

Improving Human-Machine Cooperative Visual Search With Soft Highlighting

RONALD T. KNEUSEL and MICHAEL C. MOZER, University of Colorado

Advances in machine learning have produced systems that attain human-level performance on certain visual tasks, e.g., object identification. Nonetheless, other tasks requiring visual expertise are unlikely to be entrusted to machines for some time, e.g., satellite and medical imagery analysis. We describe a human-machine cooperative approach to visual search, the aim of which is to outperform either human or machine acting alone. The traditional route to augmenting human performance with automatic classifiers is to draw boxes around regions of an image deemed likely to contain a target. Human experts typically reject this type of hard highlighting. We propose instead a soft highlighting technique in which the saliency of regions of the visual field is modulated in a graded fashion based on classifier confidence level. We report on experiments with both synthetic and natural images showing that soft highlighting achieves a performance synergy surpassing that attained by hard highlighting.

CCS Concepts: • **Human-centered computing** → **Empirical studies in visualization**; *Geographic visualization*; Heat maps; • **Computing methodologies** → **Perception**; *Supervised learning*;

Additional Key Words and Phrases: Visual search, soft highlighting, target localization

ACM Reference format:

Ronald T. Kneusel and Michael C. Mozer. 2017. Improving Human-Machine Cooperative Visual Search With Soft Highlighting. *ACM Trans. Appl. Percept.* 15, 1, Article 3 (September 2017), 21 pages.
<https://doi.org/10.1145/3129669>

1 INTRODUCTION

In professions as diverse as diagnostic radiology, satellite imagery analysis, and airport baggage screening, individuals are tasked with examining complex visual images in search of target elements. With sufficient training, individuals become skilled at such tasks, but expertise takes years to acquire and human performance is fallible due to lapses of attention. To improve the accuracy and reliability of human performance, *computer-assisted detection (CAD)* systems have been developed that digitally modulate images to increase the saliency of target elements. Transformations such as histogram equalization, smoothing, sharpening, unsharp masking, and homomorphic filtering can guide viewer attention while preserving critical image content (Bankman 2008; Kaur et al. 2011; Richards and Jia 1999; Yasmin et al. 2012).

This research was supported by NSF grants SES-1461535, DRL-1631428, SBE-0542013, and SMA-1041755.

Authors' addresses: R. T. Kneusel and M. C. Mozer, Department of Computer Science, University of Colorado, Boulder, CO 80309-0430; emails: ron@kneusel.org; mozer@colorado.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

2017 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 1544-3558/2017/09-ART3 \$15.00

<https://doi.org/10.1145/3129669>

Complementing these image-based transformations is the use of automatic pattern recognition and classification techniques (hereafter, *classifiers*) to identify candidate target locations. Classifiers have a history in medical imaging, where increased demands on radiologists necessitate means of optimizing performance. They are used both in CT and x-ray imagery and commonly in mammography and lung-nodule detection. The CAD system indicates regions of interest (ROIs) in the image where a human expert should look, either by outlining the regions with circles or boxes, pointing to the region with an arrow, or drawing a symbol in the proximity of the region to indicate the nature of the abnormality (Cheng et al. 2003; Freer and Ulissey 2001; Giger et al. 1993; Lee et al. 2000; Philpotts 2009; Samulski et al. 2009, 2010). The form of these *highlights* has not changed significantly in the past 20 years. They are in widespread use today and U.S. federal health programs reward physicians for using CAD. Although their efficacy has been demonstrated in laboratory studies (Alberdi et al. 2004; Balleyguier et al. 2005; Freer and Ulissey 2001; Sahiner et al. 2009), clinical studies are often discouraging (Philpotts 2009), and a recent large-scale study of practicing radiologists found that CAD does not improve diagnostic accuracy of mammography on any metric assessed (Lehman et al. 2016).

Highlighting is also common in remote-sensing applications involving aerial and satellite-based imagery of the earth. To assist human analysts, classifiers assign a class label to each pixel of the image conditioned on the local spatial patch centered on the pixel. These class labels indicate geographic features (e.g., mountains versus rivers versus land, agricultural versus urban versus undeveloped areas) and are represented to the human via *thematic maps* in which false pixel coloring indicates the most likely class (Anderson 1971; Van Genderen et al. 1978). The thematic map is often presented side-by-side with the original image. Various techniques have been developed to indicate classifier uncertainty, including: the use of whitening (interpolating toward white) (Hengl and Toomanian 2006) or manipulation of saturation (MacEachren et al. 2005) to represent confidence, the use of a display toggle to alternate between the thematic map and a confidence map (Van der Wel et al. 1998), and stochastic animations in which classifier probabilities dynamically resample labelings (Bastin et al. 2002).

There appears to be a dissociation in the CAD literature between (1) tasks involving search for one or a small number of visual elements in an image, typically medical images, and (2) tasks involving characterization of spatial topography with multiclass labeling of individual pixels, typically geospatial-sensing images. In the latter tasks, researchers have investigated whether human analysts benefit from interfaces that encode classifier uncertainty, yet in the former tasks, we are not aware of existing software tools that provide an explicit representation of uncertainty.

In this article, we conduct studies to determine whether observers benefit from visualizations of classifier uncertainty in *search* tasks. We contrast the *hard* highlighting that has been used to indicate candidate targets—boxes enclosing and arrows pointed at regions of interest—with *soft* highlighting that indicates not only target locations but classifier confidence. Figure 1 illustrates these two types of highlights. The soft highlights provide a graded saliency cue related to classifier confidence. We are aware of only one other study that contrasts soft and hard highlights: the very recent work of Cunningham et al. (2016). We return to this research in the General Discussion.

For visual search and identification tasks, we hypothesize that CAD systems can be significantly improved by incorporating soft highlighting over hard highlighting. To make this argument, we briefly discuss issues arising with hard highlighting.

1.1 Hard Highlighting: There is No Free Lunch

In Krupinski et al. (1993), radiology residents briefly inspected images in which circles were used to highlight an ROI in chest x-rays. The presence of highlights improved discrimination of tumors from non-tumors in the ROI (see also Freer and Ulissey (2001)). However, the highlights yielded poorer detection of tumors outside the ROI. In this study, targets appeared in the highlighted region on half of all trials. Consequently, the highlight provided no useful information about target presence. Thus, the study served to evaluate the effect of highlighting on the allocation of attention.



Fig. 1. (left) Satellite image of an urban area; (middle) soft highlighting, where the saturation indicates the likelihood of a target; (right) hard highlighting, where a box is drawn around regions highly likely to contain a target.

Zheng et al. (2001) tested board-certified radiologists and manipulated true-positive (TP) and false-positive (FP) rates of the classifier. With a highly accurate classifier—a TP rate of 0.9 and a mean of 0.5 FP highlights per image—highlights improved discrimination performance. However, with relatively poor classifiers—having either a lower TP rate or a higher FP rate—highlights decreased performance, mainly in terms of increased failure to detect non-highlighted targets.

Zheng et al. (2004) found little effect of highlighting if the highlight onset follows an initial interpretation by the radiologist. However, if the highlights appear simultaneously with the image, performance is impaired, both in a detection task and in a benign/malignant classification task. As in earlier studies, performance deficits manifest in terms of a lower rate of detecting nonhighlighted targets. Further evidence that FP highlights misdirect individuals away from discovering an unmarked target, Alberdi et al. (2004) discovered that highlighting the wrong location resulted in twice as many missed targets as failing to highlight any location.

Mirroring studies conducted with medical personnel on medical imagery, Drew et al. (2012) performed a traditional laboratory search task in which participants located a target letter T among L's in displays with cloud-like noise whose spatial-frequency distribution matched that of mammographic images. A no-highlighting condition was compared to a condition with highlights from a simulated CAD system with TP and FP rates of 0.75 and 0.10, respectively. When candidate targets were masked by the noise, highlights improved sensitivity at the highlighted locations but decreased sensitivity in non-highlighted locations. Fixation tracking confirmed that highlights narrowed the search area.

The clear message from these studies is that there is no free lunch: hard highlights have costs that can outweigh benefits (D'Orsi 2001). Trustworthiness of the CAD is key (Dzindolet et al. 2003). When the CAD appears unreliable, individuals expend cognitive resources actively ignoring the assistance; when the CAD seems reliable, individuals may become overly reliant and overlook non-highlighted targets (Parasuraman and Riley 1997). One solution is to completely automate the search task, yet despite the recent evolution of machine learning methods (LeCun et al. 2015), medical and remote-sensing experts are not yet willing or able to hand the reins to automated methods. The best we can aim for at present is human-machine synergies.

Soft highlighting has the potential to overcome two key limitations of hard highlighting. First, individuals can come to trust a well calibrated but imperfect classifier, because soft highlights indicate graded confidence not a binary decision. Second, individuals can better integrate bottom-up guidance from the highlights with their own top-down strategic search, because soft highlights provide graded saliency cues. The potential of soft highlighting depends on individuals being able to interpret and leverage confidence measures obtained from a classifier. Evidence supporting this ability is found in studies of interactive CAD systems that can be queried for




Fig. 2. Examples of synthetic digit arrays used in Experiments 1–4. (a) A control image with no highlighting of targets. (b) Highlighting with discriminability $d' = 0.75$, (c) $d' = 3.07$, (d) $d' = 5.77$.

Table 1. Three Equivalent Measures of the Quality of Stochastic Oracle-based Classifiers (SOCs) Used in Experiments 1–4

θ	1.00	1.50	1.74	1.90	2.33	4.00	9.00	19.00
EER	0.50	0.36	0.30	0.27	0.20	0.06	0.002	0.000
d'	0.00	0.75	1.00	1.25	1.69	3.07	5.77	9.24

a suspiciousness or malignancy score (Samulski et al. 2009, 2010). The availability of this information improves discrimination performance over a traditional CAD system (Hupse et al. 2013).

2 EXPERIMENTS WITH SOFT HIGHLIGHTING IN SYNTHETIC DISPLAYS

We report on a series of seven experiments investigating highlighting in visual search. The experiments are aimed at determining whether individuals exploit the information in soft highlights and whether this information leads to superior performance over the use of hard highlights and a no-highlight control. Experiments 1–4 test synthetic displays and simulated classifiers of various strengths. Experiments 5–7 test naturalistic images and an actual state-of-the-art neural network classifier.

2.1 Experiment 1: Search for Multiple Targets

In Experiment 1, participants searched a 10×10 array of handprinted digits for all instances of a target—the digit 2. Figure 2(a) shows a sample array, which consists of ten tokens of each of the digits 0–9 placed in random locations. Participants click on each target, causing the target to disappear, and a trial continues until all ten targets have been localized or until 15s have passed. We compare a control condition with no highlighting to conditions in which elements of the array are highlighted based on simulated classifiers varying in quality (Figures 2(b)–2(d)). Soft highlighting is achieved by modulating the grey level of pixels of an element. With the light-blue background, dark elements are salient and light elements are difficult to detect against the background.

We refer to our simulated classifier as a *stochastic oracle-based classifier* or SOC. The SOC generates values in $[0, 1]$ that represent the probability or confidence level that a given element is a target. The term “oracle” refers to the fact that we know the true class label of each element chosen for our displays. The term “stochastic” refers to the fact that an element’s confidence level is drawn from a probability distribution—either Beta(θ , 1) if the element is a target, or the complementary distribution Beta(1, θ) if the element is a nontarget. With $\theta = 1$, both distributions are uniform, and the level provides no information about the true class label. As $\theta \rightarrow \infty$, an element’s confidence approaches a deterministic level of 1.0 for targets and 0.0 for nontargets. The quality of the SOC is thus specified by θ . Because θ is not an intuitive measure, we use two other measures to describe the quality: the *equal error rate*, the error rate based on a decision threshold that matches TP and TN rates, and d' , a measure of discriminability based on signal-detection theory (Green and Swets 1966). Table 1 shows the three equivalent measures of quality of SOC’s used in Experiments 1–4. Hereafter, we use d' to refer to the quality, and SOC $_{d'}$ to refer to an SOC with a given d' . The d' value associated with Figures 2(b)–2(d) is presented in the

Figure caption. Note that even for values of d' traditionally considered large—say, greater than 2.0—not every target pops out and not every nontarget is concealed. The existence of a threshold that reliably discriminates targets from nontargets does not imply that targets and only targets will be perceptually salient.

In Experiment 1, our goal was to determine the classifier quality necessary to obtain a benefit of soft highlighting over no highlighting. If individuals are sensitive to graded highlights, then even relatively poor classifiers may boost performance.

2.1.1 Stimuli and Design. Displays were generated for six experimental conditions in which soft highlighting is based on $\text{SOC}_{d'}$ for $d' \in \{0.00, 0.75, 1.69, 3.07, 5.77, 9.24\}$ and for a control condition in which all elements are gray. Each display is a 10×10 array of handprinted digits consisting of 10 instances of each of the digits 0–9, randomly selected without replacement from the MNIST data set (LeCun et al. 1998). Each digit consists of 28×28 pixels whose values are thresholded such that each pixel is either off or on. The background color has RGB values on a 0–1 scale of (0.79, 1.0, 1.0). The foreground color of a digit with assigned classifier confidence $\rho \in [0, 1]$ is set to an RGB value converted from Lab space with lightness $1 - \rho$ and $a = b = 0$. Conversion from Lab to RGB for display passes through the XYZ color space; the XYZ space is based on the D55 reference illuminant. The use of the Lab-to-RGB lightness conversion is intended to match an element’s perceptual salience to the classifier confidence.¹

The experiment consisted of 42 trials in six blocks, with one trial from each of the seven conditions in each block. Trials within a block were randomized. There were no breaks between blocks.

2.1.2 Participants. In all experiments, participants were recruited via Amazon’s Mechanical Turk platform for pay, at a rate of \$8/h. Thirty five adults with normal color vision were accepted into the study and 25 completed all trials. The remainder either quit the experiment voluntarily or were rejected for failing to satisfy a requirement that they maintain focus on the browser window through the course of the experiment. We did not record the manner of termination.

2.1.3 Procedure. Each trial began by presenting the 10×10 array centered on the participant’s monitor. Participants were instructed to click on all targets (the digit 2) in the display. When a target was clicked, it was removed from the display and replaced with the background color. Clicking on a nontarget had no effect. If participants failed to find all 10 targets within 15s, then the trial terminated and they were informed how many targets they had found. At the end of each trial, a “next” button appeared below the array and the experiment paused until the participant clicked to proceed. The requirement to click on the button ensured that each trial began with the mouse in a known position relative to the array. If participants failed to find at least two targets or clicked on more than six nontargets, then they were rejected and the experiment terminated. These restrictions ensured that participants maintained focus on the experiment window until completion of the experiment.

2.1.4 Results. For each participant and each condition, we compute the percentage of the ten targets detected as a function of time from the onset of the stimulus array. Figure 3(a) shows the means across participants for the no-highlighting control condition (black) and the six experimental conditions (shaded from blue to purple in increasing order of classifier quality). A curve that rises quickly and then asymptotes at 100% is indicative of an easy search; a curve that rises slowly and doesn’t reach 100% by the end of the 15s trial is indicative of a difficult search. Error bars in the figure are corrected for between-participant variability according to the procedure in Masson and Loftus (2003). Nontarget clicks are infrequent, with a mean rate of 0.19 per trial across participants and conditions.

¹We intuited that the foreground color would become less salient against the background as $\rho \rightarrow 0$. However, a reviewer argued that increasing the luminance beyond the background level makes the display elements more, not less, salient, and suggested that it would have been better to use a contrast model in which ρ was encoded via a distance in Lab space, with a constraint on the minimum distance.

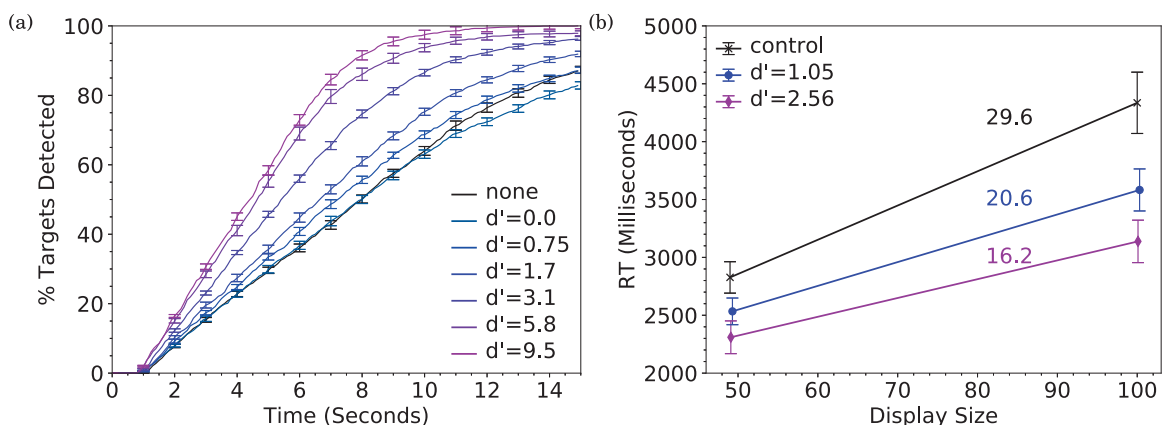


Fig. 3. (a) Experiment 1: Percentage of the ten targets detected as a function of time and condition. Error bars indicate ± 1 standard error of the mean corrected for between-participant variability according to the procedure in Masson and Loftus (2003). (b) Experiment 2: Mean target-detection latency for the two display sizes and three highlighting conditions.

Participants are sensitive to classifier quality: the soft highlighting curves do not cross over one another, indicating that the superiority of one classifier over another is consistent across time. All soft highlighting conditions with $SOC_{d'>0}$ lead to more efficient search than the no-highlighting condition. We are surprised by the finding that even $SOC_{0.75}$ is helpful; our informal polling of colleagues led to the prediction that only the highest quality classifiers would facilitate search. Participants are able to exploit the signal provided by a rather weak classifier. The pattern of results appears to rule out two alternative hypotheses: first, that relatively weak classifiers draw attention to the wrong elements and therefore suffer relative to the control; and second, that soft highlighting leads to an early advantage—due to the appropriately highlighted elements—but a late cost—due to the inappropriately highlighted elements—relative to the control.

To conduct a more quantitative analysis, we computed the median time for each participant to find *five* targets in each condition and performed paired t tests between conditions. Control, SOC_0 , and $SOC_{0.75}$ had mean times of 7.54, 7.67, and 7.38s, respectively, though none of these pairwise differences were reliably distinguishable. Using median time to find *eight* targets, $SOC_{0.75}$ is faster than control (10.67 versus 11.41s, $t(24) = 3.36$, $p = 0.0025$) and faster than SOC_0 (10.67 versus 11.52s, $t(24) = 2.56$, $p = 0.017$). Focusing instead on the number of targets detected halfway through the trial, at 7.5s, $SOC_{0.75}$ obtains a greater proportion of detections than control (0.54 versus 0.49s, $t(24) = 3.41$, $p = 0.002$) but not SOC_0 (0.54 versus 0.51s, $t(24) = 1.40$, $p = 0.17$). We have given ourselves license to fish through the data, and one should therefore interpret the findings cautiously, there does seem to be a preponderance of evidence that $SOC_{0.75}$ highlights are sufficient to facilitate search.

Because the experimental conditions are intermixed, participants cannot anticipate the trustworthiness of the classifier on any trial. However, due to our simulation of classifier outputs, an easily discernible statistic—the spread or variance of the confidence levels—is a reasonable proxy for classifier quality: the more uniform the spread, the less reliable the classifier. Thus, we do not expect results would change much with a blocked design.

2.2 Experiment 2: Search for a Single Target in Variable-Sized Displays

Traditional studies of visual search utilize displays with zero or one targets and a variable number of distractors [e.g., Treisman and Gelade (1980)]. Experiment 2 follows this paradigm with single-target displays of two sizes, 7×7 or 10×10 . This design allows for the measurement of a search *rate*—the increase in response time per additional display element. The speed up from soft highlighting observed in Experiment 1 could be due either to an increased search rate or to factors that are unrelated to display size, such as early perceptual processing,

initial segmentation of the display, or motor preparation. These two possibilities can be distinguished via an affine model relating response latency, t , to display size, s : $t = \alpha + \beta s$, where β is the search rate, and α is display-size independent processing.

2.2.1 Stimuli and Design. Each stimulus array consists of a single target—the digit 2—placed in a random cell, with the remaining cells filled at random with nontarget digits. No digit token appeared more than once in the course of the experiment. The 7×7 arrays occupied the same area as a 7×7 portion of the 10×10 arrays and were centered on the screen. The experiment consisted of 15 blocks of 6 trials, for a total of 90 trials. Each block had exactly one trial that crossed display size— 7×7 or 10×10 —with highlighting condition—control, SOC_{1.05}, and SOC_{2.56}. As in Experiment 1, order of trials within a block was randomized, and no break or indication was given at the end of a block.

2.2.2 Participants. Of 56 participants who enrolled, 37 completed all trials; the remainder were either rejected or terminated their participation voluntarily. All had normal color vision.

2.2.3 Procedure. Participants were instructed to locate and click on the target. Clicks on nontargets were ignored. To ensure adherence to instructions, participants were rejected if they failed to find the target within 45s or if they clicked on more than six nontargets in a trial. At the end of each trial, response latency was displayed and a “next” button lit up that would initiate the following trial when the participant was ready.

2.3 Results

Figure 3(b) shows the mean latency for target detection for each of the three highlighting conditions—control, SOC_{1.05}, and SOC_{2.56}—and the two display sizes— 7×7 and 10×10 . An ANOVA indicates a main effect of highlighting condition ($F(2, 72) = 14.4, p < 0.001$) and display size ($F(1, 36) = 127.6, p < 0.001$). Critically, the interaction between highlighting and display size is reliable ($F(2, 72) = 5.18, p = 0.008$). The search slope in the control condition appears steeper than in the SOC_{1.05} condition, which in turn is steeper than in the SOC_{2.56} condition. Thus, highlighting reduces the time to search for each item in the display, consistent with the notion that highlighting guides attention to relevant locations.

The slope of the best fit line to each condition is shown in Figure 3(b). Performing pairwise comparisons, both the nonparametric two-sided Wilcoxon signed-rank test and paired t-test show a marginal difference between the slopes for control versus SOC_{1.05} ($w(36) = 1550, p = 0.079; t(36) = 1.89, p = 0.067$) and a reliable difference for control versus SOC_{2.56} ($w(36) = 1652, p = 0.004; t(36) = 3.18, p = 0.003$). No difference was observed between SOC_{1.05} and SOC_{2.56} ($w(36) = 1492, p = 0.24; t(36) = 1.18, p = 0.26$). Our findings are thus consistent with the claim that highlighting increases the efficiency of search.

2.4 Experiment 3: Comparing Soft and Hard Highlighting

Having established that individuals are able to exploit the signal in soft highlights to guide attention, we turn to the key issue of how soft highlighting fares relative to hard highlighting. Three classifiers were studied: SOC_{0.75}, SOC_{1.00}, and SOC_{1.25}. We selected this set of relatively weak classifiers to further investigate the relationship between classifier quality and human performance. Soft highlighting was based on the SOC output, as in earlier experiments. Hard highlighting was based on thresholding the SOC output at 0.5, yielding displays such as that shown in the right panel of Figure 4.

2.4.1 Stimuli and Design. Experiment 3 used digit arrays identical to those used in Experiment 1, with ten targets per array. In addition to a control condition with no highlighting, six classifier conditions were formed by the Cartesian product of highlighting type—soft, hard—and classifier quality—SOC_{0.75}, SOC_{1.00}, and SOC_{1.25}. The experiment consisted of six blocks, each with one instance of the seven conditions, for a total of 42 trials. For hard highlighting, classifier confidence is thresholded at 0.5. Elements above and below this threshold are




Fig. 4. Experiment 3, sample stimulus arrays: (left) soft highlighting via $SOC_{1.05}$; (right) hard highlighting obtained by thresholding $SOC_{1.05}$ confidence levels.

set to a display intensity equivalent to a confidence of 1.0 and 0.2, respectively. The level of 0.2 was selected to ensure that the digits were readily perceptible but were less salient than the highlighted digits. We conducted pilot studies in which we performed hard highlighting by circling elements or changing the background color of the highlighted element, but decided to match the appearance of displays with soft highlighting.

2.4.2 Participants. Of 54 participants who enrolled in the experiment, 41 completed all trials. The remaining 13 either quit the experiment on their own or were rejected. We did not record the reason for their failure to complete.

2.4.3 Procedure. Experiment 3 was run in a manner identical to Experiment 1 with identical instructions.

2.4.4 Results. For each participant and condition, we compute the number of targets detected as a function of time within the 15s trial. The curves have the same general shape as the data from Experiment 1 (Figure 3(a)), but to accentuate differences between conditions, Figure 5(a) plots the difference in the number of targets detected between each highlighting condition and the control condition. The control is depicted as a black dashed line at the baseline of zero. Values above zero indicate that more targets were located by that instant than in the control condition. Qualitatively, hard highlighting appears to have a small benefit relative to the control for $SOC_{1.25}$ between 3 and 9s, though the benefit for $SOC_{1.00}$ and $SOC_{0.75}$ appears minimal. In contrast, soft highlighting appears to have a larger advantage relative to the control for all three classifier strengths and for most of the duration of the trial.

Figure 5(b) shows the difference between the mean number of detected targets in corresponding soft versus hard highlighting conditions, as a function of time. Values of the curve greater than zero indicate that at a particular instant of time, participants found more targets with soft highlighting than with hard. An advantage is observed for soft highlighting for all three SOCs.

To perform a quantitative analysis, we divided the first 12s of each trial into four bins of 3s. We computed the number of targets found within each bin and performed a three way ANOVA with participant as the random factor and classifier quality, highlighting type, and time window (0–3, 3–6, 6–9, and 9–12s) as three within-participant independent variables. We observe a main effect of highlighting condition ($F(1, 40) = 10.213$, $p = 0.003$), with soft highlighting superior to hard highlighting. We also observe a main effect of time ($F(3, 120) = 933$,




Fig. 5. Experiment 3. (a) The mean number of targets located across all participants for each condition minus the mean number of targets located across all participants for the control condition of no highlighting as a function of time. Values above zero mean that more targets were detected relative to the control condition while values less than zero mean fewer targets were detected. (b) The difference between the mean number of targets detected by that time for the soft and hard highlighting conditions for corresponding SOCs. Positive values reflect better performance with soft than hard highlighting. Error bars are calculated using a between-participant variability correction (Masson and Loftus 2003).

$p < 0.001$), which simply reflects the fact that more targets are found around 6 or 9s than around 3 or 12s. No main effect of classifier quality is observed ($F(2, 80) < 1$), not terribly surprising given the small range of d' tested. No interactions involving these three factors are significant at the 0.05 level.

The effect size of soft versus hard highlighting was assessed at 6 and 9s for each of the three classifier qualities. The effect sizes range from small to medium: Cohen’s d at 6s is 0.53, 0.24, and 0.25 for $SOC_{0.75}$, $SOC_{1.00}$, and $SOC_{1.25}$, respectively; and at 9s 0.22, 0.32, and 0.27. For a given SOC, our results indicate that soft highlighting, which leverages the classifier’s graded output, supports human visual search better than hard highlighting, which thresholds the classifier output.

2.5 Experiment 4: Search for a Variable Number of Targets

Experiments 1–3 show that even a weak classifier can boost human performance when soft highlighting is used. However, because targets are present on every trial and the number of targets is known, participants continue searching until all targets are found or time is exhausted. Consequently, we have no evidence concerning the effect of highlighting on an individual’s decision to quit searching, and therefore, on the possibility of missed targets. As discussed in the introduction, the medical literature suggests that hard highlighting can increase the rate of target misses.

To explore the effect of soft highlighting on target misses, participants in Experiment 4 searched 10×10 arrays that contained between zero and two targets. They were instructed to click on each target and to press a “done” button to terminate a trial when they were confident that no targets remained. Nontarget clicks were ignored.

2.5.1 Stimuli and Design. Each display contained 0, 1, or 2 targets, the digit “2.” The remainder of the array was filled with randomly chosen digits. Five highlighting conditions were included in the design: a control with no highlighting, SOC_0 , $SOC_{1.69}$, $SOC_{3.07}$, and $SOC_{5.77}$. The experiment consisted of 80 trials, the first five of which were practice and were not used in analyses. The practice trials included one with no targets present and two each with one and two targets present and appeared in random order. The 75 subsequent trials were presented in five blocks of 15. The trials within a block were formed through the Cartesian product of five highlighting types and from zero to two targets. The sequence of trials was such that across the experiment there was exactly

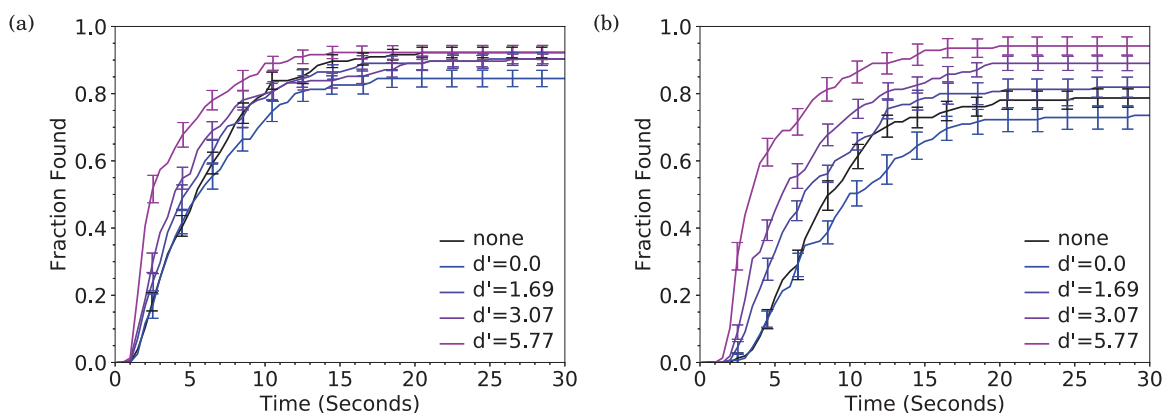


Fig. 6. Experiment 4. Mean number of targets detected by highlighting conditions for (a) one-target displays and (b) two-target displays. Error bars are calculated using a between-participant variability correction (Masson and Loftus 2003).

one trial per highlighting type C with T targets following a previous trial of type P . This led to five levels of the current type times five levels of the previous type times three levels of number of targets to arrive at 75 trials total.

2.5.2 Participants. Of 50 participants who enrolled in the experiment, 31 completed all trials and 19 quit the experiment on their own as no participants were rejected based on performance.

2.5.3 Procedure. Each trial continued until the participant clicked “done.” There was no upper limit on the time for a trial. Because this task was more difficult than in the previous experiments, we provided feedback after each trial to ensure participants remained vigilant and motivated. This feedback consisted of a three-quarter sized version of the digit array with the actual targets colored green if they had been detected or red if they had been missed. A buzzer sounded if any target was missed.

2.5.4 Results. The analysis of this experiment is more complex than the analyses of the previous experiments due to the variable number of targets and the possibility of a speed-accuracy trade off.

We first examine zero-target displays and ask how highlighting influences the time to determine that no target was present and to terminate the trial. We compute the median time to terminate for each participant and condition. Figure 7(a) shows the mean of this statistic across participants. An ANOVA reveals no main effect of condition ($F(4, 120) = 0.89$), indicating that highlighting did not influence the time to terminate a trial.

For one- and two-target displays, detection accuracy asymptotes at 85–93% and 75–95%, respectively (Figures 6(a) and 6(b)). Differences among conditions do not attain significance for one-target displays ($F(4, 120) = 2.10$, $p = 0.085$) but are reliable for two-target displays ($F(4, 120) = 6.18$, $p < 0.001$). Both figures show the same trend that higher classifier qualities yield faster detection and higher asymptotic performance. Random highlights, as embodied by SOC_0 , appear to impose a cost relative to no highlighting.

No hint of a speed-accuracy trade off is observed for target-present displays. Figures 7(b) and 7(c) indicate that time to detect all targets in one- and two-target displays decreases with increasing classifier quality (main effect of condition on one target displays: $F(4, 120) = 8.68$, $p < 0.001$; two target: $F(4, 120) = 11.69$, $p < 0.001$).

Across displays with varying numbers of targets, soft highlighting appears to be an unqualified win. No speed-accuracy trade offs are induced: participants miss fewer targets and are faster to complete the task, even with uncertainty in the presence and number of targets. Experiment 4 involves a more realistic search task than Experiments 1–3 because of the target uncertainty and because earlier experiments offered participants no means of self-terminating a trial without finding all targets. Experiment 4 strengthens the results from Experiments 1–3




Fig. 7. Experiment 4. Mean latency to terminate trial after correctly locating all targets in (a) zero-target displays, (b) one-target displays, and (c) two-target displays. Error bars are calculated using a between-participant variability correction (Masson and Loftus 2003).

by showing that, if anything, soft highlighting leads to earlier termination of search with fewer misses. However, the conclusions from Experiments 1–4 pertain to tasks in which target-distractor discrimination is easy.

2.6 Conclusions from Experiments with Synthetic Displays

We described four experiments with synthetic displays and simulated classifiers to explore the influence of soft highlighting on visual search. Although the search task we studied is simple and artificial, it provides a convenient vehicle to investigate the influence of classifier quality on performance and the value of soft versus hard highlighting. We found that even relatively weak classifiers boost human performance when the classifier output is incorporated into search displays with soft highlighting but not with hard highlighting. We observed across-the-board advantages for soft highlighting over a no-highlight control. Individuals are sensitive to subtle signals embedded in noisy classifier output and can leverage these signals to search more accurately and efficiently.

The displays that we studied, typical of synthetic images in psychological tasks, are composed of spatially isolated, segmented elements. Based on experiments using these displays, one would be reluctant to extrapolate findings to more complex, naturalistic images and real-world search tasks. In tasks such as medical and satellite imagery analysis, images are continuous, display features are not neatly segmented but rather are continuous and overlapping and varying in size, and target-distractor discrimination is challenging even for domain experts.

We therefore conducted a series of experiments involving target search using naturalistic images: satellite photos of urban and suburban landscapes. Although participants in our experiments were not experts in analyzing these images, most individuals have some familiarity with them. To ensure the task was as true-to-application as possible, we based highlights on the output of a state-of-the-art neural net classifier. In contrast to our SOCs, which generate outputs that depend only on the target class of a display element, a neural net's output and errors are systematically related to the image data. If the neural net and humans fail to detect the same targets, then highlighting is unlikely to boost human performance.

Before turning to the experiments, we first describe the domain, the classifier we constructed for the domain, and the technique we use for highlighting satellite imagery.

3 ANALYSIS OF SATELLITE IMAGERY

The analysis of satellite imagery is a critical, time-intensive task performed by government intelligence, military, forestry, and disaster-relief experts, as well as biologists and social scientists interested in characterizing human activity and agriculture. This analysis often requires searching for targets, for example, certain types of vegetation, sporting fields, missile silos.

We devised a search task that leverages a publicly available point-of-interest data base containing coordinates of fast-food restaurants operated by a popular chain (McDonald's) within the United States. Using an API from




Fig. 8. (a) 100×100 pixel patches of satellite images (a) not containing the target, and (b) containing the target.

Mapquest, we pulled sets of satellite images both with and without restaurants at the highest resolution available. We converted the pansharpened (color) images to greyscale to follow the workflow of a typical image analyst who looks almost exclusively at panchromatic imagery (like a black and white photograph). We reviewed images to ensure that a restaurant was present at the specified coordinates and we rejected restaurants that did not follow the classical design. Figure 8 presents examples of 100×100 pixel patches centered on targets and other patches not containing targets. These patches are large enough to contain the entire restaurant along with a portion of the surrounding drive-through and parking lot. Variability in the layout of the property makes identifying the restaurant challenging.

Stimuli used in our experiments were 1000×1000 images. Each contained no more than one target, and the target was randomly offset such that it could appear anywhere except in the outer 50 pixel wide border around the image. We obtained 2,992 target-present images and ten times as many target-absent images, the latter being selected randomly from the 0.15-mile vicinity around the restaurant and constrained such that no portion of a target was present in the image. This proximity-based selection ensures that the target-absent images have roughly the same density and pattern of trees, roads, and buildings as the target-present images. We further verified that target-absent images had a similar character via manual inspection.

3.1 Constructing a Classifier for Satellite Imagery

We constructed a state-of-the-art classifier for the restaurant search task. The classifier takes as input a 100×100 pixel patch of an image and outputs the probability that a target is at least partially contained in the patch. The classifier is implemented as a convolutional neural net. Convolutional nets exploit the structure of images via spatiotopically organized layers of processing units that detect features in local spatial regions of the layer below. Arranging these feature detectors hierarchically achieves a transformation from a location-specific, representation of primitive features (pixels) to a location-independent representation of object identities (Mozer 1991). We used Bayesian optimization methods (Snoek et al. 2012) to select the specific architecture, which had two convolutional layers and two fully connected layers. We augmented the data set by horizontal and vertical reflections and 90° rotations, maintaining the 1:10 target:nontarget ratio. Cross validation was used to settle on further aspects of the model and training procedure. Details of the models, data, and training procedure are described in Kneusel (2016).

The selected classifier was evaluated on a separate test set that had not been used during the training process. When thresholded at probability 0.5, the classifier achieves a TP rate of 0.886 and a TN rate of 0.995. When




Fig. 9. (a) satellite image; (b) pixel-by-pixel classifier outputs; (c) smoothed classifier outputs; (d) incorporating classifier outputs into image via soft highlighting.

thresholded to achieve equal TP and TN rates, the resulting *equal error rate* (*EER*) is 0.030. The classifier achieves an *area under the curve* (*AUC*) of 0.992. (*AUC* is an alternative measure of discriminability to d' and ranges from 0.5 for a classifier that performs no better than chance to 1.0 for a classifier that perfectly discriminates positive and negative examples; see Green and Swets (1966) for further details.) Although these statistics seem to suggest that the classifier is extremely strong, one must take into account that a target can appear anywhere in a high-resolution image, affording many opportunities for false positives.

3.2 Highlighting Satellite Imagery

Given an image like that in Figure 9(a), the classifier is presented with each 100×100 pixel image patch and the target probability obtained from the classifier is associated with the middle pixel of the patch—the pixel at (50, 50). Figure 9(b) shows the probability by pixel. (Only complete patches are classified; thus, there is a roughly 50-pixel border around the image in which targets will never be detected.) Because a target has a spatial extent that is roughly the size of a 100×100 patch, the classifier outputs are smoothed using a Gaussian filter with $\sigma = 50$ pixels, resulting in an activation map such as that in Figure 9(c).

The activation map is used to determine highlights. For soft highlights, we express each pixel in the original image using hue-lightness-saturation coordinates, and transform the pixel by setting its saturation level to the value of the corresponding pixel in the activation map (0 is minimum saturation, 1 is maximum). The lightness value is obtained from the original image, and the hue in all our experiments is fixed to be red.² This manner of highlighting preserves luminance information in the original greyscale image while enhancing the saliency of highlighted locations. For hard highlights, we developed a heuristic scheme that: (1) finds the image location of the largest value in the activation map, (2) draws in the original image a 100×100 pixel red square centered on the location, (3) resets the activation map in a 150×150 pixel region around the identified location, (4) repeats the procedure until the largest activation map value is less than a given threshold, θ . This greedy procedure ensures that hard highlights will not overlap. We selected $\theta = 0.5$, which produced between 2 and 15 highlights per image, with a mean and median of 6. See the right panel of Figure 1 for an illustration of hard highlighting.

4 EXPERIMENTS WITH SOFT HIGHLIGHTING IN SATELLITE IMAGERY

4.1 Experiment 5: Target Localization with Feedback

In Experiment 5, participants searched images that each contained exactly one target and were instructed to click where they believed the target was located. Following an incorrect click, a buzzer sounded and the trial continued until the target was correctly located. Three highlighting conditions were studied: soft, hard, and a control with no highlighting.

²Red was chosen as the highlight color after visual inspection by the authors of a range of alternatives.

4.1.1 Stimuli and Design. Stimuli were drawn from a set of 245 target-present images, each 1000×1000 pixels. The target placement was random, with the constraint that no portion of the target appear in the outer 50-pixel perimeter of the image. The images were drawn from a test set not used for training our classifier; thus, the highlights provided by the classifier reflect the quality of highlighting one would anticipate in actual usage. After the image was classified, the outer 50-pixel border was stripped off and the image was rescaled by a factor of $2/3$ to 600×600 pixels to ensure that it would fit entirely within a participant's browser window.

Participants viewed two blocks of 17 trials, with the first two trials in a block considered to be practice to familiarize them with the task and displays. All stimuli in a block were highlighted in the same manner (soft, hard, or control). Counterbalancing was performed across a set of six participants, with participants all seeing the same sequence of images but with different pairings of highlighting conditions in the two blocks. (The six pairings are: soft-hard, hard-soft, soft-control, control-soft, hard-control, and control-hard). This counterbalancing ensured that each image occurred at the same position in a block in each condition, and that each highlighting condition occurred equally often in blocks 1 and 2.

4.1.2 Participants. Participants were restricted to being from North America to ensure familiarity with the typical setting and layout of fast food restaurants. Of 289 participants who enrolled in the experiment, 84 (14 groups of 6) completed it successfully. (Fourteen groups were chosen to use each image roughly twice across the pool.) Of the 205 participants who did not complete, 100 of them quit voluntarily, and 105 were rejected for defocusing their browser window. They were instructed not to defocus during a trial, because multitasking corrupts our measure of latency.

4.1.3 Procedure. Prior to the start of the experiment, all images were preloaded on the participant's machine to ensure smooth pacing of the experiment. In the initial instructions, participants were told the name of the restaurant chain (McDonald's) and were shown 20 small example images of the restaurant (Figure 8) to familiarize themselves with the task. They were asked not to multitask as Turk workers are apt to do. At the start of each block, they were shown a representative satellite image for the highlighting condition, and they were given instructions specific to the highlighting condition. For soft highlighting, they were told, "You will now see a series of 17 black and white images that have been shaded where the computer believes a restaurant may be"; and for hard highlighting, "You will now see a series of 17 black and white images that have red boxes where the computer believes a restaurant may be. The actual restaurant may be outside any of the boxes." In both conditions, they were instructed, "The computer is not perfect but can provide assistance in locating the restaurant."

Participants continued to click on the image until the target was detected. We considered a click to be on the target if it landed within 30 pixels of recorded target location in both the horizontal and vertical directions (which is within 45 pixels of the unscaled image—roughly the full extent of the restaurant). After each trial, a "next" button appeared to continue to the following trial.

4.1.4 Results. Figure 10(a) plots the fraction of targets found by time within a trial for each condition. Better performance is indicated by a steeper curve and/or a higher asymptote. Soft highlighting appears to outperform hard, which in turn outperforms the control condition. This impression is supported by within-participant comparisons for pairs of conditions using the median target-detection latency for each participant and each condition. The means across participants are depicted in Figure 10(b). Starting with the group on the left, tested with soft highlighting and the control condition, soft is faster than the control (9.7 versus 5.4s), significant by a paired t -test and by the nonparametric two-sided Wilcoxon signed-rank test ($t(27) = 5.40$, $p < 0.001$; $w(27) = 1132$, $p < 0.001$). Soft highlighting beats hard highlighting as well (5.2 versus 7.4 sec; $t(27) = 2.73$, $p = 0.011$; $w(27) = 539$, $p < 0.001$). These analyses excluded the first two trials of a block as practice; including the two trials does not affect the outcome of the statistical tests.

The mean number of nontarget clicks per trial is 7.6, 6.1, and 3.6 in the control, hard, and soft conditions, respectively. Unpaired t -tests indicate that soft highlighting leads to significantly fewer clicks than either hard

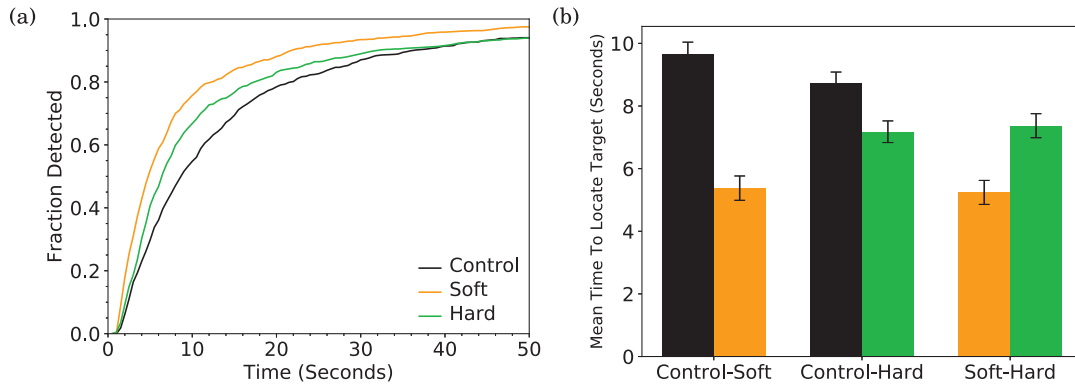


Fig. 10. Experiment 5: (a) Fraction of targets detected by time and condition; (b) within-participant comparison of mean search time for pairs of highlighting conditions.

($t(110) = 2.75$, $p = 0.007$, Cohen's $d = 0.53$) or control ($t(110) = 3.88$, $p < 0.001$, Cohen's $d = 0.77$), though the difference between hard and control is not reliable ($t(110) = 1.19$, $p = 0.24$, Cohen's $d = 0.23$).

We examined the relationship between an individual's target-detection latency and the classifier output probability at the target location. We found the strongest correlation with soft ($\rho_{\text{soft}} = -0.43$), then hard ($\rho_{\text{hard}} = -0.25$), then the control ($\rho_{\text{control}} = -0.09$), suggesting that soft highlighting is effective in communicating classifier confidence.

4.2 Experiment 6: Target Detection with Feedback

In Experiment 5, trials terminated only when a target was correctly detected; therefore, false-negative trials could not occur. This design avoided a key potential pitfall of highlighting: an increase in the probability of missing a target that is not highlighted (see Section 1.1). Experiment 6 was therefore designed to allow for the possibility of false-negative responses. Both target present and target absent trials were included, and participants were allowed to terminate a trial without finding a target.

4.2.1 Stimuli and Design. Target-present images were re-used from Experiment 5. Target-absent images were selected as explained in Section 3. The experiment consisted of 36 trials arranged in 4 blocks of 9. Each block included three images for each highlighting condition (soft, hard, control). Targets were present in two of the three images. The order of the nine trials within a block was randomized. We counterbalanced the assignment of images to conditions using a Latin square design and testing participants in groups of three. For each set of three, a unique set of images was used.

4.2.2 Participants. Of the 155 participants who enrolled in the experiment, 90 completed successfully (30 groups of 3). Of the remaining 65, 43 were rejected for changing window focus and the remainder quit voluntarily. Turk workers were excluded if outside the U.S. or if they had done previous studies.

4.2.3 Procedure. Because highlighting conditions were intermixed, initial instructions explained both soft and hard highlighting. Each trial proceeded as in Experiment 5, except that the screen included a "No restaurant present" button, which allowed participants to terminate a trial without finding the target. As in Experiment 5, when a location was clicked, the trial terminated if the location contained a target; otherwise, a buzzer sounded and the trial continued. Following each trial, a "next" button appeared to continue, and between blocks the instructions were re-displayed to break up the monotony of the trial sequence.

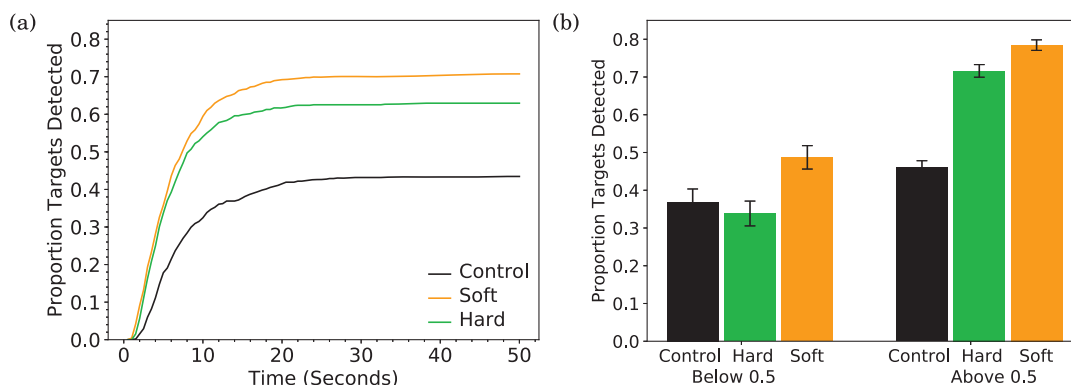


Fig. 11. Experiment 6, target-present trials: (a) Fraction of targets detected by time and condition; (b) target-detection probability by classifier output and condition. Error bars are calculated using a between-participant variability correction (Masson and Loftus 2003).

4.2.4 Results. Figure 11(a) shows the fraction of targets detected by highlighting condition as a function of time from stimulus onset. The curves in this figure do not asymptote at 1.0, because participants may give up and miss a target. The asymptotic detection rate for soft, hard, and control are 70.7%, 62.8%, and 43.9%, respectively. All pairwise differences are highly reliable by a nonparametric two-sided Wilcoxon signed-rank test (soft-control, $w(89) = 10618$, $p < 0.001$; hard-control, $w(89) = 9982$, $p < 0.001$; soft-hard, $w(89) = 7153$, $p = 0.004$), as well as by paired t tests that make stronger assumptions about the data distributions.

On target-absent trials, mean time to terminate is longer for soft (11.3s) than for hard (9.6s) or control (9.4s). Performing a t -test on the log-transformed times, both the soft-hard ($t(89) = 2.94$, $p = 0.004$) and soft-control ($t(89) = 3.47$, $p < 0.001$) differences are reliable. It’s not entirely clear how to weigh this additional time spent on target-absent trials against the improvement in both speed and accuracy on target present trials (Figure 11(a)).

A critical question this experiment was designed to address is whether non-highlighted targets are missed more often when highlighting is present. Figure 11(b) suggests that soft highlighting may potentially be a solution to the “no free lunch” dilemma. We have taken all target-present trials and broken them up by whether the classifier obtained a confidence level above or below 0.5—the threshold we use for hard highlighting.³ For high-confidence classifier outputs, soft highlighting produces more detections than hard ($t(178) = 3.12$, $p = 0.002$), and both soft and hard highlighting beat the control condition ($t(178) = 13.99$, $p < 0.0001$, $t(178) = 10.26$, $p < 0.0001$). For low-confidence classifier outputs, soft highlighting beats both hard and the control ($t(156) = 3.06$, $p = 0.003$, $t(156) = 2.38$, $p = 0.018$). We do not observe a cost for hard highlighting relative to the control ($t(156) = 0.61$, $p = 0.54$), which we expected to see based on the existing literature suggesting that hard highlighting can mask non-highlighted targets. Nonetheless, the advantage of soft highlighting over hard in Figure 11(a) seems to hold regardless of the classifier output.

4.3 Experiment 7: Target Detection Without Feedback

In Experiment 6, false-positive clicks received immediate feedback. In our final experiment, we removed all feedback during a trial, allowing participants to select one location if they believed a target was present or to indicate that no target was present. Feedback was provided after a response was submitted. (Without any feedback, we were concerned that participants would fail to learn the task and might also lose motivation.) With

³We did not control for classifier output level in each participant’s trial sequence. Consequently, some participants had no trials in certain conditions with below 0.5 confidence levels. Our classifier has high confidence for targets on 76% of all trials.




Fig. 12. Experiment 7: (a) Proportion of target-present trials in which target is detected (TP rate); (b) proportion of target-absent trials in which some location is misidentified as a target (image-wise FP rate); (c) proportion of targets detected when classifier response to target was below or above 0.5.

the possibility of both false-positive (FP) and false-negative (FN) responses, this experiment allows us to analyze the effect of highlighting on target *discriminability*, using measures such as d' and AUC.

4.3.1 Stimuli and Design. The stimuli and design were identical to Experiment 6. Participants were presented with 36 images in four blocks of 9 trials.

4.3.2 Participants. Of the 130 participants who enrolled in Experiment 8, 90 completed successfully. Of the 40 who did not complete, 16 left voluntarily and 24 were rejected for de-focusing their browser windows during a trial. Turk workers were excluded if outside the US or if they had done previous studies.

4.3.3 Procedure. Once a trial began, participants were allowed to click anywhere in the image to indicate a target. Clicking displayed a cross hair (plus sign) at the clicked location. If another location was clicked, then the cross hair moved. Thus, at most one cross hair was present on the screen. No feedback was provided as the participant clicked. Below the image were two buttons. Clicking the button labeled “no restaurant present” removed a cross hair from the display if one was present. Clicking the button labeled “submit response” terminated the trial and recorded the location of the cross hair if one was present. After each trial, participants received feedback indicating whether their response was correct or not. A reduced-scale version of the image was shown with the location of the actual target, if any. Additionally, a pleasant ding or an annoying buzzing sound was played depending on response correctness. Feedback remained on the screen until participants clicked a “next” button to continue the experiment.

4.3.4 Results. Using target-present trials, we compute the proportion of targets detected (TP rate) by condition (Figure 12(a)). Fewer targets are detected with the no-highlight control than hard highlighting (0.18 versus 0.29, $t(89) = 5.64$, $p < 0.0001$) or soft highlighting (0.18 versus 0.34, $t(89) = 6.47$, $p < 0.0001$). Fewer targets are detected with hard highlighting than soft (0.29 versus 0.34, $t(89) = 2.10$, $p = 0.039$). As in Experiment 6 (Figures 11(a) and 11(b)), lower asymptotic detection rates in Experiment 7 go hand in hand with slower detection (not shown), and the benefit of soft highlighting is observed regardless of classifier confidence (Figure 12(c)).

Using target-absent trials, we compute the proportion of trials that *some* location is selected as the target, which we refer to as the *image-wise FP rate* (Figure 12(b)). Consistent with the literature discussed in Section 1.1, highlighting—both soft and hard—leads to more locations erroneously identified as targets (soft versus control, $t(89) = 3.68$, $p < 0.001$; hard versus control $t(89) = 2.87$, $p = 0.005$); however, the difference between soft and hard is not significant ($t(89) = 0.95$, $p = 0.34$).

Although it is encouraging that the TP rate is higher for soft than hard highlighting and that the FP rates do not differ significantly, one would like to compare conditions using a measure of discriminability like d' .

Table 2. Human and Classifier Discriminability Scores First Column Obtained by Pooling Human Data Across Participants, and Using $P = 100$ and $Q = 0$. Second Column Obtained by Computing d' for Each Participant and Then Averaging, Using $P = 100$ and $Q = 1$. Standard Deviations of the Mean are Shown in Parentheses, Corrected for Inter-Participant Variability (Masson and Loftus 2003)

		pooled d'	individual d' mean	individual d' SD
human	control condition	1.90	2.07	(0.38)
human	hard highlighting	2.49	2.45	(0.37)
human	soft highlighting	2.70	2.60	(0.47)
classifier	matched FP rate	2.19	2.20	
classifier	matched TP rate	1.97	2.08	

However, to do so, we must address an issue concerning the computation of FP rates. A false positive occurs when a participant mistakenly labels a patch as a target. Each target-absent image contains many patches, and if *any* of these patches is seen as a target, it will contribute to the image-wise FP rate. Assuming that the scoring of one patch is independent of another, we can recover the patch-wise FP rate from the image-wise rate:

$$FP_{\text{patch}} = 1 - (1 - FP_{\text{image}})^{\frac{1}{P}},$$

where P is the effective number of patches. We use FP_{patch} for estimating d' . Although the choice of P has a scaling effect on d' , P does not influence any comparisons we make across conditions (i.e., relative values of d'). We used $P = 100$ to reflect the fact that targets lie in a roughly 100×100 pixel region of our 1000×1000 pixel images, yielding 10×10 nonoverlapping patches.

Table 2 (middle column) shows per-condition d' computed by pooling data across participants to obtain patch-wise TP and FP rates for the population. Soft highlighting outperforms hard, and both outperform the control. It is informative to compute d' for the classifier alone. To do so, we must pick a classifier-output threshold for judging a target as being present. We picked two different thresholds to match either the FP or the TP rate of our human population. Although these two thresholds yield slightly different d' values, both are only slightly better than humans in the control condition, and both are worse than humans in the highlighting conditions. Thus, the hybrid human-classifier judgments, obtained via soft or hard highlighting, are superior either to human judgments alone (the control condition) or classifier judgments alone.

These population-wide d' values say nothing about how reliably *individuals* discriminate targets from non-targets across conditions. To do so, we must calculate d' for each individual, but estimates are quite noisy due to the small number of trials in each condition: per participant and condition, there were only four target-absent trials to compute a FP rate, and only eight target-present trials to compute a TP rate. Because the TP and FP rates could easily be 0.0 or 1.0, yielding ill-defined d' values, we performed smoothing based on the Bayesian notion of imaginary-count priors. Essentially, we assume a constant Q that regularizes the computed rates:

$$TP = \frac{C + Q}{8 + 2Q} \text{ and } FP = \frac{E + Q}{4 + 2Q},$$

where C is the number of targets correctly detected in target-present images, and E is the number of trials in which a target is erroneously detected in target-absent images. Q represents the strength of a prior belief that the TP and FP rates are 0.5, with a larger Q smoothing observations away from the extreme rates of 0.0 and 1.0. As $Q \rightarrow \infty$, the observations matter less and all values approach 0.5 in the limit. Operationally, smoothing assumes Q additional trials in which a target is detected and Q additional trials in which no target is detected.

We arbitrarily picked $Q = 1$, which limits TP rates to lie in $[0.1, 0.9]$ and FP rates to lie in $[0.17, 0.83]$. Table 2 (right column) presents d' values for the three experimental conditions and the classifier, computed with Q smoothing. The table also shows inter-participant variability. Soft and hard highlighting both lead to larger d' than the control (soft: $t(89) = 6.13$, $p < 0.001$; hard: $t(89) = 6.65$, $p < 0.001$), and soft highlighting is marginally

superior to hard ($t(89) = 1.93$, $p = 0.057$). Effect sizes, as measured by Cohen's d are 1.27 (control vs. soft), 1.03 (control vs. hard), and 0.36 (soft vs. hard).

Although the contrasts between the two highlighting conditions and the control are insensitive to Q , the critical contrast between soft and hard is. As Q increases, the effect size becomes larger and the reliability of the t test increases. Rather than picking Q arbitrarily, we used statistical tests of non-normality to determine the smallest Q that yields a data distribution that did not deviate systematically from a Gaussian. Using the Jarque-Bera or D'Agostino test of normality to pick Q , we obtain still-marginal soft versus hard contrast ($t(89) = 2.03$, $p = 0.045$).

4.4 Conclusions from Experiments with Satellite Imagery

Soft highlighting proved robustly superior to hard highlighting on satellite imagery search in Experiments 5 and 6, which prevented false-positive responses via intra-trial feedback. In Experiment 7, when no intra-trial feedback was provided, false positives did occur, but when true and false positives are combined to obtain a d' discriminability measure, soft highlighting is marginally better than hard highlighting. In retrospect, our Experiment 7 may not have had adequate power due to the small number of trials per participant and condition, and we expect that a larger replication would obtain a more robust contrast.

Across Experiments 5–7, highlighting conditions are superior to the control condition with no highlighting. Given that our participants were relative novices at the particular search task, it seems sensible that highlights would facilitate performance. It remains an interesting open question how much domain experts would benefit from highlights, and the relative benefit that soft and hard highlights would provide. Because the same pattern of results was observed in Experiments 1–4 where participants were domain experts, we might argue that expertise does not eliminate the benefit of highlighting, soft or hard, and given the literature cited earlier on the drawbacks of hard highlighting, it would seem that soft highlighting could only help.

5 GENERAL DISCUSSION

To boost human performance on difficult image analysis tasks, techniques have been developed to augment images with boxes or arrows to indicate locations that warrant close scrutiny. However, this sort of all-or-none or *hard* highlighting has not proven robustly effective. Our findings give us optimism to suppose that graded or *soft* highlighting may turn out to be more beneficial. In both tasks where individuals are domain experts (handprinted digit search, Experiments 1–4) and where individuals are relative novices (satellite-imagery search, Experiments 5–7), we have shown a consistent advantage for soft over hard highlighting.

A very recent manuscript by Cunningham et al. (2016) also argues for the superiority of soft over hard highlighting, or—in their terminology—*analog* over *binary*. In their experiments, participants search clusters of colored dots for a target defined by the statistics of its dots. Clusters are highlighted by rendering a ring around the cluster. Hard highlighting is performed via rings that are present or absent, soft highlighting via the shading of the rings. Cunningham et al. found a benefit for soft highlighting, consistent with our results. Their controlled, artificial stimuli allow for the objective assessment of stimulus signal strength, facilitating in principle the investigation of how individuals combine signals from the stimulus and highlights. In contrast, our experiments have real-world validity, utilizing complex stimuli that are familiar and naturalistic.

Another distinction between our work and that of Cunningham et al. is the manner of highlighting. We modulate image lightness (Experiments 1–4) and saturation (Experiments 5–7) directly, whereas Cunningham et al. superimpose distinct display markers to guide attention. We conjecture that the most effective means of communicating graded signals from a classifier will be those that leverage the sensitivity of the human attentional system. The attentional system excels in integrating bottom-up stimulus saliency with top-down task guidance. Our approach, which involves modulating lightness and saturation, provides natural bottom-up cues to guide attention. Of course, other types of highlights may be equally effective, for example, contrast or display flicker.

With cuing via saliency, we have shown that individuals are sensitive to the gradations of soft highlighting, and that even poor quality classifiers can boost human visual search.

To develop even more effective techniques for soft highlighting, we suspect that it will be valuable to conduct a systematic investigation of techniques to represent classifier confidence scores in images. In our work, we aimed for a linear relationship between classifier confidence and attentional saliency. However, it may turn out that convex or concave functions will be more suitable depending on the nature of the classifier, the domain, or the highlighting manipulation.

Another interesting direction for future research concerns the complementarity of humans and machines. One expects the greatest synergy when the skills of humans and machines are orthogonal. For example, our restaurant classifier systematically inspected local image patches, which likely made it superior to humans in picking out the granular features of roofs, parking lots, and so on. In contrast, humans are able to leverage coarser contextual knowledge, such as the fact that restaurants are often located on major streets and intersections. Consequently, the human-classifier combination outperformed either agent alone.

One might question the value of human-machine cooperative methods, given that machine learning methods are likely to advance to a level where they match or surpass human expertise, and they do not suffer the attentional and motivational lapses that confound human experts. Nonetheless, humans are still necessary in the loop to provide initial training data and to refine the data set. The refinement process will be particularly critical and will involve humans working with machines in much the manner we examined in this work, to catch the subtle cases that either human alone or machine alone is likely to miss.

REFERENCES

- E. Alberdi, A. Povyakalo, L. Strigini, and P. Ayton. 2004. Effects of incorrect computer-aided detection (CAD) output on human decision-making in mammography. *Acad. Radiol.* 11, 8 (2004), 909–918.
- J. R. Anderson. 1971. Land-use classification schemes. *Photogram. Eng.* 37, 4 (1971), 379–387.
- C. Balleyguier, K. Kinkel, J. Fermanian, S. Malan, G. Djen, P. Taourel, and O. Helenon. 2005. Computer-aided detection (CAD) in mammography: Does it help the junior or the senior radiologist? *Eur. J. Radiol.* 54, 1 (2005), 90–96.
- I. Bankman. 2008. *Handbook of Medical Image Processing and Analysis*, 2nd ed. Academic Press, San Diego, CA.
- L. Bastin, P. F. Fisher, and J. Wood. 2002. Visualizing uncertainty in multi-spectral remotely sensed imagery. *Comput. Geosci.* 28, 3 (2002), 337–350.
- H.-D. Cheng, X. Cai, X. Chen, L. Hu, and X. Lou. 2003. Computer-aided detection and classification of microcalcifications in mammograms: A survey. *Pattern Recogn.* 36, 12 (2003), 2967–2991.
- C. A. Cunningham, T. Drew, and J. M. Wolfe. 2016. Analog computer-aided detection (CAD) information can be more effective than binary marks (2016). Submitted for publication.
- C. J. D’Orsi. 2001. Computer-aided detection: There is no free lunch. *Radiology* 221, 3 (2001), 585–586.
- T. Drew, C. Cunningham, and J. M. Wolfe. 2012. When and why might a computer-aided detection (CAD) system interfere with visual search? An eye-tracking study. *Acad. Radiol.* 19, 10 (2012), 1260–1267.
- M. T. Dzindolet, S. A. Peterson, R. A. Pomranky, L. G. Pierce, and H. P. Beck. 2003. The role of trust in automation reliance. *Int. J. Human-Comput. Studies* 58, 6 (June 2003), 697–718.
- T. W. Freer and M. J. Ulissey. 2001. Screening mammography with computer-aided detection: Prospective study of 12,860 patients in a community breast center 1. *Radiology* 220, 3 (2001), 781–786.
- M. L. Giger, K. Doi, H. MacMahon, R. M. Nishikawa, K. R. Hoffmann, R. A. Vyborny, C. J. Schmidt, H. Jia, K. Abe, and X. Chen. 1993. An “intelligent” workstation for computer-aided diagnosis. *Radiographics* 13, 3 (1993), 647–656.
- D. M. Green and J. A. Swets. 1966. *Signal Detection Theory and Psychophysics*. Wiley, New York.
- T. Hengl and N. Toomanian. 2006. Maps are not what they seem: Representing uncertainty in soil-property maps. In *Proceedings of the 7th International Symposium on Spatial Accuracy Assessment in Natural Resources and Environmental Sciences*, M. Caetano and M. Painho (Eds.). Instituto Geographico Portugues IGP, 805–813.
- R. Hupse, M. Samulski, M. B. Lobbes, R. M. Mann, R. Mus, G. J. den Heeten, D. Beijerinck, R. M. Pijnappel, C. Boetes, and N. Karssemeijer. 2013. Computer-aided detection of masses at mammography: Interactive decision support versus prompts. *Radiology* 266, 1 (2013), 123–129.
- M. Kaur, J. Kaur, and J. Kaur. 2011. Survey of contrast enhancement techniques based on histogram equalization. *Int. J. Adv. Comput. Sci. Appl.* 2, 7 (2011), 137–141.
- R. T. Kneusel. 2016. *Improving human-classifier interaction through enhanced highlighting techniques*. Ph.D. Dissertation. Department of Computer Science, University of Colorado at Boulder.

- E. A. Krupinski, C. F. Nodine, and H. L. Kundel. 1993. Perceptual enhancement of tumor targets in chest x-ray images. *Percept. Psychophys.* 53, 5 (1993), 519–526.
- Y. LeCun, Y. Bengio, and G. E. Hinton. 2015. Deep learning. *Nature* 521 (2015), 436–444.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- S.-K. Lee, C.-S. Lo, C.-M. Wang, P.-C. Chung, C.-I. Chang, C.-W. Yang, and P.-C. Hsu. 2000. A computer-aided design mammography screening system for detection and classification of microcalcifications. *Int. J. Med. Informat.* 60, 1 (2000), 29–57.
- C. D. Lehman, R. D. Wellman, D. S. M. Buist, K. Kerlikowske, A. N. A. Tosteson, and D. L. Miglioretti. 2016. Diagnostic accuracy of digital screening mammography with and without computer-aided detection. *JAMA Intern. Med.* 11 (2016), 1828–1837.
- A. M. MacEachren, A. Robinson, S. Hopper, S. Gardner, R. Murray, M. Gahegan, and E. Hetzler. 2005. Visualizing geospatial information uncertainty: What we know and what we need to know. *Cartogr. Geogr. Info. Sci.* 32, 3 (2005), 139–160.
- M. E. J. Masson and G. R. Loftus. 2003. Using confidence intervals for graphically based data interpretation. *Can. J. Exp. Psychol./Revue canadienne de psychologie expérimentale* 57, 3 (2003), 203.
- M. C. Mozer. 1991. *The Perception of Multiple Objects: A Connectionist Approach*. MIT Press, Cambridge, MA.
- R. Parasuraman and V. Riley. 1997. Humans and automation: Use, misuse, disuse, abuse. *Human Factors* 39, 2 (1997), 230–253.
- L. E. Philpotts. 2009. Can computer-aided detection be detrimental to mammographic interpretation? 1. *Radiology* 253, 1 (2009), 17–22.
- J. A. Richards and X. Jia. 1999. *Remote Sensing Digital Image Analysis: An Introduction*, 4th ed. Springer, Berlin.
- B. Sahiner, H.-P. Chan, L. M. Hadjiiski, P. N. Cascade, E. A. Kazerooni, A. R. Chughtai, C. Poopat, T. Song, L. Frank, J. Stojanovska, and others. 2009. Effect of CAD on radiologists' detection of lung nodules on thoracic CT scans: Analysis of an observer performance study by nodule size. *Acad. Radiol.* 16, 12 (2009), 1518–1530.
- M. Samulski, A. Hupse, C. Boetes, G. den Heeten, and N. Karssemeijer. 2009. Analysis of probed regions in an interactive CAD system for the detection of masses in mammograms. *Proceedings of SPIE* 7263 (2009), 726314–726314–9. DOI: <http://dx.doi.org/10.1117/12.813391>
- M. Samulski, R. Hupse, C. Boetes, R. D. M. Mus, G. J. den Heeten, and N. Karssemeijer. 2010. Using computer-aided detection in mammography as a decision support. *Eur. Radiol.* 20, 10 (2010), 2323–2330.
- J. Snoek, H. Larochelle, and R. P. Adams. 2012. Practical bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., La Jolla, CA, 2951–2959.
- A. M. Treisman and G. Gelade. 1980. A feature-integration theory of attention. *Cogn. Psychol.* 12, 1 (1980), 97–136.
- F. J. M. Van der Wel, L. C. Van der Gaag, and B. G. H. Gorte. 1998. Visual exploration of uncertainty in remote-sensing classification. *Comput. Geosci.* 24, 4 (1998), 335–343.
- J. L. Van Genderen, B. F. Lock, and P. A. Vass. 1978. Remote sensing: Statistical testing of thematic map accuracy. *Remote Sens. Environ.* 7, 1 (1978), 3–14.
- M. Yasmin, M. Sharif, S. Masood, M. Raza, and S. Mohsin. 2012. Brain image enhancement—A survey. *World Appl. Sci. J.* 17, 9 (2012), 1192–1204.
- B. Zheng, M. A. Ganott, C. A. Britton, C. M. Hakim, L. A. Hardesty, T. S. Chang, H. E. Rockette, and D. Gur. 2001. Soft-copy mammographic readings with different computer-assisted detection cuing environments: Preliminary findings 1. *Radiology* 221, 3 (2001), 633–640.
- B. Zheng, R. G. Swensson, S. Golla, C. M. Hakim, R. Shah, L. Wallace, and D. Gur. 2004. Detection and classification performance levels of mammographic masses under different computer-aided detection cueing environments 1. *Acad. Radiol.* 11, 4 (2004), 398–406.

Received December 2016; revised June 2017; accepted June 2017