Preprint version. Accepted for publication in Psychological Science.

Eliminating the Low Prevalence Effect in Visual Search with a Remarkably Simple Strategy

J. Eric T. Taylor^{1,2}, Matthew D. Hilchey³, Blaire J. Weidler⁴, & Jay Pratt⁵

1. The Vector Institute for Artificial Intelligence, Toronto, Ontario
2. University of Guelph, School of Engineering, Guelph, Ontario
3. University of Toronto, Rotman School of Management, Toronto, Ontario
4. Towson University, Department of Psychology, Towson, MD
5. University of Toronto, Department of Psychology, Toronto, Ontario

Corresponding Author:

J. Eric T. Taylor

j.eric.t.taylor@gmail.com

The Vector Institute for Artificial Intelligence

MaRS Centre, West Tower

661 University Ave, Suite 710

Toronto, ON, Canada

M5G1M1

Author Contributions: The idea was conceptualized by MDH & JETT. All authors contributed to the design. The experiment was coded by JETT. Data was collected by BJW, MDH, & JETT. Results were analyzed by MDH & JETT. The paper was written by JETT in collaboration with MDH, BJW, & JP.

Abstract

The low prevalence effect in visual search occurs when rare targets are missed at a disproportionately high rate. This effect has enormous significance in health and public safety and has proven resistant to intervention. In three experiments ($Ns = 41, 40, 44$), we document a dramatic reduction of the effect using a simple cognitive strategy requiring no training. Instead of asking participants to search for the presence or absence of a target, as is typically done in visual search tasks, we asked participants to engage in “similarity search” – to identify the display element most similar to a target on every trial, regardless of whether a target is present. Under normal search instructions, we observed strong low prevalence effects. Using similarity search, we failed to detect the low prevalence effect under identical visual conditions across three experiments.

Statement of Relevance

The low prevalence effect in visual search occurs when rare targets go unnoticed at disproportionately high rates. This phenomenon poses serious harm in applied settings, including medical imaging and security screening, where signs of cancer or perilous threats go undetected because they are rare. The effect has also been historically resistant to intervention. In this study, we were able to mitigate the low prevalence effect in a laboratory setting using a simple strategy that requires no training.

Introduction

The low prevalence effect (LPE) in visual search refers to the finding that observers miss rare targets at inordinately high rates (Wolfe et al., 2005). This behaviour has deleterious consequences in circumstances where important targets are rare, including radiological scanning (Evans et al., 2011; Evans et al., 2013) and airport security screening (Wolfe et al., 2013). Developing successful strategies to mitigate the LPE could translate directly into lives saved. In this paper, we failed to detect the LPE after a small change in task instruction and no training, effectively eliminating the LPE across three experiments.

The LPE is a fascinating problem for applied psychologists because it is well-researched and has significant costs to the public yet remains resistant to intervention. Some attempts to reduce the LPE have been motivated by the finding that low target prevalence is correlated with early search termination on target-absent trials and that such searches involve motor or decisional biases toward the 'absent' response (Rich et al., 2008). Surprisingly, deliberately slowing responses, working with a partner (Wolfe et al., 2007), slowly revealing parts of the search display (Kunar et al., 2010), and adjusting expectations of target presence (Lau & Huang, 2010; Reed et al., 2011; Ishibashi et al., 2012) have little or no effect on the LPE. Other interventions have shown mixed results or are not ideal for practical reasons. Correctible search (Fleck & Mitroff, 2007) can mitigate the LPE by allowing participants to reverse their last decision, but this intervention has shown mixed results (Van Wert et al., 2009). Inserting a block of high prevalence searches with feedback, or misleading participants with erroneous feedback can help searchers readjust their decision criteria, improving hit rate, but concomitantly

increasing costly false alarms (Wolfe et al., 2007; Schwark et al., 2012). There is currently no effective intervention for the LPE. The present study poses a novel solution.

Experiment 1

Given that the LPE is thought to be caused by modulation of the target-absent response threshold or the accumulation of information toward this threshold, we hypothesized that it should be possible to mitigate the effect by requiring a target-present response in every search episode. Rather than asking participants to report the presence or absence of a target, as is typically done in search tasks, we required participants to point out the most target-like stimulus on each display. This simple change in search instruction should maintain the expectation for a target on each trial. Rather than manipulating the physical search conditions, our intervention manipulates the rule and response defining a successful search.

Method

Participants

We aimed to analyze data from 40 participants. This sample size is greater than other common demonstrations of the LPE (e.g., Wolfe et al., 2005; 2007; Fleck & Mitroff, 2007) and is sufficient to detect a 20% LPE with 0.95 power and a 0.05 tolerance for Type I errors using a paired-samples *t*-test. Estimates for effect size and standard deviation of the LPE were drawn from Figure 1 of Wolfe et al. (2005), 10% and 50% prevalence conditions on a display with 18 objects, as these were the conditions most similar to those in our study. Sample size forecasts made with G*Power v3.1 (Faul et al., 2007).

Forty-four undergraduate students from the University of Toronto participated in exchange for course credit or \$10. All students gave informed consent according to the

university's IRB. All participants had normal or corrected-to-normal vision. All participants were naïve to the purpose of the study and its hypotheses. Data from four participants were rejected prior to analysis for producing incomplete data sets.

Apparatus & Materials

Stimuli were presented on a Dell computer with 24" LED monitors, using a 2560 x 1440 resolution graphics mode and a 120 Hz refresh rate. The experiment was programmed using MATLAB software with Psychtoolbox (Brainard, 1997). Viewing distance was held constant at 57 cm with a chin rest.

Procedure

The search condition manipulation (present/absent vs. similarity; see below) was a between-subject factor. Participants were assigned to either condition in alternating order. The target prevalence manipulation (low - 10% vs. high - 50%) was blocked within subjects and the order of the blocks was counterbalanced across subjects. There were 250 trials in the low prevalence block and 50 trials in the high prevalence block (i.e., there were 25 target-present trials in each prevalence condition). All participants completed 20 practice trials at the start of the experiment and again before the target prevalence switched to adjust to the target frequency. During practice trials, the words "Practice Trial" appeared centered at the top of the screen. Participants were informed of the target prevalence before each block. In the low prevalence block, participants were encouraged to take a break every 50 trials by an interrupting display.

In both search conditions, every trial contained a search display with 16 elements that were pseudo-randomly positioned on an invisible 8x8 grid to ensure that no array element

occluded another. Element position was randomly selected without replacement from the 8x8 grid and random positional jitter was applied to each by an amount no greater than half of the distance between the adjacent grid position. To an observer, the display appeared as a randomly assorted but uncluttered array on every trial.

All stimuli were presented on a dark grey background (RGB: 40 40 40). Array elements differed in shade of grey (randomly generated on every trial for every element for some RGB value between 100 and 160) and orientation (randomly rotation of 0, 90, 180, or 270 deg). Participants were informed that the shade of grey and orientation were irrelevant.

In both search conditions, every display contained 15 near-T distractors and a 16th element, which was either a perfect T or a near-T that was nearer to perfect than the other 15 elements. Critical stimuli were “perfect” Ts when the stem perfectly bisected the branch of the T at its midpoint. Thus, “near-Ts” occurred when the stem appeared offset from the perfect midpoint of the branch (as generated by random offset between the stem and branch of the T). The offset was applied independently for each distractor and was a minimum of one quarter of the length of the branch off-centre and a maximum of the entire length (resulting in an L). On target-absent trials, instead of a perfect T, the array contained one element that had less offset jitter than the other 15 near-Ts: the offset was equal to 1/11 of the length of the branch (see Figure 1).

Participants were given feedback after every trial. There was a 500ms intertrial interval. The only differences between search conditions were the instructions, response mode, and feedback, described below.

Target Present/Absent Search.

RUNNING HEAD: SIMILARITY SEARCH ELIMINATES LOW PREVALENCE

Participants were instructed that they would be searching for a perfect T in an array of near-Ts, and were shown an example of a display with and without a target. Consistent with traditional LPE visual search tasks, participants were instructed to press 'p' when the target was present and 'a' when the target was absent.

On target-absent trials, a feedback message appeared at the centre of the display saying "Correct! / There was no target in this array" or "Error! / There is no target in this array." "Correct!" and "Error!" were presented in green (RGB: 0 180 20) and red (RGB: 180 0 0) font, respectively. On target-present trials, feedback occurred in the form of a green (correct) or red (incorrect) circle surrounding the target element.

Similarity Search.

Unlike traditional LPE search tasks, participants in this condition were instructed to find the item that most closely resembled a T on every trial. Participants were instructed to press the 'p' key once they had found the "item that most closely resembles a T", which revealed a mouse cursor. Participants then clicked on the item that they had chosen.

Click accuracy was determined by whether the response was made near the correct element (i.e., the perfect-T or the nearest-to-perfect-T) within a radius of 1.5 times its size. If participants correctly identified the most T-like item (a perfect T on target-present trials and a near-T on target-absent trials), a green circle appeared around the element they clicked on. If they identified any other element in the display, a red circle appeared around the correct element.

RUNNING HEAD: SIMILARITY SEARCH ELIMINATES LOW PREVALENCE

Figure 1. Illustrated response and stimuli parameters for all experiments. For the target present/absent condition, a perfect T is the target, and anything else is a distractor. In the similarity search condition, the target is the singular element that is closest to a perfect T (and on 10% or 50% of trials, depending on prevalence block, there *is* a perfect T). Critically, on trials with a perfect T, the search is physically identical between search conditions, allowing a comparison of missed target frequency. Required responses are illustrated below the displays. On present/absent search, the observer simply presses 'p' or 'a'. For similarity search, participants press 'p', revealing a cursor, then mouse over and click on the element they have chosen. In Experiment 2, the stimuli parameters and the similarity search condition were identical, but the present/absent search was modified such that if participants believed the target was present, they were to press 'p', then mouse over and click the target, and if the target was absent they were to press 'p' and mouse over to the edge of the screen. Experiment 3 was a replication of Experiment 1, except that trials without a perfect T contained 16 randomly determined near-Ts, with the one closest to a T being the target. Consequently, observers could not rely on a stored target template for the near-T. Image and display elements are not to scale.

Results

The critical analysis concerns whether perfect Ts are missed at different rates across prevalence depending on the search instruction conditions (i.e., the search instruction by prevalence interaction). Miss rates on trials where a perfect-T was present were analyzed with a 2 (search instruction: present/absent or similarity; between-subjects) X 2 (target prevalence: 10% or 50%; within-subjects) mixed-factors analysis of variance (ANOVA). For target present/absent search, a miss occurs when there is a perfect T in the display and participants indicate the target is absent. For similarity search, a miss occurs when there is a perfect T in the display and participants indicate that another element was the most T-like. Critically, optimal performance in both conditions requires that participants find the perfect T, so performance is comparable between search conditions on perfect-T present trials.

There were main effects of search condition, $F(1,39) = 13.29, p < .001, \eta_p^2 = 0.25$, and target prevalence, $F(1,39) = 19.10, p < .001, \eta_p^2 = 0.33$. Critically, there was a significant interaction, $F(1,39) = 9.99, p = .003, \eta_p^2 = 0.204$, indicating that the LPE depended on search condition (see Figure 2). To characterize this interaction, we ran post hoc paired-samples two-tailed t -tests that compared miss rates between high and low prevalence at each level of search condition, applying Bonferroni corrections for multiple comparisons. These allow us to evaluate the magnitude of the LPE in each search condition and, importantly, to evaluate whether we observed an LPE in the similarity search condition. As expected, the results showed a robust LPE for target present/absent search (20.50% difference), replicating established findings: $t(23) = 6.02, p < .001, \alpha = .025$. The LPE effect was reduced to 3.3% in the similarity condition, which was not statistically significant: $t(16) = 0.76, p = .46, \alpha = .025$. For performance

data from trials without a perfect T and a note on response times, please consult the Supplementary Materials.

Figure 2. (A) Mean miss rate of perfect Ts as a function of target prevalence and search condition. In Experiment 1, we observed a robust LPE for typical target present/absent searches, which was not present for similarity search. Despite having added a mouse interaction and localization response to more closely match response demands of the similarity search, we still observed the LPE in the modified present/absent search of Experiment 2 while replicating the mitigated LPE in similarity search. In Experiment 3, we replicated the effect despite removing the unique nearest-to-perfect T in trials without a perfect T. Error bars represent one standard error of the mean. *** $p < .001$; ** $p < .01$; *n.s.* indicates not significant. (B) Effect of search Condition (Present/Absent vs. Similarity) on the Low Prevalence Effect (LPE) in each experiment and meta-analysis (Overall). Error bars are 95% confidence intervals (95%CI). Similarity search reduced the LPE by 13.64 percentage points [95%CI = 7.15%, 20.12%; see “Overall” Line]. Overall, we estimate an LPE of 16.89% (95%CI = 11.96% - 21.83%) in the standard present/absent search task. We estimate an LPE of 3.25% (95%CI = -0.93% - 7.43%) in similarity search, which was not statistically significant.

Experiment 2 sought to reproduce the mitigated LPE with similarity search and to determine whether differences in response modes between search conditions – localized mouse responses versus a ‘p’/‘a’ button-press – may have played a role.

Method

Participants

Forty-four additional undergraduate students from the University of Toronto who met the same criteria outlined in Experiment 1 participated in exchange for course credit or \$10. Data from three participants were removed prior to analysis for producing incomplete data sets.

Procedure

The procedure was identical to Experiment 1 with the following exception: In the present/absent search condition, when a perfect T was present, participants were required to press the ‘p’ button then mouse over and click the target (exactly as in the similarity search condition); if the target was absent, participants were required to press ‘p’, then move the mouse to the edge of the screen. We call this the “present-click/absent-flick” condition going forward (see Figure 1). No changes were made to the similarity search condition. Critically, these search and response conditions are identical when a perfect T is present in the display.

Results

Data from one participant were removed prior to analysis for responding randomly during present/absent search (48% accuracy in the high prevalence condition). We conducted the same analyses as in Experiment 1. There were main effects of search condition, $F(1,38) = 30.15$, $p < .001$, $\eta_p^2 = 0.44$, and target prevalence, $F(1,38) = 6.53$, $p = .015$, $\eta_p^2 = 0.15$. The interaction between search condition and target prevalence was marginally significant, $F(1,38) =$

3.82, $p = .058$, $\eta_p^2 = 0.09$, indicating that the LPE depended on search condition, although that interpretation should be made with caution (see Figure 2). The critical test is for the presence of LPEs in either condition. We performed post hoc paired-samples two-tailed t -tests comparing miss rates between high and low prevalence at both levels of search condition, applying Bonferroni corrections for multiple comparisons. Replicating Experiment 1, the similarity search condition showed virtually no LPE (1.6% difference in miss rates between high and low prevalence), which did not reach statistical significance: $t(19) = 0.43$, $p = 0.67$, $\alpha = .025$. The present-click/absent-flick condition revealed a LPE, (12.0% difference), $t(19) = 3.10$, $p = .006$, $\alpha = .025$.

Experiment 3

In Experiments 1 and 2, when a perfect T was not present, there was a near-T that was nearer-to-perfect than the rest of the distractors (its branch offset at 1/11 of the length of the T branch). An astute reviewer noted that the presence of this element was a signal that observers might use differently between search conditions. Observers in similarity search could use a template for this nearest-to-perfect-T to find their target, whereas observers in the present/absent search might use the same template to correctly terminate a search. Essentially, observers could complete the task with two templates in mind (e.g. Irons et al., 2012; Roper & Vecera, 2012), one of which was always present, effectively making this a 100% prevalence task. To address this, we repeated Experiment 1 under conditions where observers could not store a template for the nearest-to-perfect T. On trials without a perfect T, all elements were fully randomly generated, and the one with the least offset was designated the nearest-to-perfect T (and hence the target in similarity search; see Figure 1). By introducing variability to the

nearest-to-perfect-T stimuli, we made it impossible to complete the task by maintaining two target templates.

Method

Participants

Forty-four new participants were recruited online via Prolific in exchange for £5GPB/hr. The same inclusion criteria were used from Experiments 1 and 2. The experiment was coded in Javascript and hosted on a custom server.

Procedure

The procedure was identical to Experiment 1 with the following exceptions. First, whereas in Experiments 1 and 2 on trials without a perfect T there were 15 randomly generated distractors and one distractor that was uniquely close to a T (1/11 offset), in Experiment 3, all 16 distractors were randomly generated and the one with the least offset was designated the nearest-to-perfect-T (and therefore the target in the similarity search condition). There was only ever one nearest-to-perfect T. Second, the experiment was moved online due to the COVID-19 pandemic. This also afforded us the opportunity to replicate our findings outside of the laboratory in a population of non-students.

Results

We conducted the same analyses as in Experiment 1. There was a main effect of target prevalence, $F(1,42) = 12.58$, $p < .001$, $\eta_p^2 = .23$. There was no effect of task, $F(1,42) = 0.215$, $p = 0.645$. The interaction between search condition and target prevalence was significant, $F(1,42) = 4.13$, $p = .048$, $\eta_p^2 = .09$, replicating the effect that the LPE depended on search condition (see Figure 2). We performed post hoc paired-samples two-tailed t -tests comparing miss rates

between high and low prevalence at both levels of search condition, applying Bonferroni corrections for multiple comparisons. Replicating Experiments 1 & 2, the similarity search condition did not elicit a reliable difference between high and low prevalence (4.7% difference in miss rates): $t(21) = 1.49, p = .152, \alpha = .025$. The present/absent condition revealed a LPE, (17.4% difference), $t(21) = 3.24, p = .004, \alpha = .025$. The increased error rates bear mentioning. Because of the COVID-19 pandemic, Experiment 3 was conducted online out of necessity. Given our experience with many other online studies, and from what we have heard from colleagues, higher than laboratory error rates are ubiquitous, and we are sure this is the case here. Notwithstanding, we stand by the result given that the predicted pattern of the interaction was replicated despite higher overall error rates. Indeed, it would be difficult to account for why the change to distractor settings would allow for the LPE to be preserved in present/absent search but not similarity search for any other reason. Importantly, even when baseline miss rates between present/absent and similarity search are incidentally similar (as in Exp 3), only the present/absent search shows a robust LPE.

Discussion

The present study documents a remarkably simple solution to the LPE in visual search in a laboratory setting. Whereas past efforts have tried to reduce the LPE by changing physical search conditions or the context of the search, we changed the search task and response requirements. In addition to a typical present/absent search, which elicited the usual LPE, we asked another group of participants to search for the most target-like object on *every* trial. This simple strategic adjustment eliminated the LPE¹. All three experiments provided evidence that

¹ Obviously, we can never confirm a null result. But we eliminated the LPE insofar as we detected it reliably during present/absent search and consistently failed to detect it under equal power in three experiments in similarity search.

RUNNING HEAD: SIMILARITY SEARCH ELIMINATES LOW PREVALENCE

similarity search reduces the LPE relative to present/absent search, and planned comparisons failed to find evidence for an LPE in similarity search every time and when pooling data across experiments to increase statistical power.

The similarity search task differs from the typical present/absent paradigm in some important ways. Of course, similarity search involves mouse localization responses whereas the typical paradigm involves a binary keypress. However, Experiment 2 rules out the possibility that mouse localization responses eliminated the LPE. Experiment 3 replicated the effect under conditions where the target on one trial could subsequently be a distractor, ruling out the possibility that a singular template for the near-T could guide search.

We suspect that similarity search mitigated the LPE by either maintaining a high search termination threshold, reducing the accumulation of evidence toward the threshold, or both. Our suspicion can be understood in the context of Wolfe and Van Wert's (2010) two-process model for the LPE. The first process is perceptual, and it concerns the accumulation of evidence for the presence or absence of a target (i.e., a signal detection problem; MacMillan & Creelman, 2004). In this model, it is possible that similarity search affects the position of a decision criterion with respect to the target-present signal. The second process is decisional. A quitting signal accumulates when perceptual comparisons fail to produce a target-present response; when the threshold is exceeded, the search is terminated (e.g., Moran et al., 2013). In this model, similarity search would diminish the accumulation of the quitting signal or shift the quitting threshold upward. Sequential target-absent decisions cannot shift the quitting threshold downward if there are no target-absent decisions. Recall that the task in similarity search is not to decide *whether* a target is absent but *which* element deserves a response.

This intervention has the potential to eliminate the LPE in high-stakes, real-world visual search scenarios. Similarity search could be easily and inexpensively applied to any search task where the goal can be reframed as a graded rather than binary comparison: What element of this mammogram looks most like a pre-cancerous lesion? What object in this baggage x-ray looks most like a weapon? However, there are potential limitations to our intervention that could diminish its real-world applicability. For example, security screening involves a search with heterogeneous target templates (e.g., knife, explosive substance, unsanctioned liquid) and heterogeneous distractors, whereas our laboratory study involved a search for a single target template (the “perfect” T) in a comparatively homogenous set of distractors. The current findings do not speak to whether observers can judge similarity across more abstract templates (e.g., “most dangerous”), or whether similarity search would be effective at reducing the LPE under such conditions (e.g. Godwin et al., 2010).

Considerations also arise in the application of similarity search to medical imaging. When scanning for signs of cancer or injury, identifying similarity is not sufficient; medical experts must then decide whether the stimulus poses a genuine threat (i.e., it may be cancer-like, but is it really cancer?). Likewise, similarity search must be shown to eliminate the LPE without a concomitant rise in false alarms. In practice, a two-tiered decision procedure could alleviate both concerns. Following an initial judgment of similarity, a second step would ask observers to re-examine each suspicious element (e.g., the element designated most similar to a lesion or threat) to determine whether it poses a genuine threat. In the second step however, there is no guarantee that the LPE will be absent or that the false alarm rate will be acceptable. In addition, it is worth noting that our one-hour experiment was brief compared to

the prolonged vigilance demanded of security and medical experts (despite even lower target prevalence rates, <1% vs. 10% in our study).

Other important empirical questions ought to be explored. It is unclear, for example, whether similarity search depends on having a singular pseudo-target; would it still work in cases where observers view multiple most-similar items, as in a multiple-target search task? Understanding these and other potential issues in implementation, including the human resource costs, will be important for extrapolating this promising laboratory-based intervention to applied settings. Although work remains to be done, this study is the important first step of identifying a promising solution to a stubbornly persistent behavioural phenomenon that has wide ranging impact beyond the laboratory.

References

- Brainard, D. H. (1997) The Psychophysics Toolbox. *Spatial Vision*, 10, 433-436.
- Evans, K. K., Tambouret, R. H., Evered, A., Wilbur, D. C., & Wolfe, J. M. (2011). Prevalence of abnormalities influences cytologists' error rates in screening for cervical cancer. *Archives of pathology & laboratory medicine*, 135(12), 1557-1560.
- Evans, K. K., Birdwell, R. L., & Wolfe, J. M. (2013). If you don't find it often, you often don't find it: Why some cancers are missed in breast cancer screening. *PloS one*, 8(5), e64366.
- Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior research methods*, 39(2), 175-191.
- Fleck, M. S., & Mitroff, S. R. (2007). Rare targets are rarely missed in correctable search. *Psychological science*, 18(11), 943-947.
- Godwin, H. J., Menneer, T., Cave, K. R., Helman, S., Way, R. L., & Donnelly, N. (2010). The impact of relative prevalence on dual-target search for threat items from airport X-ray screening. *Acta psychologica*, 134(1), 79-84.
- Irons, J. L., Folk, C. L., & Remington, R. W. (2012). All set! Evidence of simultaneous attentional control settings for multiple target colors. *Journal of Experimental Psychology: Human Perception and Performance*, 38(3), 758.
- Ishibashi, K., Kita, S., & Wolfe, J. M. (2012). The effects of local prevalence and explicit expectations on search termination times. *Attention, Perception, & Psychophysics*, 74(1), 115-123.
- Kunar, M. A., Rich, A. N., & Wolfe, J. M. (2010). Spatial and temporal separation fails to counteract the effects of low prevalence in visual search. *Visual cognition*, 18(6), 881-897.
- Lau, J. S. H., & Huang, L. (2010). The prevalence effect is determined by past experience, not future prospects. *Vision Research*, 50(15), 1469-1474.
- Macmillan, N. A., & Creelman, C. D. (2004). Detection theory: A user's guide. Psychology press.
- Manassi, M., Kristjánsson, Á., & Whitney, D. (2019). Serial dependence in a simulated clinical visual search task. *Scientific Reports*, 9(1), 1-10.
- Moran, R., Zehetleitner, M., Müller, H. J., & Usher, M. (2013). Competitive guided search: Meeting the challenge of benchmark RT distributions. *Journal of Vision*, 13(8), Article 24.

- Peltier, C., & Becker, M. W. (2017). Target-present guessing as a function of target prevalence and accumulated information in visual search. *Attention, Perception, & Psychophysics*, 79(4), 1064-1069.
- Reed, W. M., Ryan, J. T., McEntee, M. F., Evanoff, M. G., & Brennan, P. C. (2011). The effect of abnormality-prevalence expectation on expert observer performance and visual search. *Radiology*, 258(3), 938-943.
- Rich, A. N., Kunar, M. A., Van Wert, M. J., Hidalgo-Sotelo, B., Horowitz, T. S., & Wolfe, J. M. (2008). Why do we miss rare targets? Exploring the boundaries of the low prevalence effect. *Journal of Vision*, 8(15), 15-15.
- Roper, Z. J., & Vecera, S. P. (2012). Searching for two things at once: Establishment of multiple attentional control settings on a trial-by-trial basis. *Psychonomic Bulletin & Review*, 19(6), 1114-1121.
- Schwark, J., Sandry, J., & Dolgov, I. (2013). Evidence for a positive relationship between working-memory capacity and detection of low-prevalence targets in visual search. *Perception*, 42(1), 112-114.
- Schwark, J., Sandry, J., MacDonald, J., & Dolgov, I. (2012). False feedback increases detection of low-prevalence targets in visual search. *Attention, Perception, & Psychophysics*, 74(8), 1583-1589.
- Van Wert, M. J., Horowitz, T. S., & Wolfe, J. M. (2009). Even in correctable search, some types of rare targets are frequently missed. *Attention, Perception, & Psychophysics*, 71(3), 541-553.
- Wolfe, J. M., Brunelli, D. N., Rubinstein, J., & Horowitz, T. S. (2013). Prevalence effects in newly trained airport checkpoint screeners: Trained observers miss rare targets, too. *Journal of Vision*, 13(3), 33-33.
- Wolfe, J. M., Horowitz, T. S., & Kenner, N. M. (2005). Cognitive psychology: rare items often missed in visual searches. *Nature*, 435(7041), 439.
- Wolfe, J. M., Horowitz, T. S., Van Wert, M. J., Kenner, N. M., Place, S. S., & Kibbi, N. (2007). Low target prevalence is a stubborn source of errors in visual search tasks. *Journal of Experimental Psychology: General*, 136(4), 623.
- Wolfe, J. M., & Van Wert, M. J. (2010). Varying target prevalence reveals two dissociable decision criteria in visual search. *Current Biology*, 20(2), 121-124.

Supplementary Materials

It is not possible to compare directly between search instruction conditions on trials without a perfect T. This is because a trial without a perfect T in the present/absent search results in either a false alarm or correct rejection, whereas a trial without a perfect T in similarity search results in either a hit or a miss. This asymmetry precludes meaningful cross-condition comparisons. For completeness, Table 1 provides the descriptive statistics for perfect-T absent trials separately for both conditions.

We attempted to operationalize how long it took to find a target as the time from onset of the search array until participants pressed the 'p' button in the similarity search condition, or until they pressed 'p' or 'a' in the present/absent search condition. Unfortunately, participants in the similarity search condition adopted strategies that made this measurement meaningless. For example, some participants pressed 'p' almost immediately after the onset of the search display (i.e., many RTs were < 200 ms). Others held down the 'p' button between trials, as evidenced by some 0 ms RTs. These response times do not genuinely reflect how long it took to find the most t-like stimulus.

RUNNING HEAD: SIMILARITY SEARCH ELIMINATES LOW PREVALENCE

<u>Present/Absent Search</u>	<u>Target Prevalence</u>	<u>Perfect T Present</u>		<u>Perfect T Absent</u>	
		<u>H</u>	<u>FA</u>	<u>H</u>	<u>FA</u>
Experiment 1	High (50%)	0.75	--	--	0.13
	Low (10%)	0.55	--	--	0.03
Experiment 2	High (50%)	0.69	--	--	0.05
	Low (10%)	0.57	--	--	0.02
Experiment 3	High (50%)	0.67	--	--	0.17
	Low (10%)	0.51	--	--	0.05
<u>Similarity Search</u>					
Experiment 1	High (50%)	0.81	--	0.68	--
	Low (10%)	0.78	--	0.62	--
Experiment 2	High (50%)	0.82	--	0.65	--
	Low (10%)	0.80	--	0.63	--
Experiment 3	High (50%)	0.63	--	0.61	--
	Low (10%)	0.58	--	0.55	--

Table 1. Hit and False Alarm rates in all experiments. The table is divided into trials with and without a perfect T to better illustrate the difference in task demands between present/absent and similarity search. Specifically, Present/Absent search can be separated into the classic formulation of the signal detection matrix (hits and false alarms), whereas similarity search has two different kinds of hits – they have a ‘target’ on every trial. Note that the tasks in Similarity Search and both versions of the Present/Absent Search are fundamentally different when there is no perfect T in the display. In Present/Absent Search without a perfect T, participants are asked to make a negative judgment – to declare that the display does *not* contain a target. A false alarm in this case is if participants declared that the display contained a perfect T when there were none present. In Similarity Search without a perfect T, participants are asked to identify the array element that most closely resembles a perfect T. This task is clearly very different and its difficulty is reflected in the performance, which is graded as a hit when they identify the element closest to a T. There is no response that qualifies as a false alarm in this context.