested three likely candidate features (hue, shape, and orientation), examined whether these features are weighted in accord with whether those features are known to change or not, and how these features are used in combination. Not only could this data be directly applied to modeling efforts, providing a set of empirically-established parameters for the top-down component of guidance, but the results can inform the further development of search theory.

First, findings describing a mismatch effect, in which differences between a preview and the search target cause a decrease in performance, were replicated and extended, showing that these effects also occur in feature dimensions not previously explored (hue and shape). I also found that these mismatch effects in manual RT are due to both changes in search guidance and due to other processes (i.e. distractor rejection or target verification). Because eye movement measures were not previously used (Bravo & Farid, 2009; Vickery et al., 2005) it was possible that these mismatch effects were being driven purely by verification-type processes (such as distractor rejection) and that, as a result, previous conclusions regarding the search template were incorrect. In fact, Vickery et al. (2005, experiment 5B) did not find any reliable interaction between set size and orientation mismatch, which (in the absence of eye movement measures) may have been interpreted as mismatch not affecting guidance processes. The present study, however, confirmed that these mismatch effects can occur in guidance processes, demonstrating that this paradigm is useful in exploring the target template used in search. In addition, a strong similarity was found between patterns of mismatch effects in guidance measures and verification time, providing converging evidence for some recent suggestions of a common representation in guidance and verification (Maxfield & Zelinsky, 2012; see also Navalpakkam & Itti, 2005). Specifically, Maxfield and Zelinsky suggested that the representation used to guide search might be “locked in” and used in verification even when a different representation might be more useful for recognition than for guidance. In the present study, hue dominated the representation used for guidance, and hue mismatch effects were prominent in verification times as well, despite suggestions in the recognition literature that shape features typically dominate over color in object recognition (see Wichmann, Sharpe, & Gegenfurtner, 2002 for a discussion). In fact, many models of object recognition do not include the use of color at all, relying primarily on shape features (such as “geons” or “wire-frames”) instead. On the other hand, there are reasons to believe that color may dominate over shape features in at least some recognition tasks (e.g. Sanocki, Bowyer, Heath, and Sarkar, 1998). I speculate that—consistent with Maxfield and Zelinsky (2012)—the representation used to guide search (in this case, a target template based primarily on hue) carries over to the verification stage of the search task, and that verification and guidance may be far more similar than is typically believed. As a result, the features used in recognition may be highly task-dependent, and very different features may be found in recognizing objects in object-naming tasks than in tasks where other factors are involved (such as search). It does, however, remain possible that shape features other than aspect ratio are involved in this task, and caution should be taken when generalizing the present results to claims about shape in general. While the aspect ratio manipulation did involve a wide variety of shape features (e.g. height, width, and curvature) there are many shape features (e.g. area, the number of corners, and the number of line terminators) which were not assessed in the current work and which may yield much larger mismatch effects.

Another potential concern with the current findings is that these results may not generalize to previews of objects in non-canonical orientations or with non-canonical shapes or colors. In the current experiments, unaltered objects are likely to be of canonical shape and color, and may be in canonical orientations, simply because that is how objects are likely to be photographed. As detailed in section 3.3, the objects were selected from the Hemera Photo-Objects collection. The three volumes used contain a combined 150,000 images, allowing for a wide variety in the kinds of objects used, and many of those objects are very atypical. However, to the extent that typical, highly canonical objects were selected, any change in feature values is a change away from canonicality. It is therefore possible that some of the results may be due in part to changes away from canonicality rather than changes away from the previewed values (see Ballaz, Boutsen, Peyrin, Humphreys, & Marendaz, 2005, for explorations of this in the orientation dimension). It is possible that this concern could be addressed by manipulating targets not just in the search display, but at preview as well, but this creates many additional concerns. Most importantly, if a preview deviates from its canonical appearance, subjects might create a target template that more closely matches the object in its canonical appearance, which would result in a decrease in guidance even when the preview matched the search display exactly. One means of testing the effect of canonicality would be to assess the canonical position of each target in each feature space. For example, a hue wheel could be displayed and subjects could indicate what specific hue is canonical for a given object, and could rotate each object and indicate the canonical orientation, and could stretch each object and indicate its canonical aspect. These ratings could then be correlated with mismatch effects, and/or conditions could be made in which the mismatch moved either towards or away from canonicality. While it is possible to dissociate these factors, it is considered to be outside the scope of the current work. The purpose of these experiments was to test how features are used to guide search in the real world, in which searchers are likely to use a canonical representation of targets, as (by definition) that is the most common way that the objects are seen.

I also demonstrated that orientation and aspect mismatch effects only occur under conditions where hue is not available, and only begin to take effect after the first few eye movements, strongly suggesting that color dominates visual search from the very first few eye movements. Note that this pattern of effects is in-line with correlational studies demonstrating that fixations are more likely to land on areas with the target color than areas with the target shape (e.g. Williams, 1966) or target-similar orientations and spatial frequencies (e.g. Hwang et al., 2009; Rutishauser & Koch, 2007), providing converging evidence that color dominates search guidance. One possible explanation is that hue may be easier to discriminate in the periphery than aspect ratio and orientation, making hue a much more useful feature for guiding search. The first few eye movements would be driven by hue due to its discriminability, and only after the eye had moved closer to the target would orientation and aspect ratio become discriminable and begin to have an effect. However, even after aspect and orientation were sufficiently discriminable to have an effect (in time-to-target and later measures), these effects were not found in full-color displays: Hue dominates over other features even when those features were available. This dominance of hue in initial guidance has strong implications both for color-blind observers and for searching performed in the dark (for example, at night). In the context of the real-world or search through scenes (rather than discrete arrays of objects), this dominance of hue in early search processes might be even more striking, as global orientation and aspect ratio of objects cannot be accurately determined until those objects are segmented from the background, while hue often reliably signals a discontinuity between objects and can act as a cue in performing that segmentation. Even in the grayscale versions of the experiment, participants did not direct their initial eye movements based on aspect ratio or orientation information. While it is possible that this could change with large amounts of experience or training (such as the color-blind would have), it is clear that persons with color vision abnormalities are at a large disadvantage. The present data suggest that not only are those with abnormal color vision worse at tasks which obviously relate to red/green color abnormalities, such as searching for flowers or berries among other foliage (Cole & Lian, 2006) and at noticing street signs and traffic signals (O'Brien, Cole, Maddocks, & Forbes, 2002), but those with abnormal color vision are likely to have difficulty in search tasks in general, even those which can be performed using other features.

The lack of effects of orientation- or aspect-change when hue was available, together with guidance in some hue-change conditions continuing to decrease up to almost 180º of hue change (e.g. Figure 11), may suggest a surprising invariance (or lack of precision) in the target template, at least for these dimensions. This qualifies the findings of Vickery et al. (2005) and Bravo and Farid (2009), who concluded that the target template (in previewed search) uses detailed visual information and that orientation mismatch affects search performance. Instead, this appears to be the case only when hue is either not available (in grayscale experiments; e.g. Vickery et al., 2005) or when hue is made less useful by the experimental context, as in Bravo and Farid (2009) where the background consisted of brightly-colored coral reefs. When hue is available, the target template appears to be less precise in nature, and is invariant to changes in orientation and aspect.

Current models of visual search assume that attention and gaze are guided through the use of specific visual features from the target. Some image-based models that are designed to handle complex targets similarly extract features from a preview (e.g. Pomplun, 2006; Zelinsky, 2008) and suggest that many feature dimensions are used to guide search. However, even though participants were previewed with the target image, it is apparent from the current data that the precise visual information gained from the preview went largely unused. Target templates were either completely or partially orientation- and aspect ratio-invariant, and large changes to hue were needed to eliminate guidance based on hue information. This may reflect a memory limitation; if a finite number of feature values can be encoded into the target template, and hue information is far more useful than other feature dimensions, orientation and aspect information may not be included in the target template in order to maximize the encoding of hue information.

Another speculative possibility is that participants may rely primarily on a categorical representation of the target, perhaps elaborated with some visual features from the preview. When previewed with a red sports car, participants might search for "any red sports car" rather than the specific car shown at preview. In the context of the present experiments, the use of an elaborated categorical representation would allow search to be effectively guided to the target without relying on very specific hue, orientation, or aspect ratio information. The present data strongly suggest this possibility, as very little information from the preview appears to be used. This would still be consistent with previous research showing that previews are superior to text labels (e.g. Bravo & Farid, 2009; Castelhano, Pollatsek, & Cave, 2008; Malcolm & Henderson, 2009; Maxfield & Zelinsky, 2012; Vickery et al., 2005; Wolfe, Horowitz, Kenner, Hyle, & Vasan, 2004; Yang & Zelinsky, 2009), since a template of "a red sports car" would still be a better, more precise target template than can be created with the simple text labels (such as "car") which are typical in those experiments (see Schmidt & Zelinsky, 2009, for evidence that more precise text labels benefit guidance). The relationship between mechanisms involved in categorical search and previewed visual search is as yet unclear, with the only firm distinction made in the literature being that categorical search likely relies on long-term memory, while previewed visual search likely relies on visual working memory (see Yang & Zelinsky, 2009, for a discussion). If, in fact, previewed visual search relies on categorical knowledge to a smaller but still meaningful extent than categorical search, this would suggest that theories and models of visual search may be far easier to extend to categorical search tasks than one might expect. Some recent work does suggest that categorical search operates in a qualitatively similar way to previewed search (e.g. Alexander & Zelinsky, 2011), lending support to the possibility that categorical templates are involved in both forms of search tasks. The preview-target mismatch paradigm may be particularly useful in exploring the transition between the use of visual working memory and long-term memory as delays are introduced between the preview and the search display. As delays increase, the features used should be biased towards the features of the category to which the target belongs (retrieved from the participant's long-term memory). This could be reflected by an improvement to guidance, rather than a decrease, when the preview-target mismatch is changed in the direction of the categorical average. For example, if apples are typically red, and a green apple is used at preview, then a hue mismatch that changes the target to the color red would hurt performance when there is little to no delay between the preview and the search display (as the preview does not match the information available to guide search). With longer delays, however, this mismatch effect should decrease, disappear, and—with sufficient delay—might reverse and actually improve guidance as participants begin to use the visual features of the category rather than the preview. Another possibility is that when the representation shifts to long-term memory, performance may simply become more invariant to mismatch: Because categories are learned by viewing objects from different viewpoints and in different lightings, long-term representations of target categories are perhaps more likely to be invariant to aspect and orientation (which change with viewpoint) and hue (which changes with lighting). Future work varying the delay between preview and the search display may shed light on these possibilities.

The ability of this paradigm to test the precision of an individual's target template, and the features that a participant is using, provides some intriguing possibilities for future work. Exploring individual differences in precision and feature weightings could have strong implications for clinical populations such as persons with autism. Specifically, individuals with autism have been shown to have superior visual search performance (e.g. Joseph, Keehn, Connoly, Wolfe, & Horowitz, 2009), though the mechanisms behind this search advantage have only begun to be explored. One possibility is that individuals with autism may rely on a more precise target representation, allowing a stronger match to the target and limiting the distractor features that overlap with the target template. Neurotypical participants might rely more on conceptual or categorical information about the target, causing a loss in this precision. Further studies will be necessary to address such possibilities.

The pattern of the observed mismatch effects was also informative: Small amounts of preview-target mismatch resulted in large decreases in performance, but the performance decline decelerated (and/or plateaued) at larger amounts of mismatch (e.g. Figure 22A). Note that a similar pattern was found in Experiment 5B in Vickery et al. (2005), in which orientation mismatch was manipulated with grayscale, realistic objects: Orientation mismatch effects increased steeply up to around 60-90º of change, followed a "less orderly" effect in which the effects appear to pleateu (Vickery et al, 2005, p. 90). This decelerating pattern is surprising because there are a number of reasons to suspect that the search process would discount small amounts of deviance from the target template (which should result in *accelerating* effects of mismatch). First, real-world objects typically contain a range of values on a given feature dimension. This is particularly obvious when dealing with color: A preview "banana" image would contain a wide range of colors, not just a single yellow. Changes in depth or intrinsic shadows result in color differences across the image even if the local color of the object is identical across its entire surface. Encoding the likely range of values into the target template (rather than a single feature value) would be an effective strategy, allowing small changes in that dimension to still fall within the defined target range, which would result in no effect of mismatch at small amounts of change. Second, perceptual constancies (size, color, and shape constancy) by definition result from small changes in the stimuli being systematically discounted and may also have been expected to reduce mismatch effects at small amounts of change. Third, targets with small amounts of mismatch were likely to still be the most target-like objects in the search display. While a ten degree change in orientation between the preview and the search target would result in a feature mismatch, if the orientation of the target is a closer match to the preview than the orientation of any distractor objects, the target is still the best match on that dimension. In addition, when the orientation is changed (as in Experiment 1), the shape, color, and numerous other features still match the preview precisely. If, as most models of search predict, search is guided to areas with the most target-like features, guidance should have been relatively unaffected by small amounts of mismatch: Smaller feature changes that do not make the target less similar to the preview than distractor items should be largely irrelevant, as the largest guidance signal would still be at the target location. Fourth, intentionally discounting small differences from the preview might be an optimal strategy for this task due to the nature of the distractors: If a participant is trying to decide whether a currently fixated object matches the preview, simply responding to highly similar objects (without regard to small deviations from the expected feature values) would result in good performance since random object distractors are unlikely to be highly target-similar. Only larger feature changes should have required participants to compare targets to the preview more precisely. Despite all of these factors, the opposite pattern was found.

This pattern suggests that search may only be minimally (if at all) guided to the most target-like features when those features are below some threshold of similarity to the target template, and may then either be directed randomly or may rely only on other feature values. Note that many models of search (e.g. Zelinsky, 2008; Pomplun, 2006) would not predict this effect, instead assuming that all feature values in the display (however dissimilar) are included in the comparison process that guides eye movements. One possible explanation is that features may be coded categorically (e.g. Wolfe, Friedman-Hill, Stewart, & O'Connell, 1992). The most obvious prediction from a categorical coding of feature values, and one which is *not* supported by the current results, is that preview-target mismatch effects may only occur when the mismatch changes the feature value of the dimension past a categorical boundary. If, for instance, orientation is coded categorically, targets with a "steep" orientation should be invariant to small changes in orientation that do not cross the category boundary: A five degree change would still result in a "steep" orientation, matching the preview category. Instead, I speculate that category boundaries may play a role in the initial segmentation of a scene, and selection of objects that serve as potential distractors. An initial, early stage of search may identify which objects have features falling within the category boundary, followed by a comparison between those selected objects and the exact feature values in the target template. As a result, a feature dimension with too great a mismatch would not contribute any guidance signal at all, and only feature dimensions with little-to-no mismatch would generate a guidance signal. This exclusion of features that are outside of a certain range in values would be consistent with a suggestion by Neider and Zelinsky (2008) that only target-similar items are considered as potential distractors.

In experiment 4, I demonstrated that switching costs (in which performance is worse following a trial in which a different feature dimension was manipulated) only occur following hue-change trials. The dimension-weighting literature (e.g. Müller et al., 1995; Found & Müller, 1996; Müller et al., 2003), has shown that subjects bias processing towards features which *define* the target on previous trials. Switching cost analyses in Experiment 4 suggest that subjects did not, however, generally weight features differently due to what feature was changed on a previous trial. Some evidence of switching costs was found in Experiment 4, but only when switching *from* a hue-change trial to a trial in which a different feature dimension changed (an effect which may simply reflect changes in participant expectancies about the difficulty of the experiment). Assuming that the dimension-weighting theories of Müller and colleagues apply to complex, real-world objects, these switching costs should have been found when switching from a hue-change trial to a no change trial. The general lack of switching costs in the current study may suggest that weights are only switched on a trial-by-trial basis with simpler stimuli. In the experiments of Müller and colleagues, targets were always either singleton targets (defined by a feature which was not present on any distractors) or simple conjunction targets (defined by a unique combination of two features, among distractors that never shared that combination of features). When only one or two features are task-relevant, switching those features trial-by-trial may either be easier or more effective, as weight (a finite, limited resource) is not taken up by any additional feature dimensions beyond the two task-relevant features. In the present study, though participants did appear to rely primarily on hue, other features were also being used (as evident from the fact that guidance never reached chance in Experiment 2). Those other features—whatever they may be—may have taken up sufficient weight that either weights were not reassigned, or were changed to too small a degree to be detected. This, together with the relative invariance to aspect- and orientation-changes, may suggest a relatively low limit to top-down biasing signals involved in guidance.

I also demonstrated that mismatch effects are more pronounced when there is uncertainty in which feature dimension will change, but only in the context of the full-color experiments (Figure 22A). Participants appear to have weighted hue heavily in Experiment 4, since hue was valid on most trials, and deweighted hue information in Experiment 2, where hue was invalid on almost all trials. Similarly, aspect ratio and orientation appear to have been deweighted in Experiments 1 and 3, in which those dimensions were generally invalid, while the weights were more evenly distributed in Experiment 4 (at least for the full-color trials). This suggests that, although reweighting of dimensions did not occur on a trial-by-trial basis, participants can use different weights according to context: Feature dimensions that prove to be consistently invalid can be deweighted or ignored, improving performance.

Finally, in Experiment 5, I tested the assumption of linear summation that underlies most models of visual search, which had not previously been directly tested. I conclude that the assumption of linearity appears to be valid in grayscale experiments: Orientation and aspect ratio mismatch both summed linearly, in guidance and in most other measures. This lack of nonlinearities, however, may be due to the dominance of color in visual search: Minimally-guiding features may sum linearly, or may have too small an effect for the nonlinearities to emerge. When hue is involved, something more than a simple linear sum is involved. Not only did combined conditions that included a hue manipulation (hue*aspect, orientation*hue, and orientation*hue*aspect) have superadditive mismatch effects, larger than would be expected by linear summation, but the orientation*aspect condition also showed some nonlinearities in the full-color version of the task. It is apparent that a linear summation is not sufficient to explain how feature dimensions are combined, contrary to the assumptions in many models of search (e.g. Cave, 1999; Eckstein, Thomas, Palmer, & Shimozaki, 2000; Navalpakkam & Itti, 2005; Palmer et al., 2000; Wolfe, 1994; Wolfe, Cave, & Franzel, 1989). Instead, it appears that targets which mismatch on multiple dimensions are treated as being more different than one would expect from linear summation, and models of search should be modified to account for these findings. One possible explanation is that the maximum use (weighting) of a feature may only occur in the presence of mismatch in other features. When none of the colors in the display match the target template, and hue is therefore unreliable, more weighting may be given to shape and orientation, resulting in superadditive mismatch effects when those dimensions are also unreliable.

Mismatch effects were replicated and extended to two new feature dimensions: Hue and shape), providing a set of empirically-established parameters for the top-down component of guidance, and allowing a number of conclusions to be drawn about the representation used to guide search. First, orientation and aspect mismatch effects only occur under conditions where hue is not available (or is not useful), strongly suggesting that color dominates visual search from the very first few eye movements. Second, and consistent with Maxfield and Zelinsky (2012), the representation used to guide search appears to carry over to the verification stage of the search task, resulting in mismatch effects in both search guidance and verification processes, and suggesting that the same underlying representation is used at both stages of the search task. Third, the extensive invariance in the target template to orientation and aspect mismatch effects demonstrates that very little information from the preview is actually used, and participants may rely on largely categorical representations of the target even when previews are available. Fourth, the mismatch cost decelerated (and/or plateaued) at larger amounts of mismatch, suggesting that search may only be minimally guided to the most target-like features when those features are below some threshold of similarity to the target template (possibly determined by category boundaries). Fifth, the present data generally did not support the presence of switching costs (e.g. Müller et al., 1995; Found & Müller, 1996; Müller et al., 2003), which may suggest that dimension-weighting only occurs on a trial-by-trial basis with simpler stimuli. Sixth, mismatch effects are more pronounced when there is uncertainty in which feature dimension will change, demonstrating that participants can use different weights according to context, deweighting or ignoring feature dimensions that are consistently invalid. Finally, I conclude that the assumption of linearity (made by many models of visual search) appears to only be valid in grayscale experiments. When hue is involved, feature dimensions are combined superadditively, with targets mismatching on multiple dimensions being treated as more different than one would expect from linear summation.

Literature Cited

Alexander, R. G., & Zelinsky, G. J. (2011). Visual similarity effects in categorical search. *Journal of Vision*, *11*, 1-15.

Alexander, R. G., & Zelinsky, G. J. (2012). Effects of part-based similarity on visual search: The Frankenbear experiment. *Vision Research, 54*, 20-30.

Alvarez, G. A., & Cavanagh, P. (2004) The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science*, *15*, 106-111

Ashby, F. G., Prinzmetal, W., Ivry, R., & Maddox, W. T. (1996). A formal theory of feature binding in object perception. *Psychological Review, 103*, 165-192.

Ballaz, C., Boutsen, L., Peyrin, C., Humphreys, G., & Marendaz, C. (2005). Visual search for object orientation could be modulated by canonical orientation. *Journal of Experimental Psychology: Human Perception and Performance, 31*, 20-39.

Bauer, B., Jolicoeur, P. & Cowan, W. B. (1996). Visual search for colour targets that are or are not linearly-separable from distractors. *Vision Research, 36,* 1439–1466.

Bravo, M.J. & Farid, H. (2009). The Specificity of the Search Template. *Journal of Vision, 9*, 1-9.

Brockmole, J. R., & Henderson, J. M. (2006). Recognition and attention guidance during contextual cueing in real-world scenes: Evidence from eye movements. *Quarterly Journal of Experimental Psychology, 59,* 1177-1187.

Carter, R. C. (1982). Visual search with color. *Journal of Experimental Psychology: Human Perception & Performance*, *8*, 127-136.

Castelhano, M.S., Pollatsek, A., & Cave, K.R. (2008). Typicality Aids Search for an Unspecified Target, but Only in Identification, and not in Attentional Guidance. *Psychonomic Bulletin and Review, 15*, 795-801.

Cave, K. R. (1999). The FeatureGate model of visual selection, *Psychological Research, 62,* 182-194.

Chen, X., & Zelinsky, G. J. (2006). Real-world visual search is dominated by top-down guidance, *Vision Research, 46*, 4118-4133.

Cole, B. L., & Lian, K.-Y. (2006). Search for coloured objects in natural surroundings by people with abnormal colour vision, *Clinical and Experimental Optometry, 89*, 144–149.

Daoutis C.A., Pilling M., Davies, I.R.L. (2006). Categorical effects in visual search for colour. *Visual Cognition. 14,* 217–240.

Deubel, H., & Schneider, W. X. (1996). Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Visual Research, 36*, 1827-1837.

Donk, M. (1999). Illusory conjunctions are an illusion: The effects of target-nontarget similarity on conjunction and feature errors. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 1207-1233.

Donk, M. (2001). Illusory conjunctions die hard: A reply to Prinzmetal, Diedrichsen, and Ivry (2001). *Journal of Experimental Psychology: Human Perception and Performance, 27*, 542-546.

Duncan J. (1989). Boundary conditions on parallel processing in human vision. *Perception, 18,* 457 – 469.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96,* 433-458.

D'Zmura, M. (1991). Color in visual search. *Vision Research, 31*, 951-966.

Eckstein, M. P. (1998). The lower visual search efficiency for conjunctions is due to noise and not serial attentional processing (an SDT model), *Psychological Science, 9*, 111-118.

Eckstein, M. P., Drescher, B., & Shimozaki, S. S. (2006). Attentional cues in real scenes, saccadic targeting and Bayesian priors. Psychological Science, 17, 973-80.

Eckstein, M. P., Thomas, J. P., Palmer, J., & Shimozaki, S. S. (2000). A signal detection model predicts the effects of set size on visual search accuracy for feature, conjunction, triple conjunction, and disjunction displays, *Perception & Psychophysics, 62*, 425-451.

Ehinger, K. A., Hidalgo-Sotelo, B., Torralba, A., & Oliva, A. (2009). Modeling search for people in 900 scenes: A combined source model of eye guidance, *Visual Cognition, 17,* 945-978.

Findlay, J. M., & Gilchrist, I. D. (2003). Active Vision: The Psychology of Looking and Seeing. Oxford, UK: Oxford University Press.

Foster, D. H., Savage, C. J., Mannan, S., & Ruddock, K. H. (2000). Asymmetries of saccadic eye movements in oriented-line-target search. *Vision Research, 40,* 65-70.

Foster, D. H., & Ward, P. A. (1991). Asymmetries in oriented-line detection indicate two orthogonal filters in early vision. *Proceedings of the Royal Society of London Series B, 243*, 75-81.

Foulsham, T., & Underwood, G. (2011). If saliency affects search then why? Evidence from normal and gaze-contingent search tasks in natural scenes. *Cognitive Computation, 3,* 48-63.

Found, A., & Müller, H. J. (1996). Searching for unknown feature targets on more than one dimension: Investigating a "dimension-weighting" account. *Perception & Psychophysics, 58,* 88-101.

Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review, 96*, 267-314.

Green, D. M., & Swets, J. A. (1966). Signal Detection Theory and Psychophysics. New York: Wiley.

Guest, D., & Lamberts, K. (2011). The time course of similarity effects in visual search. *Journal of Experimental Psychology: Human Perception and Performance, 37,* 1667-1688.

Hayakawa, S., Kawai, N., & Masataka, N. (2011). The influence of color on snake detection in visual search in human children. *Scientific Reports, 1*, 1–4.

Huang, L. & Pashler, H (2005). Quantifying object salience by equating distractor effects. *Vision Research, 45*, 1909-1920.

Hwang, A. D., Higgins, E. C., & Pomplun, M. (2009). A model of top-down attentional control during visual search in complex scenes. *Journal of Vision, 9*, 1-18.

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

Johnston, J. C., & Pashler, H. (1990). Close binding of identity and location in visual feature perception. *Journal of Experimental Psychology: Human Perception and Performance, 16*, 843-856.

Jolicoeur, P. (1985). The time to name disoriented natural objects, *Memory and Cognition 13*, 289-303.

Joseph, R., Keehn, B., Connolly, C., Wolfe, J.M., & Horowitz, T.S. (2009). Why is visual search superior in autism spectrum disorder? *Developmental Science, 12(*6), 1083-96.

Klein, R. M., & Farrell, M. (1989). Search performance without eye movements. *Perception & Psychophysics, 9*, 476-482.

Kotowicz, A., Rutishauser, U, & Koch, C. (2010). Time course of target recognition in visual search, *Frontiers in Human Neuroscience, 4,* 1-11.

Kristjánsson, Á. & Tse, P. U. (2001). Curvature discontinuities are cues for rapid shape analysis, *Perception & Psychophysics, 63*, 390-403.

Levin, D. T., Takarae, Y., Miner, A. G., & Keil, F. (2001). Efficient visual search by category: Specifying the features that mark the difference between artifacts and animals in preattentive vision, *Perception & Psychophysics, 63*, 676-697.

Luck, S. J. & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature, 390*, 279-281.

Malcolm G.L. & Henderson J.M. (2009) The effects of target template specificity on visual search in real-world scenes: Evidence from eye movements. *Journal of Vision, 9, 1-13.*

Maxfield, J. T., & Zelinsky, G. J. (2012) Searching through the hierarchy: How level of target categorization affects visual search. *Visual Cognition, 20(10)*, 1153-1163.

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

Müller, H. J., Reimann, B., & Krummenacher, J. (2003). Visual search for singleton feature targets across dimensions: Stimulus- and expectancy-driven effects in dimensional weighting. *Journal of Experimental Psychology: Human Perception & Performance, 29*, 1021–1035.

Murray, J. E. (1997). Flipping and spinning: Spatial transformation procedures in the identification of rotated natural objects, *Memory and Cognition 25*, 96-105.

Murthy A., Thompson K.G., Schall J.D. (2001). Dynamic dissociation of visual selection from saccade programming in frontal eye field. *Journal of Neurophysiology* 86, 2634-2637.

Nagy, A. L., Sanchez, R. R., & Hughes, T. C. (1990) Visual search for color differences with foveal and peripheral vision, *Journal of the Optical Society of America A, 7*, 1995-2001.

Najemnik, J. & Geisler, W.S. (2005) Optimal eye movement strategies in visual search. *Nature, 434*, 387-391.

Navalpakkam, V., & Itti, L. (2005). Modeling the influence of task on attention. *Vision Research, 45*, 205-231.

Navon, D., & Ehrlich, B. (1995). Illusory conjunctions: Does inattention really matter? *Cognitive Psychology, 29,* 59-83.

Neider, M. B., & Zelinsky, G. J. (2006). Scene context guides eye movements during search. *Vision Research*, *46*, 614-621.

Neider, M. B., & Zelinsky, G. J. (2008). Exploring set size effects in scenes: Identifying the objects of search. *Visual Cognition*, *16*, 1-10.

Nothdurft, H. C. (1991). Texture segmentation and pop-out from orientation contrast. *Vision Research, 31,* 1073-1078.

Nothdurft, H. C. (2000). Salience from feature contrast: Additivity across dimensions. *Vision Research, 40,* 1183-1201.

O'Brien, K.A., Cole, B.L., Maddocks, J.D., Forbes, A.B. (2002). Color and defective color vision as factors in the conspicuity of signs and signals, *Human Factors, 44*, 667–677.

Oliva, A., Torralba, A., Castelhano, M.S., & Henderson, J. M. (2003). Top-down control of visual attention in object detection. In *Proceedings of the IEEE international conference on image processing* (Vol. I, pp. 253–256). Barcelona, Spain: IEEE.

Oliva, A., Wolfe, J.M., & Arsenio, H.C. (2004). Panoramic Search: The Interaction of Memory and Vision in Search Through a Familiar Scene. *Journal of Experimental Psychology: Human Perception and Performance, 30*, 1132-1146.

Palmer, J. (1995). Attention in visual search: Distinguishing four causes of a set-size effect, *Current Directions in Psychological Science, 4*, 118-123.

Palmer, J., Verghese, P, & Pavel, M. (2000) The psychophysics of visual search, *Vision Research, 40,* 1227-1268.

Parkhurst, D. J., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision Research,42*, 107–123.

Peterson, M. S., Kramer, A. F., & Irwin, D. E. (2004). Covert shifts of attention precede involuntary eye movements. *Perception & Psychophysics, 66*, 398-405.

Poirier, F.J.A.M., & Gurnsey, R. (1998). The effect of eccentricity and spatial frequency on the orientation discrimination asymmetry. *Spatial Vision, 11*, 349-366.

Pomplun, M. (2006). Saccadic selectivity in complex visual search displays.*Vision Research, 46,* 1886–1900.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32,* 3–25.

Sanocki, T., Bowyer, K., Heath, M. & Sarkar, S. (1998). Are edges sufficient for object recognition? *Journal of Experimental Psychology: Human Perception and Performance, 24*, 340-349.

Rao, R., Zelinsky, G., Hayhoe, M., & Ballard, D. (1996). Modeling saccadic targeting in visual search. In D. Touretzky, M. Mozer, & M. Hasselmo (Eds.), *Advances in neural information processing systems*(Vol. 8, pp. 830–836). Cambridge, MA: MIT Press.

Rao, R., Zelinsky, G., Hayhoe, M., & Ballard, D. (2002). Eye movements in iconic visual search. *Vision Research, 42,* 1447–1463.

Reijnen E., Wallach D., Stöcklin M., Kassuba T., & Opwis K. (2007). Colour similarity in visual search. *Swiss Journal of Psychology, 66,* 191-199.

Rensink, R. A., and Enns, J. T. (1998) Early completion of occluded objects. *Vision Research, 38*, 2489-2505.

Rosenholtz, R. (2001). Search asymmetries? What search asymmetries? *Perception & Psychophysics, 63,* 476-489.

Rutishauser, U., & Koch, C. (2007). Probabilistic modeling of eye movement data during conjunction search via feature-based attention. *Journal of Vision, 7*, 1-20.

Schmidt, J., & Zelinsky, G. J. (2009). Search guidance is proportional to the categorical specificity of a target cue. *Quarterly Journal of Experimental Psychology, 62,* 1904-1914.

Simmons, D. R., & Foster, D. H. (1992). Segmenting textures of curved-line elements. In G. A. Orban, Nagel, H H (Ed.), Artificial and Biological Vision Systems, (pp. MS): Springer-Verlag.

Smallman, H. S., & Boynton, R. M. (1990). Segregation of basic colors in an information display, *Journal of the Optical Society of America A*, *7,* 1985-1994.

Torralba, A., Oliva, A., Castelhano, M. and Henderson, J. M. (2006). Contextual guidance of attention in natural scenes: The role of global features on object search. *Psychological Review, 113*, 766-786.

Townsend, J. T. (1990). Serial vs. parallel processing: Sometimes they look like Tweedledum and Tweedledee but they can (and should) be distinguished, *Psychological Science, 1*, 46-54.

Treisman, A. (1988). Features and objects: The Fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology, 40*(A), 201–237.

Treisman, A. (1991). Search, similarity, and integration of features between and within dimensions. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 652–676.

Treisman, A. M., & Gelade, G. (1980). A feature integration theory ofattention. *Cognitive Psychology*, *12,* 97–136.

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

Treisman, A., & Sato, S. (1990). Conjunction search revisited, *Journal of Experimental Psychology: Human Perception and Performance, 16*, 459-478.

Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General, 114*, 285-310.

Vickery T.J., King L.-W., & Jiang Y. (2005). Setting up the target template in visual search. *Journal of Vision, 5*, 81-92.

Vighneshvel, T. & Arun, S.P. (2013). Does linear separability really matter? Complex search is explained by simple search. *Journal of Vision, 13*, 1-24.

Weidner, R., & Müller, H. J. (2009). Dimensional weighting of primary and secondary target-defining dimensions in visual search for singleton conjunction targets. *Psychological Research, 73,* 198–211.

Wichmann, F.A., Sharpe, L.T. & Gegenfurtner, K.R. (2002) The contributions of color to recognition memory for natural scenes. *Journal of Experimental Psychology: Learning, Memory & Cognition, 28,* 509-520.

Williams, L. G. (1966). The effect of target specification on objects fixated during visual search. *Perception & Psychophysics, 1*, 315-318.

Williams, C. C. (2010). Not all visual memories are created equal. *Visual Cognition, 18*, 201-228.

Williams, C. C., Henderson, J. M., & Zacks, R. T. (2005). Incidental visual memory for targets and distractors in visual search. *Perception & Psychophysics, 67*, 816-827.

Wolfe, J.M. (1992) "Effortless" Texture Segmentation and "Parallel" Visual Search are not the Same Thing.*Vision Research , 32*, 757-763

Wolfe, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review, 1*, 202-238.

Wolfe, J. M. (1998). What can 1 million trials tell us about visual search? *Psychological Science, 9*, 33-39.

Wolfe, J.M. & Bennett, S.C. (1997) Preattentive Object Files: Shapeless bundles of basic features. *Vision Research*, 37, 25-44.

Wolfe, J.M., Cave, K. R., & Franzel, S.L. (1989). Guided Search: An Alternative to the Feature Integration Model for Visual Search. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 419-433.

Wolfe, J.M., Friedman-Hill, S.R., Stewart, M I., & O'Connell, K. M. (1992) The Role of Categorization in Visual Search for Orientation. *Journal of Experimental Psychology, 18,* 34-49.

Wolfe, J.M., & Horowitz, T.S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, *5,* 1-7.

Wolfe, J.M., Horowitz, T.S., Kenner, N., Hyle M., & Vasan, N. (2004). How fast can you change your mind? The speed of top-down guidance in visual search. *Vision Research, 44,* 1411-1426.

Wolfe, J.M., Klempen, N., & Shulman, E. (1999) Which End is Up? Two Representations of Orientation in Visual Search. *Vision Research*, *39,* 2075-2086

Wolfe, J. M., Reijnen, E., Horowitz, T. S., Pedersini, R., Pinto, Y., & Hulleman J. (2011). How does our search engine see the world?, *Attention, Perception, & Psychophysics*, *73,* 1054-1064.

Wolfe, J. M., Yee, A., & Friedman-Hill, S. R. (1992). Curvature is a basic feature for visual search. *Perception*, *21*, 465-480.

Wolfe, J. M., Yu, K. P., Stewart, M. I., Shorter, A. D., Friedman-Hill, S. R., & Cave, K. R. (1990). Limitations on the parallel guidance of visual search: color x color and orientation x orientation conjunctions. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 879-892.

Yang, H., & Zelinsky, G. J. (2009). Visual search is guided to categorically-defined targets. *Vision Research*, *49*, 2095-2103.

Yang, H., Chen, X., & Zelinsky, G. J. (2009). A new look at novelty effects: Guiding search away from old distractors. *Attention*, *Perception*, & *Psychophysics*, *71*, 554-564.

Yokoi, K., & Uchikawa, K. (2005). Color category influences heterogeneous visual search for color. *Journal of the Optical Society of America A, 22*, 2309-2317.

Zelinsky, G. J. (2008). A theory of eye movements during target acquisition. *Psychological Review*, *115*, 787-835.

Zelinsky, G. J., & Loschky, L. C. (2005). Eye movements serialize memory for objects in scenes. *Perception & Psychophysics*, *67*, 676-690.

Zelinsky, G. J., Rao, R. P. N., Hayhoe, M. M., & Ballard, D. H. (1997). Eye movements reveal the spatiotemporal dynamics of visual search. *Psychological Science, 8(6)*, 448-453.

Zelinsky, G. J., & Sheinberg, D. L. (1997). Eye movements during parallel-serial visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 244-262.

Zelinsky, G. J., & Schmidt, J. (2009). An effect of referential scene constraint on search implies scene segmentation. *Visual Cognition*, *17*, 1004-1028.