

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Information Attracts Attention: A Probabilistic Account of the Cross-Race Advantage in Visual Search

Permalink

<https://escholarship.org/uc/item/52v909x5>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 29(29)

ISSN

1069-7977

Authors

Lingyun, Zhang
Tong, Matthew H.
Cottrell, Garrison W.

Publication Date

2007

Peer reviewed

Information Attracts Attention: A Probabilistic Account of the Cross-Race Advantage in Visual Search

Zhang Lingyun Matthew H. Tong Garrison W. Cottrell
(lingyun, mhtong, gary@ucsd.edu)

Department of Computer Science and Engineering
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92037 USA

Abstract

People are better at recognizing faces from their own race than from different races (Shapiro & Penrod, 1986; Bothwell, Brigham, & Malpass, 1989), an effect commonly known as the other-race effect. The causes of this effect have been attributed to the fact that people have more experience with faces from their own race during development (Feingold, 1914; Chance, Turner, & Goldstein, 1982; Shepard, 1981; Valentine, Chiroro, & Dixon, 1995). However, in visual search tasks, cross-race (CR) faces are found faster than same-race (SR) faces (Levin, 1996, 2000). This advantage of CR faces in visual search tasks seems at first to be inconsistent with the advantage of the SR face in recognition tasks. To account for this discrepancy, Levin proposed that there is a race feature, which is active only for CR faces. By explicitly assuming this feature, the face search data fits into a visual search paradigm in which the search asymmetry can be explained. In this paper, we will present an alternate explanation of the CR face advantage in visual search which is consistent with the SR face advantage in recognition without making additional assumptions. The probabilistic visual search model we developed predicts that CR faces should be found faster than SR faces, given that people have more experience with SR faces. This model was developed based on an intuitive assumption of the goal of the visual system, yet, with no extra assumptions, it fully accounts for the CR advantage based on the same mechanism believed to be responsible for the SR advantage in recognition.

Keywords: visual search; other race effect; search asymmetry; saliency; self-information

Introduction

People are better at recognizing faces from their own race than from different races (Shapiro & Penrod, 1986; Bothwell et al., 1989). This is known as the other-race effect, cross race effect or own race bias. Most researchers agree that this phenomenon is caused by the learning history; the fact that people have more experience with their own race faces during development results in more accurate recognition. As minorities usually do not have an other race effect for the majority, it might be more appropriate called an unfamiliar facial structure effect. It was reported that while six year old children did not show a significant other-race effect, older participants show a degree of the effect positively correlated with age (Chance et al., 1982). Feingold suggested that the degree of the other-race effect a person shows is dependent on how much contact he has with people from other races (Feingold, 1914). Later studies argued that contact for individuation is more important than mere contact (Shepard, 1981; Valentine et al., 1995). This leads to the hypothesis of optimal feature selection. The key idea is that an optimal feature set is se-

lected for identifying individuals that a person comes in contact with. As people are generally more exposed to faces from their own race, the features developed are tuned accordingly and are thus less optimal for faces from other races. This hypothesis has been computationally modelled using principal component analysis, which successfully accounts for the other-race effect. (O'Toole, Deffenbacher, Abdi, & Bartlett, 1991; O'Toole, Deffenbacher, Valentin, & Abdi, 1994).

On the other hand, Levin found that people who show the other-race effect are faster in searching for a cross-race (CR) face among same-race (SR) faces than searching for an SR faces among CR faces (Levin, 1996, 2000). Figure 1 shows an example stimulus from his experiment. In this stimulus, there is one African face (the bottom one) and seven images of a Caucasian face. The faces were preprocessed to have identical shading, hair, ears and jaw lines. The only difference between faces from different races are thus the internal features (e.g. eyes, mouth, and nose). The subjects' task is to detect whether the target face, i.e. the African face in this stimulus, is present or not. For a Caucasian subject, this requires looking for a CR face in SR faces.



Figure 1: A sample stimulus in the visual search task. The target in this stimulus is an African face (the bottom one). The rest, i.e. the distractors, are Caucasian faces.

Levin's finding, referred to as the advantage of CR faces in visual search, seems to be in conflict with the decreased accuracy of CR faces in recognition; it is counterintuitive that ex-

tensive experience with SR faces will lead to a faster search of CR faces. The optimal feature selection hypothesis does not seem to provide a straightforward explanation either, since intuitively a good feature set should aid in search. Yet, since subjects who do not show the other race effect also do not show the search asymmetry, these two phenomena must be tightly related (Levin, 2000).

To accommodate this new observation, Levin proposed the addition of a race feature which captures whether a face was in or out of one's racial group. This race feature is positive for CR faces. Treisman showed that visual search for a feature-positive target among feature-negative distractors is faster than the converse (Treisman & Gormican, 1988). With the assumption that CR faces are feature-positive for race, the visual search for CR face among SR should be faster.

In computational models, the feature-positive advantage is often achieved using a front-end parallel feature processing system, as in (Wolfe, 1994, 2001b; Itti, Koch, & Niebur, 1998; Itti & Koch, 2000). As the feature-positive item elicits a stronger response from the feature detector, it attracts more attention than the feature-negative item. Thus, not every single feature has the quality of causing search asymmetry. Only certain basic features that are processed in parallel in the low level visual system are eligible (see section II of (Wolfe, 1998) for a review of basic features). The so called race feature, seems to be too abstract and high level to be processed in parallel by the low-level visual system. It is therefore not as straightforward as it seems to generalize the search asymmetry Treisman and Gormican found in their experiments to the cross-race advantage in facial visual search.

Furthermore, the CR recognition deficit vanishes or is greatly reduced with certain kinds of repeated exposure (Levin, 2000); for example, white basketball fans who must frequently differentiate between black basketball players show no deficit (Li, Dunning, & Malpass, 1998). What happens to the race feature as experience with CR faces alters the CR deficit? Without incorporating the effects of experience, it seems that the race feature explanation is incomplete.

In this paper, we present a probabilistic visual search model based on a simple assumption of the visual system's goal. We show that the CR advantage in visual search is a natural prediction of our model, given that people experience more SR faces during the development. It is not due to an advantage or disadvantage in processing of CR face, but a natural result of the fact that CR faces occur with low probability during development. This model, coupled with the compatible hypothesis of optimal feature selection, is thus capable of explaining both the CR recognition deficit and the the CR search advantage. Features are selected to be optimal for recognition (giving CR faces a disadvantage) and the fact that CR faces are outliers in the feature space leads to them being faster to find (see (Haque & Cottrell, 2005) for an alternate formulation of this claim).

Probabilistic Search

In this section we introduce our probabilistic visual search model. First we will define bottom-up saliency, i.e. how strongly a visual feature attracts attention. We will show that the visual system should direct attention more to visual features that occur with low probability. Then we will discuss how attention is directed in classical visual search paradigms. Our model assumes that attention is directed probabilistically according to the saliency distribution in the visual field. With some simplifying assumptions, the expected search time can be derived in closed form. The search time is positively correlated with the relative saliency of the distractors to the target.

Saliency is Information

Our definition of saliency is derived from a single assumption: a goal of the visual system is to find potential targets that are important for the organism's survival, e.g. food and predators. Under this assumption the visual system should direct its attention to locations where visual features suggest a high probability of such targets. To achieve this, the visual system must actively estimate the probability of an interesting object being at a particular location given the scene's visual features. We denote the occurrence of an interesting object at a point z in the visual field with the binary variable C_z , the perceptual representation of a point's visual features employed by the visual system with F_z , and the location of a point with L_z . The proposed theory makes no claims as to the nature of F_z or to how high- or low-level these representations may be. The probability of interest is $p(C_z = 1|F_z = f, L = l)$. In this paper, we will concentrate on how visual features affect attention and thus assume location invariance, dropping L from the equation and resulting in $p(C_z = 1|F_z = f)$ being the probability to be estimated. This probability can then be calculated using Bayes' rule:

$$p(C_z = 1|F = f) = \frac{p(F = f|C_z = 1) \cdot p(C_z = 1)}{p(F_z = f)} \quad (1)$$

The visual system compares relative probabilities over visual features to decide where to direct attention. The prior $p(C_z = 1)$ can therefore be dropped from the formula because it is independent of features. This leaves us with:

$$p(C_z = 1|F_z = f) \propto \frac{p(F_z = f|C_z = 1)}{p(F_z = f)}. \quad (2)$$

Note that estimating $\log p(C_z = 1|F_z = f)$ is as good as estimating the probability itself because the logarithm is monotonically increasing. Therefore the ordering will be kept.

$$\log p(C_z = 1|F_z = f) = \log \frac{p(F_z = f|C_z = 1)}{p(F_z = f)} + const. \quad (3)$$

$$= -\log p(F_z = f) + \log p(F_z = f|C_z = 1) + const. \quad (4)$$

The first term $-\log p(F_z = f)$ is dependent only on the visual feature of point z but not on what the target class is.

In information theory, $-\log p(F_z = f)$ is known as the *self-information* of random variable F_z when it takes the value f . Self-information increases when the probability of a feature decreases. The second term $\log p(F_z = f | C_z = 1)$ is dependent on the target class and is known as *log-likelihood*. It biases features that are consistent with the target with higher saliency, facilitating the search. The sum of these two terms is also known as *pointwise mutual information*.

When the organism is not actively searching for any particular target (free viewing), the target class is unknown. But to best survive, if any potential target such as prey or predators appear in the visual field, attention should be directed to it. In this case, the animal should not assume the potential target is more likely to bear one feature value than another (the target is as likely to be green as blue). This is equivalent to assuming a uniform distribution on the likelihood term. Hence only the first term of self-information is relevant.

We propose that the first term $-\log p(F_z = f)$, or the self-information contained by the visual features, is the bottom up saliency of z in the visual field. The visual system actively calculates it over the visual field regardless of whether there is a particular target or not. Features with small probability have high saliency and are more attractive to attention.

A Probabilistic Scan Procedure in Visual Search

We make the assumption that attention is directed *stochastically*. The probability of attending to a particular location, feature, or object is equal to its saliency, modified by an activation function g , and normalized, i.e. $p_z = \frac{g(s_z)}{\sum_z g(s_z)}$, where p_z is the probability of attending to point z and s_z is the salience of point z . When g is an exponential function, this reduces to a soft max function. The probability map is virtual in that it does not need to be normalized explicitly, which would require global communication. Neither the mechanism through which stochastic behavior arises nor the specific activation function are crucial to our conclusions, so they remain unspecified. We follow the intuition that salient features are more likely to be attended, so the activation function should be monotonically non-decreasing. Less salient features may also be attended but with lower probability. The relative advantage of high salience depends on the steepness of the activation function. For example, an exponentially increasing activation function will tend to suppress locations except for the most salient ones.

We can apply this to the visual search paradigm and derive search time quantitatively. We use five assumptions to simplify our analysis (several adapted from (Wolfe, 1994)). They provide a closed form solution for our qualitative analysis, but can be relaxed without changing our qualitative conclusions.

- Subjects scan items in discrete time steps. In each step, exactly one item is attended. An item is abstracted to be a point z as used in the earlier saliency inference.
- Inhibition of return is strict, i.e. attended items shall not be attended again.

- Saliency does not change over time. I.e. the saliency map is not updated by eye movements and shifts of attention.
- The saliency of objects outside the subject's display is zero, i.e. external objects do not compete for attention.
- The saliency of all distractors is the same.
- The subjects scan until they find the target.

Let s_{targ} denote the saliency of the target and s_{dist} denote that of a distractor. When there are n distractors, the target is attended on the first time step with probability

$$\frac{g(s_{targ})}{g(s_{targ}) + n \cdot g(s_{dist})}. \quad (5)$$

A distractor is attended with probability

$$\frac{n \cdot g(s_{dist})}{g(s_{targ}) + n \cdot g(s_{dist})}. \quad (6)$$

Now we can infer our model's behavior. We denote the *distractor strength* as

$$x = \frac{g(s_{dist})}{g(s_{targ})} \quad (7)$$

This can be thought of as the relative salience of the distractors vs. the target. For a classical feature target search, the target is highly salient compared to the distractors, so the distractor strength x is very small. For classical conjunction target search, the distractors are as salient as the target and the distractor strength x is approximately 1. We can divide out the salience of the target and rewrite the probability in equation 5 in terms of the distractor strength x . Letting D_t denote the detection of the target on scan step t , we have:

$$P(D_1) = \frac{g(s_{targ})}{g(s_{targ}) + n \cdot g(s_{dist})} \quad (8)$$

$$= \frac{1}{1 + n \cdot g(s_{dist})/g(s_{targ})} \quad (9)$$

$$= \frac{1}{1 + n \cdot x} \quad (10)$$

Intuitively, larger x implies that the distractors are more salient and the target is less likely to be found in one fixation. Similarly, a greater number of distractors reduces the probability that the target will be found on the first scan.

We are now ready to derive the expected number of steps to find the target. Let $E_{n,x}$ denote the expected number of steps when n distractors are present with distractor strength x . The target will be found on the first step with some probability, $P(D_1)$. If the target is not found on the first step, the assumption of strict inhibition of return dictates that we repeat the procedure but with $n - 1$ distractors. Thus the expected number of steps for n distractors can be found by considering these two cases.

$$E_{n,x} = P(D_1) \cdot 1 + (1 - P(D_1)) \cdot (1 + E_{n-1,x}) \quad (11)$$

$$= 1 + \frac{n \cdot x}{1 + n \cdot x} \cdot E_{n-1,x} \quad (12)$$

The target will always be found on the first scan step when there are no distractors, i.e.:

$$E_{0,x} = 1 \quad (13)$$

Equations 12 and 13 can also be equivalently described as:

$$E_{n,x} = 1 + \frac{x}{1+x} \cdot n \quad (14)$$

This formula says that the expected number of scan steps increases linearly with the number of items, with a slope of $\frac{x}{1+x}$. When the distractor strength is small, the slope approaches zero and the target “pops out” and is usually found in the first step regardless of the number of distractors. As the distractor strength x increases, the slope increases and a greater response time is required.

Search Asymmetry

Search asymmetry refers to the phenomenon that the response time slope typically changes when the role of the target and distractor are switched (see (Wolfe, 2001a) for a review). For example, a bar tilted 15 degrees will “pop out” among vertical bars, but a sole vertical bar among tilted bars seems to produce a serial search, where the search time goes up linearly with the number of distractors. In our model, the search slope is determined by the distractor strength x , which is in turn a function of s_{dist} and s_{targ} . If the roles of the target and distractors switch, then (ignoring for the time being any changes to contrast features) s_{dist} and s_{targ} swap, x becomes its reciprocal, and the search time slope changes.

Recall that the salience of a feature is proportional to $\frac{1}{p(F)}$. If item A is more unusual than item B , A has a higher salience than B . Hence A makes a stronger distractor for target B ($x > 1$) than B does for A ($x < 1$). Searching for an A in a field of B s will be more efficient than searching for a B among A s. In other words, the less encountered item is easier to find.

This observation easily explains many classical search asymmetries, and we show two examples here:

- The absence of a feature does not pop out. Figure 2 shows an example that a “Q” stands out in “O”s but not vice versa (Treisman & Souther, 1985). It is intuitive that a feature is typically less encountered than its absence. Thus searching for a missing feature is more difficult than searching for the presence of that feature.
- “Prototypes” do not pop out. Prototypes are specific features that are more commonly encountered. For example, the color red, straight lines, and stationary objects are prototypical while magenta, curves, and moving objects are not. In each case, the atypical object pops out in a field of




Figure 2: “Lack of feature” fails to pop out. To find the “Q” in “O”s (left) is easier than to find the “O” in “Q”s (right).

prototypical objects, but not the reverse. Figure 3 shows an example that a tilted line stands out in vertical lines but not vice versa (Treisman & Gormican, 1988). Since prototypes are by definition common, our model readily predicts this asymmetry.




Figure 3: “Prototypes” do not pop out. Finding the tilted (atypical) bar among vertical (prototypical) ones (left) is easier than the reverse (right).

In the face search scenario, this dictates that looking for cross-race faces among self-race faces should be faster than the reverse because CR faces are less frequently encountered. This explains the advantage of the CR faces in visual search directly. It is notable that our model has a straightforward account for the asymmetry in the face search without any parameters or task specific assumptions.

We can provide a quantitative fit of our model to the response time data. The human data presented here are from Figure 9 in (Levin, 1996) and Figure 3 in (Levin, 2000) which examines the response time in face search for white participants. This fit verifies two claims of our model. First, a linear model matches the data well. Secondly, the values assigned to the distractor strength, x , behave as expected, with a distractor strength lower than 1 when searching for the CR face.

Here we concentrate on the situation when the target is present. Our model predicts that in target trials the expected time to locate the target is $E_{n,x} = 1 + \frac{x}{1+x} \cdot n$, with n representing the number of distractors and x the distractor strength. Let t_u denote the unit processing time, i.e. the time in milliseconds to scan each object. Let t_0 represent an additional constant time cost, accounting for other factors such as the time needed for pressing the response button. The response

time in milliseconds under our model is:

$$RT^{model} = E_{n,x} \cdot t_u + t_0 \quad (15)$$

$$= \frac{x}{1+x} \cdot t_u \cdot n + t_u + t_0 \quad (16)$$

Let x_{black} represent the distractor strength when the target is black. Then the distractor strength when the target is white is $x_{white} = \frac{1}{x_{black}}$. Note that $\frac{x_{white}}{1+x_{white}} = \frac{1}{1+x_{black}}$. Assume that search for black and white targets share the same unit process time (t_u) and have their own constant cost ($t_{0,black}$ and $t_{0,white}$). We set different constant costs for these two conditions because when there are no distractors, it boils down to a classification problem and experiments have shown that White subjects classify Black faces faster than White faces (Valentine & Endo, 1992). Now the model predicts the search time:

$$RT_{black}^{model} = \frac{x_{black}}{1+x_{black}} \cdot t_u \cdot n + t_u + t_{0,black} \quad (17)$$

$$RT_{white}^{model} = \frac{x_{white}}{1+x_{white}} \cdot t_u \cdot n + t_u + t_{0,white} \quad (18)$$

$$= \frac{1}{1+x_{black}} \cdot t_u \cdot n + t_u + t_{0,white} \quad (19)$$

where x_{black} , t_u , $t_{0,black}$ and $t_{0,white}$ are free parameters. Our model gives a linear fit to the data. Given a linear model:

$$RT_{black}^{data} = a_{black} \cdot n + b_{black} \quad (20)$$

$$RT_{white}^{data} = a_{white} \cdot n + b_{white} \quad (21)$$

where a_{black} , a_{white} , b_{black} and b_{white} are the slopes and the y-intersections of the best linear fit to the human response time in searching for the black face and white face respectively, the corresponding model parameters are related as follows:

$$x_{black} = \frac{a_{black}}{a_{white}} \quad (22)$$

$$t_u = a_{white} \cdot (1 + x_{black}) \quad (23)$$

$$t_{0,black} = b_{black} - t_u \quad (24)$$

$$t_{0,white} = b_{white} - t_u \quad (25)$$

Figure 4 shows the fitting results, and the corresponding parameters are shown in Table 1. Note that the distractor strength in black search x_{black} is less than 1, which means that the black face (target) is more salient than the white faces (distractor) to white subjects.

Discussion

The face search asymmetry is of special interest to us for two reasons. First, the race effect in visual search is an interesting phenomena in the domain of face processing. Second, in visual search, the direction of this search asymmetry is dependent on prior experience. This asymmetry is therefore a




Figure 4: Our model is fit to the reaction times of searching for white and black face targets for white participants. Human data is from (Levin, 1996) and (Levin, 2000). The fit parameters are given in table 1.

Table 1: The fitted parameters.

	data@1996	data@2000
a_{black}	53.00	61.77
b_{black}	584.67	721.62
a_{white}	76.88	93.35
b_{white}	657.29	801.08
distractor strength x_{black}	0.69	0.66
distractor strength x_{white}	1.45	1.51
unit scan time t_u (ms)	129.88	155.12
constant cost $t_{0,black}$ (ms)	454.79	566.5
constant cost $t_{0,white}$ (ms)	527.42	645.96

useful set of human data with which to test visual search theories and models.

Our probabilistic model predicts that less commonly occurring visual features are more likely to attract attention. This gives rare items an advantage in visual search, which leads to a search asymmetry; it is faster to find a rare target among common distractors than vice versa. This accounts for many classical search asymmetries and we showed two examples: 1) it is easier to find a feature positive target among feature negative distractors than the reverse 2) it is easier to find a non-prototypical target in prototype distractors. This model also explains the search asymmetry that it is faster to find the CR face among SR faces than the reverse; this occurs for the simple reason that people experience SR faces much more than CR faces. Comparing this model to Levin's hypothesis, which introduces a race feature, reveals that our model is more principled in that it makes no specific assumptions for a particular experiment. Our model also explains how the CR search advantage and CR recognition deficit would vanish given exposure of the right sort to CR faces. Because no face specific assumptions are needed, this suggests that faces are not special in the visual search domain, but follow the same principles as simple stimuli used in classical visual search paradigms.

It is important to note that although our model dictates that rare items have advantages in visual search, we did not specify at which level the CR faces are pre-attentively perceived (as described by F) to be rare in the face search, resulting in

the attraction of attention. Could it be at the face level? This is the highest possible level; here, the race of the faces are perceived first and attention is then attracted to the CR face. Or could it be at the internal feature level, the intermediate level? Did the subjects notice the eyes and the mouth are special in the CR faces? Potentially, the perceptual representation could be lower still, such as with the frequency domain, which can be processed in parallel in low-level visual systems. Do CR faces reside in a different area in the frequency space, resulting in attention being attracted to these outliers? We think that the current data is insufficient to fully answer this question, although some other search asymmetries related to familiarity seem to suggest that low level features can not account for everything (Wang, Cavanagh, & Green, 1994; Shen & Reingold, 2001).

At the same time, our model is the only visual search model that relates the search speed to the visual system's experience. Thus we believe this set of human data where the search asymmetry is dependent on the race, or rather, subjects' experience with race, favors our model over others; our model accounts for the asymmetry straightforwardly, while other models need extra assumptions (such as an additional race feature) to accommodate the data.

Acknowledgments

The authors would like to thank Daniel Levin for sharing the experiment data. We also thank Dan N. Hill, Honghao Shan and Tim K. Marks for helpful discussions, and everyone in GURU (Gary's Unbelievable Research Unit) for comments. This work is supported by NIMH grant MH57075 to GWC.

References

- Bothwell, R., Brigham, J., & Malpass, R. (1989). Cross-racial identification. *Personality and Social Psychology Bulletin*, 15(1), 19.
- Chance, J., Turner, A., & Goldstein, A. (1982). Development of differential recognition for own-and other-race faces. *Journal of Psychology*, 112(1st), 29–37.
- Feingold, G. (1914). The influence of environment on identification of persons and things. *Journal of the American Institute of Criminal Law and Criminology*, 5(1), 39–51.
- Haque, A., & Cottrell, G. W. (2005). Modeling the other-race advantage with pca. In *Proceedings of the 27th annual cognitive science conference*. La Stresa, Italy: Mahwah: Lawrence Erlbaum.
- Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, 40(10-12), 1489-1506.
- Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11), 1254-1259.
- Levin, D. (1996). Classifying faces by race: The structure of face categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1364–1382.
- Levin, D. (2000). Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *Journal of Experimental Psychology: General*, 129(4), 559–574.
- Li, J., Dunning, D., & Malpass, R. S. (1998). Cross-racial identification among european-americans: Basketball fandom and the contact hypothesis. In *Paper presented at the american psychology - law society*. Redondo Beach, CA.
- O'Toole, A., Deffenbacher, K., Abdi, H., & Bartlett, J. (1991). Simulating the "other-race effect" as a problem in perceptual learning. *Connection Science*, 3(2), 163–178.
- O'Toole, A., Deffenbacher, K., Valentin, D., & Abdi, H. (1994). Structural aspects of face recognition and the other-race effect. *Memory and Cognition*, 22(2), 208–24.
- Shapiro, P., & Penrod, S. (1986). Meta-analysis of face identification studies. *Psychological Bulletin*, 100, 139–156.
- Shen, J., & Reingold, E. (2001). Visual search asymmetry: The influence of stimulus familiarity and low-level features. *Perception & Psychophysics*, 63(3), 464–75.
- Shepard, J. (1981). Social factors in face recognition. In G. Davies, H. Ellis, & J. Shepard (Eds.), *Perceiving and remembering faces* (p. 55-79). London, UK: Academic Press.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: evidence from search asymmetries. *Psychological Review*, 95(1), 15-48.
- Treisman, A., & Souther, J. (1985). Search asymmetry: a diagnostic for preattentive processing of separable features. *Journal Experimental Psychology: General*, 114(3), 285-310.
- Valentine, T., Chiroro, P., & Dixon, R. (1995). An account of the own-race bias and the contact hypothesis based on a "face space" model of face recognition. *Cognitive and computational aspects of face recognition: Explorations in face space*, 69–94.
- Valentine, T., & Endo, M. (1992). Towards an exemplar model of face processing: the effects of race and distinctiveness. *Quarterly Journal of Experimental Psychology: Human Experimental Psychophysics*, 44(4), 671–703.
- Wang, Q., Cavanagh, P., & Green, M. (1994). Familiarity and pop-out in visual search. *Perception & Psychophysics*, 56(5), 495–500.
- Wolfe, J. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, 1(2), 202-238.
- Wolfe, J. (1998). Visual search. In H. Pashler (Ed.), *Attention* (pp. 13–73). London, UK: University College London Press.
- Wolfe, J. (2001a). Asymmetries in visual search: An introduction. *Perception & Psychophysics*, 63(3), 381–389.
- Wolfe, J. (2001b). Guided search 4.0: A guided search model that does not require memory for rejected distractors. *Journal of Vision*, 1(3), 349-349.