

Attention Allocation Aid for Visual Search

Arturo Deza*
Dynamical Neuroscience
UC Santa Barbara
deza@dyns.ucsb.edu

Jeffrey R. Peters
Mechanical Engineering
UC Santa Barbara
jrpeters@engineering.ucsb.edu

Grant S. Taylor
U.S. Army
Aviation Development
Directorate, USA
grant.s.taylor.civ@mail.mil

Amit Surana
United Technologies Research
Center, USA
suranaa@utrc.utc.com

Miguel P. Eckstein
Psychological and Brain
Sciences, UC Santa Barbara
eckstein@psych.ucsb.edu

ABSTRACT

This paper outlines the development and testing of a novel, feedback-enabled *attention allocation aid* (AAAD), which uses real-time physiological data to improve human performance in a realistic sequential visual search task. Indeed, by optimizing over search duration, the aid improves efficiency, while preserving decision accuracy, as the operator identifies and classifies targets within simulated aerial imagery. Specifically, using experimental eye-tracking data and measurements about target detectability across the human visual field, we develop functional models of detection accuracy as a function of search time, number of eye movements, scan path, and image clutter. These models are then used by the AAAD in conjunction with real time eye position data to make probabilistic estimations of attained search accuracy and to recommend that the observer either move on to the next image or continue exploring the present image. An experimental evaluation in a scenario motivated from human supervisory control in surveillance missions confirms the benefits of the AAAD.

ACM Classification Keywords

H.1.2 [Models and Principles]: User/Machine Systems—Human factors, Human information processing; H.4.2 [Information Systems Applications]: Types of Systems—Decision Support; H.m [Miscellaneous]; I.6.4 [Simulation and Modeling]: Model Validation and Analysis; General Terms: Experimentation, Human Factors, Verification

Author Keywords

Attention; cognitive load; decision making; visual search.

⁰All authors of the paper are also affiliated with the Institute for Collaborative Biotechnologies.

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 ACM 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.

CHI 2017, May 6–11, 2017, Denver, CO, USA.

Copyright © 2017 ACM ISBN 978-1-4503-4655-9/17/05 ...\$15.00.

<http://dx.doi.org/10.1145/3025453.3025834>

INTRODUCTION

The maturation of visual sensor technology has steadily increased the amount of real-time data that is available in modern surveillance mission scenarios ranging across military, homeland security and commercial applications. In many cases, it is the job of a human operator to ensure that this data is processed quickly and accurately. For example, *supervisory systems* involving collaboration between human operators and unmanned vehicles often require the sequential processing of imagery that is generated by the autonomous vehicles' on-board cameras for the purpose of finding targets, analyzing terrain, and making key planning decisions [39]. The incredible volume of data generated by modern sensors, combined with the complex nature of modern mission scenarios, makes operators susceptible to information overload and attention allocation inefficiencies [9], which can lead to detrimental performance and potentially dire consequences [44]. As such, the development of tools to improve human performance in visual data analysis tasks is crucial to ensuring mission success.

This article focuses on the development and experimental verification of a novel, attention allocation aid that is designed to help human operators in a sequential visual search task, which requires the detection and classification of targets within a simulated landscape. Our study is primarily motivated by surveillance applications that require humans to find high value targets within videos generated by remote sensors, e.g., mounted on unmanned vehicles; however, the presented method is applicable to a variety of application domains.

Specifically, the main contribution of this paper is the introduction and experimental verification of a real-time and feedback-enabled *attention allocation aid* (AAAD), which optimizes the operator's decision speed when they are engaging in target search, without sacrificing performance. The motivating observation is that humans have imperfect awareness of the time required to acquire all task-relevant visual information during search, and thus are generally inefficient at administering their time when scrutinizing the data sets. The proposed aid makes real-time automated search duration recommendations based on three key metrics: 1) visual search time, 2) number of eye movements executed by the observer, and 3) an estimated

target detectability based on prior measurements of both the target visibility across the visual field and the observer's fixations during search. In particular, these metrics are used by the aid to estimate the time required for the operator to acquire the visual information that is necessary to support the search decisions, and subsequently indicate when this time has elapsed via a simple indicator on the user interface. We experimentally evaluate the AAAD in a simulated surveillance scenario motivated by human supervisory control, and found a factor of $\times 1.5$ increase in user efficiency in detecting and classifying targets in realistic visual imagery from a slowly moving sensor. The AAAD pipeline is generic and can readily be extended and applied to other sources of images, i.e., satellite images, astronomical images [6], x-rays for medical imaging [1, 13] or security scanning [5], and video surveillance [45] which include a human-in-the-loop [53, 52].

Our rigorous development of the AAAD also includes a number of secondary contributions, including: definition of detectability surfaces based on eye-tracking measurements, incorporation of image clutter effects, creation of a composite exploration map, and utilization of a probabilistic framework for decision making when computing overall search satisfaction based on time, eye movements, and detectability scores.

PREVIOUS WORK

Research in human computer interaction has capitalized on basic vision science research by including samples of the visual environment through eye movements (active vision) within a larger modeling framework that includes cognitive, motor, and perceptual processing involved in visual search (Tseng and Howes, 2015 [48] and Halverson and Hofson, 2011 [20] which expands on the work of Kieras and Meyer, 1997 [32]). Such models have been proposed for potential use for computer interface evaluation and design.

Another line of research has focused on how to augment human capabilities in coordinating multiple tasks. For example, models have been used to optimize how observers split their attentional resources when simultaneously conducting two different visuo-cognitive tasks [36].

Attention allocation aids have been studied in the context of human supervisory control of large data acquired by multiple automated agents (e.g., [46, 9]). Such scenarios present the challenge of a human having to inspect large data sets with possible errors due to visual limitations, attentional bottlenecks, and fatigue. The use of advanced physiological sensing through eye-tracking technology has become a viable option for both the assessment of the operator cognitive state, and the evaluation of operator performance in a number of realistic applications, e.g. [49]. One line of research attempts to use eye-tracking measurements to detect physiological and cognitive precursors to behavior such as perceived workload, fatigue, or situational awareness. Indeed, objective measures such as blink rates [51], pupil diameter [50, 31], and fixation/saccade characteristics [2], all have correlations to cognitive processing, although the use of such measurements as reliable indicators of operator mental states is not fully understood [12]. If undesirable states can be accurately anticipated with physiological measures, then they can be used to drive

automated aids that mediate operator resources through, e.g., optimization of task schedules [38] or adaptive automation schemes [22, 43].

Other researchers have utilized eye tracker technology to show increased efficiency of human search by relying on a "divide and conquer" collaborative search [57, 7]. In such schemes, multiple observers (usually two) engage in target search simultaneously with real-time updates of their partners' gaze and access to voice communication.

A novel approach investigated in the current work is the design of an attention allocation aid that uses eye-tracking data to make real time inferences of the attained target detection accuracy and critically, the time to achieve asymptotic accuracy. Such estimates, which we will refer to as search satisfaction time, are utilized by the attention allocation aid to recommend that the user end the current search and move on to the next data set. In addition, if the observer completes search prior to the search satisfaction time, the eye position data can also be utilized to assess whether some area of the image remains unexplored, and suggest that the observer to further explore that area.

The success of the proposed approach requires an adequate understanding of the relation between fixational eye-movements and the accumulation of sensory evidence supporting task performance. A critical component to understanding the contribution of eye movement to task performance is the dependence of target detectability with its distance from the point of fixation, commonly referred to as *retinal eccentricity* [11, 37, 33, 10]. Indeed, this relationship can be used to build attention-based models for predicting performance [24]. Often, dynamic sensory evidence accumulation models are also dependent upon the nature of the stimuli. Our attention allocation aid relies on a set of experiments measuring how accuracy in detecting the target of interest varies with distance from fixation (retinal eccentricity) and as a function of the presentation time of the image data. These measurements are then used to implement the AAAD and validate its utility in optimizing search. To our knowledge the current approach for the AAAD and thorough experimental validation is novel to the field.

We also note that a key difference of our work in comparison to existing literature, is that our attention allocation aid is essentially a back-end search optimizer, which tells the observer *when* to stop search; rather than advising the observer *where* to look (it does not compute fixation cue's or saliency-like maps [23, 8]).

MOTIVATION AND HYPOTHESIS

Humans have difficulty assessing when adequate visual information has been acquired during challenging search tasks and optimally allocating their fixations over different parts of the image [55]. The purpose of the AAAD is to utilize in real time the temporal dynamics of the eye-position data and the information acquisition process to recommend to the observers that either all information has been acquired and search can be terminated, or further exploration of the image is required. The AAAD is expected to reduce both premature image search termination and long periods of image search

when no target is present without compromising the search task performance, i.e. detection and false alarm rates. Thus, the AAAD should ideally improve observer's efficiency in completing more sequential search tasks in a given allotted time period with a level of detection accuracy that is as good or better than search without the AAAD.

In what follows, we present a series of experiments in order to develop, calibrate, and test the effectiveness of the AAAD in a visual target search/classification task. Subjects were asked to search for people within simulated aerial images, and subsequently classify whether or not the person was holding a weapon. The stimuli and eye-tracking apparatus used in all of our experiments are described below.

Stimuli Creation: A total of 273 videos were created, each with a total duration of 120 seconds, where a 'birds eye' point-of-view camera rotated slowly around the center. While the video was in a rotating motion, there was no relative motion between any parts of the video. From a repeated subset of the original 273 videos, a total of 1440 different short clips were created, which were subsequently divided into the 4 groups (stimuli sets) that were used in subsequent experiments. Half of the clips had person present, while the other half had person absent. These short and slowly rotating clips were used instead of still images in our experiment, to simulate imagery from a moving sensor in a surveillance scenario. All clips were shown to participants in a random order. The stimuli used in all our experiments present varying levels of zoom (high, medium, low) and clutter (high, medium, low).

Apparatus: An EyeLink 1000 system (SR Research) was used to collect eye-tracking data at a frequency of 1000 Hz. Each participant sat at a distance of 76 cm from a LCD screen on gamma display, so that each pixel subtended a visual angle of 0.022 deg/px. All video clips were rendered at 1024 × 760 px (22.5 deg × 16.7 deg) and a frame rate of 24 fps. Eye movements with velocity over 22 deg/s and acceleration over 4000 deg/s² were qualified as saccades. Every trial began with a fixation cross, where each subject had to fixate the cross with a tolerance of 1 deg.

DEVELOPING THE ATTENTION ALLOCATION AID (AAAD)

This section presents details about the development and calibration of the AAAD. Here, preliminary experiments were run to estimate perceptual performance curves (PPCs) (describing the relationship between detection accuracy and each of three metrics: time, number of eye movements, and detectability) under various levels of image complexity (clutter and zoom levels). These PPCs allow the identification of asymptotic detection accuracy, which is the primary criterion used by the AAAD to estimate expected detection accuracy in real time.

Experiment 1: Psychometric Data Collection

We performed two preliminary studies to generate the time, eye movement, and detectability PPCs: a forced fixation search (no eye movements allowed) and a free search experiment. The free search data is directly used to compute the relevant PPCs, while the forced fixation data is used to calculate detectability surfaces that allow for the computation of the detectability PPC. See Figure 1 for experimental flow.

Forced Fixation Search

A total of 13 subjects participated in a *forced fixation search* experiment where the goal was to search within the visual periphery to identify if there was a person present or absent (yes/no task; 50% probability of person presence) and, in addition, to identify if there was a weapon present or absent (yes/no task; 50% probability of weapon presence contingent on person present). Participants had variable amounts of time (100, 200, 400, 900, 1600 ms) to view each clip. Clips were presented in a random order, with the person at a variable degree of eccentricity (1 deg, 4 deg, 9 deg, 15 deg) from point of fixation. Subjects were not made aware of the eccentricity values used in each trial. They were then prompted with a Likert scale that required them to rate from 1-10 (by clicking on a number) their confidence of person presence. A value of 1 indicated strong confidence of person absent, and a value of 10 indicated a strong confidence of person present – intermediate values represented different levels of uncertainty. Values of 1-5 were classified as person absent, and 6-10 were classified as person present. A second rating scale (identical to the first) was then presented, requiring the subject to rate their confidence regarding weapon presence. Participants had unlimited time for making their judgments, although no subject ever took more than 10 seconds per judgment. There was no response feedback after each trial.

Each subject participated in 12 sessions that consisted of 360 clips each. There were 4 stimuli sets (each set consisted of unique images), and each participants viewed each set 3 times in random order without being aware that the images were repeated (4 sets × 3 times = 12 sessions). Every set also had the images with aerial viewpoints from different vantage points (Example: set 1 had the person at 12 o'clock – as in North, while set 2 had the person at 3 o'clock – as in East). To mitigate fixation bias, all subjects had a unique fixation point for every trial associated with each particular eccentricity value. All clips were rendered with variable levels of clutter. Each session took approximately one hour to complete. The person, i.e. search target, was of size 0.5 deg × 0.5 deg, 1 deg × 1 deg, 1.5 deg × 1.5 deg, depending on the zoom level. If a subject fixated outside of a 1 deg radius around the fixation cross during the trial, then the trial was aborted.

Free Search

A total of 11 subjects participated in a *free search* experiment where the goal was to detect and classify the person. Although eye movements were allowed, subjects were not explicitly told to foveate at the person (although they usually chose to do so). Participants had twice the amount (200 ms, 400 ms, 800 ms, 1800 ms, 3200 ms) of time than in the Forced Fixation Search. All observers began each trial with a fixation at center of the screen. They then proceeded to scan the scene to find a person and determine if the person was holding a weapon. Once the trial time was over, they were prompted with a "Person Detection" and "Weapon Detection" rating scale, and had to rate from 1-10 by clicking on a number reporting how confident they were on detecting/classifying the person. Similar to the forced fixation experiment, participants had unlimited time to make their judgments and did not receive any feedback after each trial. No trials were aborted.

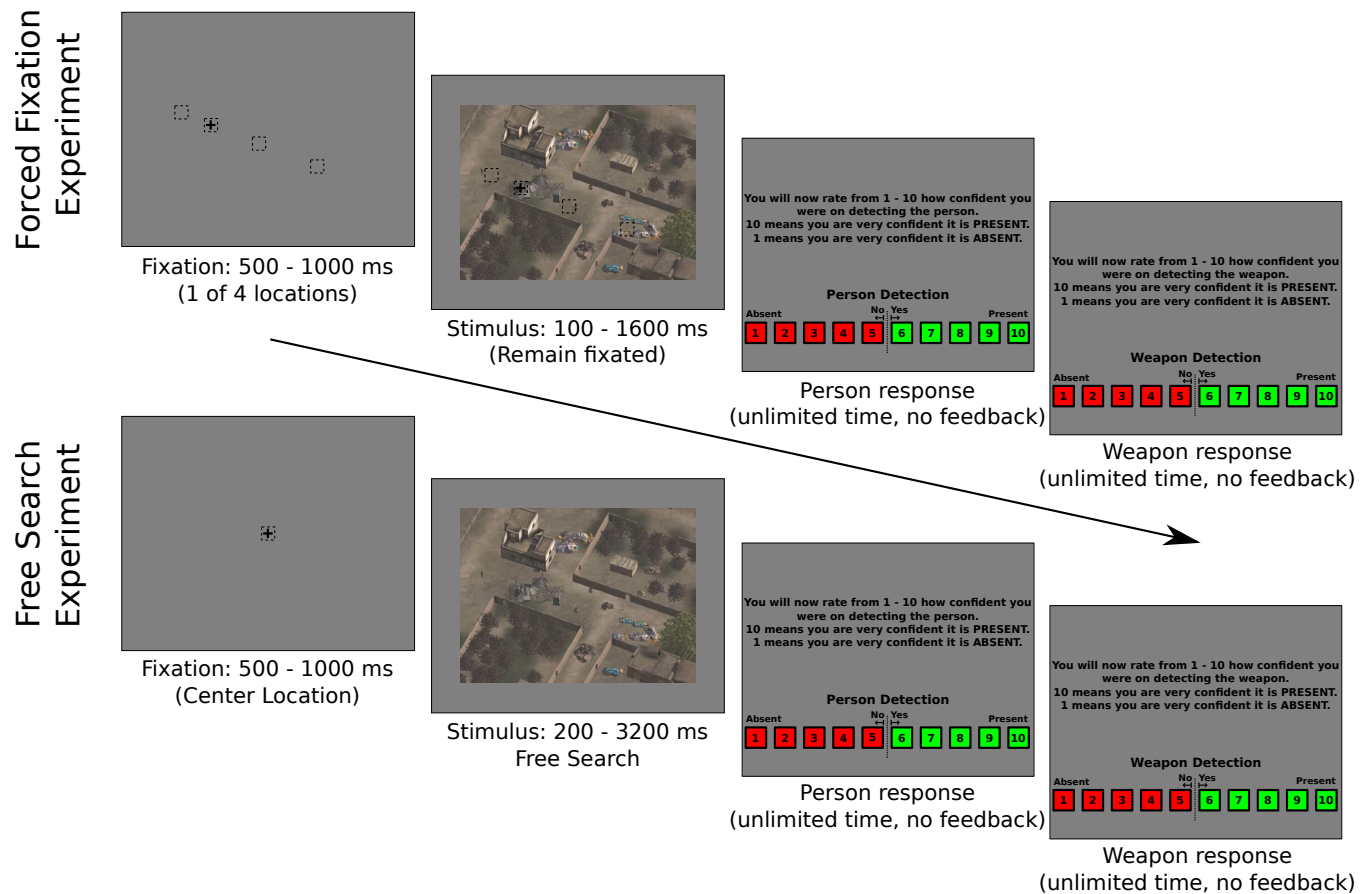


Figure 1. Experiment 1: The forced fixation (top) and free search (bottom) experiments to obtain the time, eye movements, and detectability PPCs.

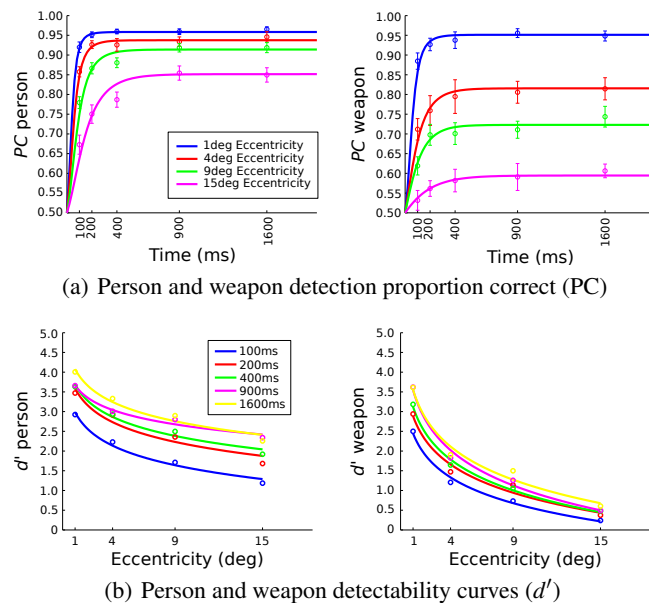


Figure 2. Person and weapon detection performance in proportion correct (PC) and d' space from the forced fixation search experiment. Notice that (a) and (b) are dual representations of each other. The bottom curves in d' space will be used to generate a detectability surface.

Each subject in the free search experiment participated in 6 sessions that consisted of two sets of 360 unique images. In these sessions, each subject viewed one of the two sets of images, and each set was presented 3 times leading to a total of 6 sessions. Subjects were not made aware that the sessions were repeated.

Fitting Perceptual Performance Curves (PPCs)

Motivation of PPCs

PPCs were constructed to relate performance to each of three different metrics. The first metric is visual search time, since it is well known that time affects visual search accuracy [15] – the main intuition being that the more time a subject spends scanning an image, the higher the likelihood of detecting the target (person or weapon). The second metric is the number of eye movements a subject performs while engaging in target search. Typically, time will pass on as more eye movements are produced, but there are some cases where scrutiny in classifying or detecting a target is needed by spending long periods of fixation. As an example, one could imagine an *exploitation vs exploration* search scenario where a subject spends 1000 ms on a single fixation, given the difficulty to *classify* the target (exploitation), as opposed to a scenario where the same subject makes 3 sparse and exploratory fixations in the same 1000 ms time window to *find* the target (exploration). For this

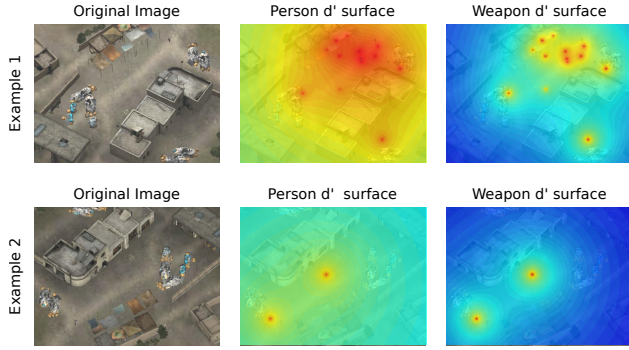


Figure 3. Sample person and weapon fixation maps generated from the forced fixation search experiment (Fig. 2(b)). These fixation maps are projections of Detectability Surfaces as described in the Supp. Mat.

reason we chose to make time and eye movements independent metrics for our AAAD system. The third and last metric is detectability. Here, a detectability score is constructed by generating a pixel-wise map that quantifies localized information aggregation in different parts of the image (as indicated by eye movements), and subsequently combining the result to quantify the target’s overall detectability. Following our previous example, one could imagine that even if an observer spends the allotted 1000 ms searching for a target and making e number of eye movements in a small spatial vicinity, it might not be a good strategy compared to spreading fixations across the image. See Figure 3 for an example of such fixations overlaid on different images.

We are interested in successful observer detection of the person and the weapon targets. Given that our results show that the weapon requires more time to detect than the person, the AAAD recommendation to end search was based on the PPCs corresponding to the detection of the weapon. Basing the AAAD on the PPCs for the person detection would likely compromise the detection of the weapon.

Computing PPCs

To model the target detection accuracy, we use the observer hit rate (the proportion of trials that the observer indicated that a target is present, given that a target is actually present in the trial stimuli) and false alarm rate (the proportion of trials that the observer indicated that the target is present, given that no target is actually present in the trial stimuli). Hit rates and false positive were represented as an empirical detectability index (d') and a decision criterion (λ) using an equal variance normal Signal Detection Theory (SDT) model (Green & Swets) [18]. We then fit the resulting data with curves to model the functional relationship between the detectability indices and each of the relevant performance metrics. The best fit functions were then utilized with the equal variance SDT model to generate estimates of attained accuracy in terms of proportion correct (PC). Notice that proportion correct and hit rate are different since proportion correct takes into account both the hit rate and the correct rejection rate (proportion of trials in which the observer correctly decided that the target is absent).

For a fixed condition and setting (assuming Gaussian signal and noise distributions), the general equations to compute (d' , λ) are [54]:

$$d' = Z(\text{Hit Rate}) - Z(\text{False Alarm Rate}) \quad (1)$$

$$\lambda = -Z(\text{False Alarm Rate}) \quad (2)$$

where $Z(\circ)$ is the inverse of the normal cumulative Gaussian distribution, and the hit/false alarm rates are calculated at the given experimental condition and setting. Consider as an example a condition and setting for the forced fixation search experiment: Condition = (4 deg, 200 ms), Setting = (high zoom, low clutter). Likewise, a sample condition and setting for the free search experiment: Condition = 400 ms, Setting = (medium zoom, high clutter).

For the Time and Eye-Movements PPCs, we used a straight-forward regression to find an exponential relation of the form $d'(x) = \alpha(1 - e^{-\beta x})$, where $x : \{\text{Time, Eye Movements}\}$ and β is constant, to obtain a continuous approximating function for the collection of points d' within a given setting.

To compute final Time and Eye Movements PPC curves, recall that there exists a function $g(\circ)$ that estimates PC, i.e., $PC(x) = g(d'(x), \lambda(x))$, where:

$$PC(x) = m(x)(\text{Hit Rate}) + n(x)(1 - (\text{False Alarm Rate})) \quad (3)$$

and

$$\text{False Alarm Rate}(x) = Z^{-1}(-\lambda(x)) \quad (4)$$

$$\text{Hit Rate}(x) = Z^{-1}(d'(x) - \lambda(x)) \quad (5)$$

where $m(x)$ and $n(x)$ are variables that are contingent on the number of signal (i.e. person/weapon) present and signal (i.e. person/weapon) absent trials ($m(x) + n(x) = 1, \forall x$), and the hit/false alarm rates are estimates of the true values at x . Here, curves are fitted in d' space, rather than directly from PC space, to deal with possible unbalanced datasets with signal present/absent trials (See Discussion).

In order to obtain the detectability PPC (Figure 4 (right)), we perform a binned regression across all trials between the respective PC performance from the free viewing task in Experiment 1 and the values of a composite detectability score D' for each image. The score D' in each trial was computed by first generating a *detectability surface* from the user’s fixations and the detectability curves (Figure. 2(b)), and subsequently performing a spatial average. We note that the detectability PPC is the only curve that is regressed directly to PC given that the argument of our regression function is in a d' -like space, thus $\lambda, m(x), n(x)$ need not be calculated. The training data we have for this regression is from the free search experiment.

Further details regarding the generation of the Time PPC, Eye Movements PPC, and Detectability PPC is provided in the Supplementary Material.

Performance Criteria and AAAD Functionality

The AAAD was designed to integrate three different inputs to compute search satisfaction. The three previous PPC in-




Figure 4. Perceptual Performance Curves (PPCs) for time (left), eye movements (center), and detectability (right) for a (low zoom, high clutter) setting. In general it takes higher time, eye movements, and detectability scores to achieve asymptotic performance for weapon detection vs person detection. Also shown are the error bars along each curve. Recall they have been fitted in d' space, and have been re-plotted in PC space (Eq. 3).

puts can be seen as individual metrics on their own, and are computed independently in the system.

Search Satisfaction Model

For computing general search satisfaction we require that all three of the following conditions are simultaneously satisfied:

$$Pr[(PC_{max}^T - PC(t)) < \epsilon] > \eta, \quad (6)$$

$$Pr[(PC_{max}^E - PC(e)) < \epsilon] > \eta, \quad (7)$$

$$Pr[(PC_{max}^D - PC(D')) < \epsilon] > \eta, \quad (8)$$

where $PC_{max}^T, PC_{max}^E, PC_{max}^D$ are the (fixed) asymptotic values of PC with respect to the time, number of eye-movements, and detectability PPCs given a (zoom, clutter) setting, resp. $PC(t), PC(e), PC(D')$ are the current estimated values of PC as calculated by the time, number of eye-movements, and detectability PPCs, respectively, and ϵ, η are fixed thresholds.

An image is only said to have been adequately searched if *all* criteria (6), (7), and (8) are simultaneously satisfied. Also notice that the criteria (6), (7), and (8) are all non-decreasing in their respective arguments (which are monotonically increasing in time); thus a condition will never revert to being “unsatisfied” after being satisfied. We will refer the above search satisfaction criterion as *PC general satisfaction*.

Our motivation for using the above mentioned probabilistic framework is to take into account the error bars of each time, eye movements, and detectability curves that are zoom and clutter level dependent. We include these error bars as Gaussian standard deviations σ in our probabilistic computation:

$$Pr[(PC_{max} - PC) < \epsilon] \rightarrow 1 - Z^{-1}\left(\frac{PC_{max} - PC - \epsilon}{\sigma}\right). \quad (9)$$

The above strategy can be thought of centering a gaussian $(\mu_x, \sigma_x) := (PC, \sigma)$ at every point in the PPC curves, and computing how far away the asymptotic performance PC_{max} is from every point in the curve. Thus, we will find and select the minimum point in the curve that fulfills this condition for each time, eye movements, and detectability PPC. These are the threshold PPC’s that once all of them are reached in real-time on the AAAD system, the AAAD will trigger “On”. A value of $\eta = 0.025, \epsilon = 0.02$ was selected for our experiments.

Attention Allocation Aid Design

The Attention Allocation Aid system is designed to run in the background of the simulation interface (see Figure 5). The AAAD system starts by computing the clutter and zoom level of the input image from the slowly evolving video clip in each trial. It is assumed that the level of zoom (high, medium, or low) can be obtained from ground truth settings, given that a pilot can control a camera’s zoom level, while the clutter level can be computed from the input image/video [56, 41, 35]. For our main experiments, we assumed that ground truth was provided to classify images based on clutter, since our main goal is to prove that the AAAD system works under ideal conditions¹. A thorough investigation of the use of clutter models as ground truth predictors is left to future work.

As the trial progresses, the three PC vs {Time, Eye Movements, Detectability} curves (Figure 4) are updated in real-time using gaze location, fixation time, and saccade information obtained from the eye-tracker. The Time PPC is updated at each frame. The Eye Movements PPC is updated with an interruption-based paradigm – contingent on the eye-tracker detecting an eye movement. The Detectability PPC is updated only after every eye movement event, given that we have the fixation position and time. Each of the PC Satisfaction conditions (see Eq. (6, 7, 8)) is monitored independently, and once all criteria are satisfied, the AAAD systems switches from an “Explore” to a “Move On” state, where observers are encouraged to cease search, and make a decision.

Parallel to this, an *Exploration* map is computed in real-time in the back-end. The goal of the Exploration map is to inform the searcher where he/she has already searched, and to indicate the highly cluttered regions where a person is likely to be. The Exploration map has no knowledge of a target present/absent, and only uses image clutter and observer fixations.

If an observer attempts to advance to the next image while the AAAD system state is in “Explore” state, the Exploration map appears for ~ 120 ms. The map is weighted by previously explored regions (computed via the detectability surface; see Supplementary Material), such that, highly

¹ If we did *not* assume *oracle-like* inputs for zoom and clutter levels, then not finding a significant effect, could be due to poor clutter modeling, rather than poor AAAD system design.

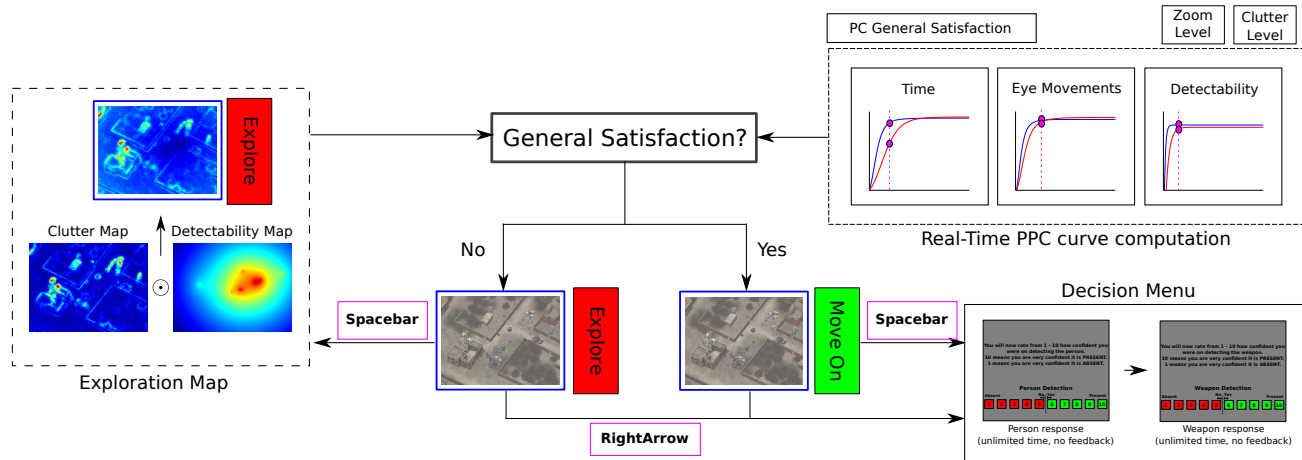


Figure 5. Attention Allocation Aid (AAAD) system diagram. From the start of each trial, the PPC curves are updated in real-time in the back-end, waiting for General Search Satisfaction to be achieved. Possible user inputs (highlighted in magenta) are the space bar and the right arrow. The right arrow takes the user directly to the decision menu terminating the trial, regardless of Satisfaction. The spacebar will lead the observer to the Exploration Map, and later back to the original stimuli if Satisfaction is not achieved, or to the decision menu otherwise

cluttered and non-explored regions are highlighted. The Exploration map is computed as follows: $\text{Exploration map} = \text{FC} \odot (1 - \text{Detectability Surface})$, where FC is the *feature congestion* [41] dense clutter map, \odot is the element-wise multiplication operator, and the detectability surface is normalized to lie in the interval $[0, 1]$.

Notice that the only inputs (to the system) that the observer can produce while performing a trial (besides passively providing eye movements), are by pressing the right arrow key, which forces the trial to terminate and the subject to make a decision, irrespective of PC general satisfaction being achieved (for both AAAD, and non-AAAD experimental sessions); or by pressing the space bar, which activates the Exploration map if PC general satisfaction is not achieved, and terminates the trial if otherwise.

EVALUATING THE ATTENTION ALLOCATION AID

In this section we summarize a second experiment used to evaluate the effectiveness of the AAAD. Two experimental conditions were considered: A person and weapon search and classification experiment with and without AAAD. The goal of this experiment is to objectively measure any improvements in the search task performance when the subjects are assisted by the AAAD. Since it was the first time for all of our second group of subjects to participate in an eye tracking experiment, we decided to add two additional practice sessions (twenty minutes each) where we would verbally explain the non-AAAD and AAAD system.

After the practice sessions, half of the subjects were tested starting with the AAAD condition and the other half started without the AAAD condition. We counterbalanced our participants to reduce possible learning effects. The group of participants involved in Experiment 2, did not participate in and were not aware of Experiment 1. The completion of the first 2 practice sessions plus Experiment 2 with both conditions (counterbalanced) took an estimate of 2 hours for each subject. The same person and weapon present/absent statistics

of Experiment 1 were used for Experiments 2 and 3. Figure 6 illustrates the design of the search task without (top stream) and with (bottom stream) the AAAD.

Experiment 2: AAAD Evaluation

Condition 1: Target search without AAAD

A total of 18 subjects participated in a non-AAAD target search experiment where the goal was to complete as many trials as possible in a 20 minute interval without sacrificing task performance, where the task per trial was to correctly detect and classify the target in the minimum amount of time.

Target detection involved reporting if the person was present or absent in the scene, and target classification involved reporting whether or not the person had a weapon. Although eye movements were allowed, subjects were not told explicitly to foveate at the target (although this was usually the case). In other words, it was possible for the subject to move on to the next trial by detecting the target in the periphery [10]. A fixation cross was placed at the center of the screen for uniform starting conditions across participants. After terminating search, subjects were prompted with a “Person Detection” rating scale where they had to rate their confidence in a person’s presence on a scale from 1-10 by clicking on a number. A “Weapon Detection” rating scale then appeared where subjects also had to rate their weapon detection confidence from a scale from 1-10. Participants had unlimited time for making their judgments, though no subject took more than 10 seconds per judgment. There was no response feedback, i.e., whether their detection responses were correct after each trial.

Condition 2: Target search with AAAD

The same 18 subjects participated in a target search experiment in presence of the AAAD where the goal was same as in the previous condition. In this experiment, the AAAD was visibly turned on for the participants. They saw a text message above the center stimuli with a caption: “Explore” or “Move On”, and there was a colored square below the stimuli that

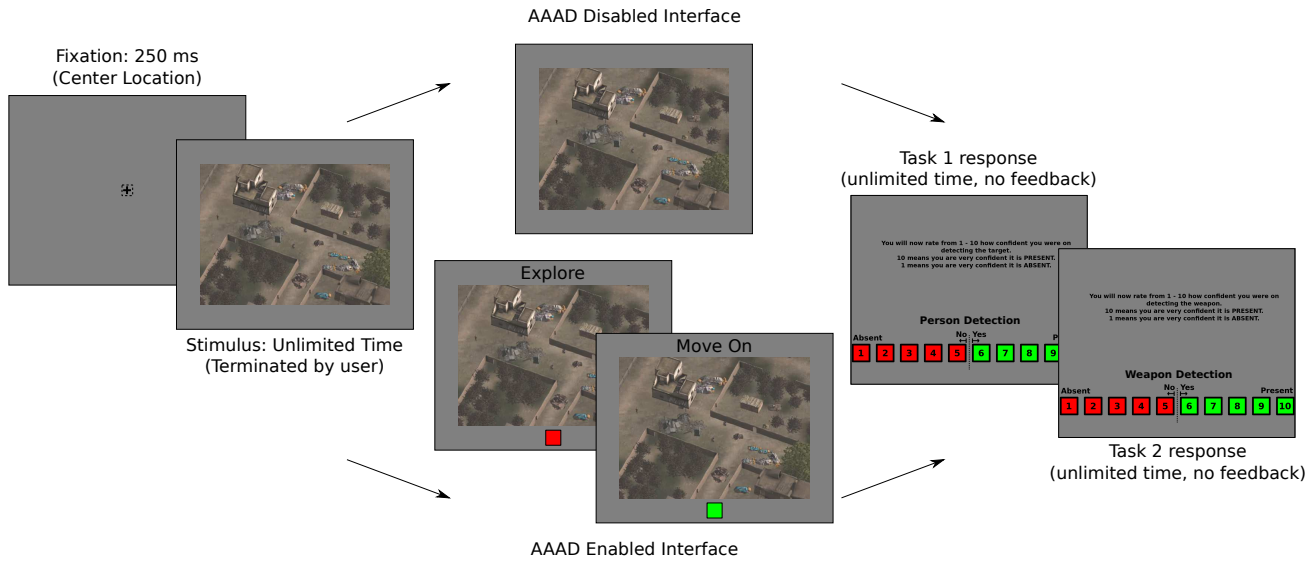


Figure 6. Experimental flow for Experiment 2 where we test the Attention Allocation Aid (AAAD) for condition 1 (AAAD off; top) and condition 2 (AAAD on; bottom). The top stream illustrates the non-AAAD condition, while the bottom stream illustrates the AAAD condition. Notice that the text and color in the AAAD enabled interface are coupled together: “Explore” = red, “Move On” = green. The red and green colors in the decision menus are independent of the AAAD colors.

	System Evaluation	
	Non-AAAD	AAAD
Average Trial Number (#)	54.33 ± 2.26	57.94 ± 1.74
Average Mean Time per Trial (s)	2.88 ± 0.42	1.94 ± 0.18
Person Hit Rate (%)	96.39 ± 1.42	96.94 ± 0.88
Weapon Hit Rate (%)	89.37 ± 2.92	92.35 ± 2.30
Person False Alarm Rate (%)	2.88 ± 2.63	2.42 ± 2.04
Weapon False Alarm Rate (%)	8.19 ± 5.54	7.15 ± 4.70
Person Miss Rate (%)	3.61 ± 1.42	3.06 ± 0.88
Weapon Miss Rate (%)	10.63 ± 2.92	7.65 ± 2.30
Person Correct Rejection Rate (%)	97.12 ± 2.63	97.58 ± 2.04
Weapon Correct Rejection Rate (%)	91.81 ± 5.54	92.85 ± 4.70
Mean Trial time vs Time Trigger (s)	2.32 ± 0.10	1.37 ± 0.05
Mean Trial time vs EyeMvmt Trigger (s)	2.35 ± 0.10	1.37 ± 0.05
Mean Trial time vs Detect. Trigger (s)	2.70 ± 0.13	1.10 ± 0.06
Mean Trial time vs General Trigger (s)	2.70 ± 0.13	1.10 ± 0.06

Table 1. General results of the systems evaluation without and with AAAD. It should be noted that subjects were counterbalanced (half-split) to start with or without the AAAD during evaluation to compensate for learning effects. Average refers to the mean computed across observers.

was colored red or green depending on the AAAD status. Participants were told to think of the AAAD as a spotlight: when it was red they should keep looking for the person/weapon, and should only move on to the next trial if the light turned green or if they were confident that they had either found the person/weapon or there was no person/weapon present.

RESULTS

Table 1 summarizes the results of both conditions in the second experiment. There is a significant increase in number of trials per person ($M = 3.61$, $SD = 5.69$, $t(17) = 2.813$, $p = 0.015$, two-tailed), as well as a significant decrease in mean trial time ($M = -0.31$, $SD = 1.36$, $t(17) = -2.613$, $p = 0.009$, two-tailed) between AAAD conditions. Overall performance is stable, though slight improvement is seen with the AAAD.

Significant differences of trial vs trigger times are also found. The AAAD was ran in the back-end (but not visible to the observers) in the non-AAAD condition, to compute these measures. Notice that there is a virtual speed-up factor of: $\times 1.5$, in terms of average mean time per trial across observers when using the AAAD.

In addition we performed a related samples t-test (for person t_P and weapon t_W detection) between the hit rates ($M_P = 0.54$, $SD_P = 3.65$, $t_P(17) = 0.497$, $p = 0.626$, two-tailed; $M_W = 2.97$, $SD_W = 12.42$, $t_W(17) = 1.016$, $p = 0.324$, two-tailed), false alarm rates ($M_P = -0.46$, $SD_P = 2.98$, $t_P(17) = -0.655$, $p = 0.521$, two-tailed; $M_W = -1.03$, $SD_W = 6.42$, $t_W(17) = -0.684$, $p = 0.503$, two-tailed), misses ($M_P = -0.54$, $SD_P = 4.65$, $t_P(17) = -0.497$, $p = 0.625$, two-tailed; $M_W = -2.97$, $SD_W = 12.42$, $t_W(17) = -1.016$, $p = 0.323$, two-tailed), and correct rejections ($M_P = 0.46$, $SD_P = 2.98$, $t_P(17) = 0.655$, $p = 0.521$, two-tailed; $M_W = 1.03$, $SD_W = 6.42$, $t_W(17) = 0.684$, $p = 0.503$, two-tailed), and found no significant differences between non-AAAD and AAAD conditions. This last finding is somewhat ideal as the AAAD is intended to either preserve or improve these measures.

Finally, we decided to compare the trigger times of the time, eye movements, detectability, and general satisfaction conditions of the non-AAAD sessions with the AAAD sessions. For comparison, we subtract the final trial time minus the respective trigger time. As such, these times can be thought of as offsets. Note that although the non-AAAD condition does not show any visible assistant to the observer, the PPCs are still being computed in the back-end for the purposes of comparative data analysis. We performed four independent samples t-tests (Time: t_T , Eye Movements: t_{EM} , Detectability: t_D , General: t_G) collapsing all trials across all observers for these times and found significant differences for all trigger case scenarios, supporting the utility of the AAAD: $t_T(1920) =$

-8.46 , $p < 0.0001$, two-tailed; $t_{EM}(1900) = -9.03$, $p < 0.0001$, two-tailed; $t_D(1192) = -11.52$, $p < 0.0001$, two-tailed; $t_G(1190) = -11.56$, $p < 0.0001$, two-tailed.

DISCUSSION

Extensions of model validity beyond current scenarios:

Our results show the potential of a new approach in attention allocation aids that optimizes human search performance by utilizing real time fixational eye movements with prior measurements of target visibility across the visual field and as a function of time. However, there are various potential questions about the generalization of the model across search scenarios.

Our development of the AAAD assumed a target that is present in 50% of the images. A logical question arises as to whether the framework can generalize to real scenarios in which the target is present less frequently. Our model fits herein are performed using signal detection theory metrics (Green and Swets, 1967 [18]) that partition performance into an index of detectability, which is invariant to target prevalence, and a decision criterion, which has an optimal value (maximizing proportion correct) that varies with target prevalence. The model curves, which are utilized to make recommendations to the user, specify proportion correct as a function of time, eye movements, etc. and will vary with target prevalence. However, the model can generalize such curves to varying estimated target prevalence assuming an optimal decision criterion for the given prevalence. Although, we have not tested the generalization experimentally, the theory accommodates such scenarios and generalizations.

For simplicity, our current work used a single target when developing the AAAD. Another natural question to ask is whether the proposed AAAD can still be used if there is the possibility of multiple targets within a given image. Indeed, it could be the case that multiple targets could change how the aid operates within a given application. In some multiple-target scenarios where the detection of *even one* target is sufficient to trigger a decision, our strategy may still apply. For example, for some medical applications, such as screening mammography, finding at least one suspicious target triggers a follow-up diagnostic mammogram. In other applications, localization of each individual targets is important and might require additional development of multiple target model curves to use in conjunction with a prior distribution of the number of targets within the images.

Impact of computer vision developments on proposed


AAAD framework: The recent advances in computer vision might seem to diminish the contributions of the proposed scheme if one assumes that all human search will eventually be replaced by machines. This is yet another reasonable question to ask, since vision has thrived in recent years, in part due to significant advances of Deep Learning [27, 26]. State of the art object recognition algorithms [47, 40, 21] have achieved high performance on certain datasets (MNIST [28], CIFAR [25], ImageNet [42]). However, the images in these datasets typically present ideal scenes with large objects at the center of an image and, currently, the ability of state of the art algorithms to find small or occluded objects in cluttered

scenes (MSCOCO [29]) remains well-below that of humans. Moreover, computers often show glaring errors that humans would not make in what have been called *adversarial* examples [17] in the computer vision community (e.g. by rigging individual pixel values in an image which ‘hacks’ a classifier, a computer can wrongly predict that a white noise-like image is a school bus with 99% confidence [34]). Furthermore, there is still a fundamental lack of understanding with regard to the effects of computer-aided detection aids as a substitute for human observers in many application domains. For example, computer automated detection is prevalent in some countries to flag potential locations for radiologists scrutinizing x-ray mammograms. Yet there is no consensus about its contributions to improving radiologists’ diagnosis accuracy (e.g. Eadie, Taylor, & Gibson, 2012 [14]). As a result of these deficiencies, human observation is still heavily relied upon in a number of applications. As a result, there are many ongoing efforts to reduce errors and optimize human visual search in life-critical tasks from military surveillance [3], to security baggage screening [19], and medical imaging [4].

What is quickly becoming prevalent across many applications is the use of a computer aid that assists humans in localizing potential targets [16]. The proposed AAAD framework does not take into account the presence of a computer aid flagging potential target locations. In some cases, the presence of a computer aid is known to guide search with the risk of leading to over-shortened searches and missed targets that are not flagged by the computer aid [30]. The underlying model in the proposed AAAD allows calculation of an estimated observer accuracy given a pattern of fixations, time and the target visibility across the visual field. In principle, the model could be used to predict if an observer is short-cutting their search (due to the presence of the computer aid) and to alert the observer to further search the image/s. Thus, the developed AAAD framework could be potentially integrated with a computer aid, although its main contribution would likely shift from reducing search times to reducing missed targets.

Potential contribution beyond current application: Although the presented work introduces an AAAD within the context of a very specific task and images, our work serves as a proof of concept for a decision aid design approach that can potentially be applied to a variety of other applications including baggage screening and medical imaging. The model within the AAAD predicts performance on any given trial as a function of time and pattern of fixations and could be potentially used for quantifying the probability that a target was missed on a given image given the observers’ search pattern. Such probabilities of error could be stored with the images and used later to identify images that require secondary inspection by an additional 2nd human observer.

Arguably, the main limitation of the AAAD is that the model relies on empirically measured curves describing the detectability of the target across the visual field and as a function of time. We are currently investigating how to predict target detectability by analyzing image properties such as clutter in real time, which would greatly benefit the application of the model to broader domains.



(a) Sample Person Present and Weapon Present Stimuli.

Figure 7. Sample Stimuli of Experiments 1 and 2. Left: we show random samples of person present, and weapon present in multiple clutter and zoom conditions. Right: we show a magnified version of the (Low Zoom, High Clutter) setting. The box in green has been overlaid on each subimage to reveal the location of the person and weapon when applicable.

CONCLUSION

Our experiments show evidence that our real-time enabled support decision system dubbed AAAD optimizes user efficiency in terms of an increase in the number of trials done as well as a decrease in time spent per each trial, while maintaining performance such as target hit rate and false alarm rate. Thus, the AAAD system has successfully integrated asymptotic performance of search time, eye movements and target detectability. We have described how to fully implement such system through an initial psychometric experiments to find perceptual performance curves for target search, as well as two consequent experiments that verify the benefits of the AAAD. Future computer-human interaction based systems could benefit from implementing AAAD-like systems where having a human-in-the-loop is critical to finding a target even beyond surveillance systems, e.g., medical imaging, astronomical data imagery and remote sensing.

ACKNOWLEDGEMENTS

This work has been sponsored by the U.S. Army Research Office and the Regents of the University of California, through Contract Number W911NF-09-D-0001 for the Institute for Collaborative Biotechnologies, and that the content of the information does not necessarily reflect the position or the policy of the Government or the Regents of the University of California, and no official endorsement should be inferred.

REFERENCES

1. Craig K. Abbey, Frank W. Samuelson, Adam Wunderlich, Lucretiu M. Popescu, Miguel P. Eckstein, and John M. Boone. 2015. Approximate maximum likelihood estimation of scanning observer templates. In *SPIE Medical Imaging*. International Society for Optics and Photonics, 94160O–94160O.
2. Ulf Ahlstrom. 2005. Subjective workload ratings and eye movement activity measures. (2005).
3. Sea Applications Center Air, Land. 2006. FM 3-04.15, Tactics, Techniques, and Procedures for the Tactical Employment of Unmanned Aircraft Systems (UAS). (2006).
4. Kate Alexander. 2010. Reducing error in radiographic interpretation. *The Canadian Veterinary Journal* 51, 5 (2010), 533.
5. Adam T. Biggs and Stephen R. Mitroff. 2015. Improving the efficacy of security screening tasks: A review of visual search challenges and ways to mitigate their adverse effects. *Applied Cognitive Psychology* 29, 1 (2015), 142–148.
6. Katherine L. Bouman, Michael D. Johnson, Daniel Zoran, Vincent L. Fish, Sheperd S. Doleman, and William T. Freeman. 2015. Computational Imaging for VLBI Image Reconstruction. *arXiv preprint arXiv:1512.01413* (2015).
7. Susan E. Brennan, Xin Chen, Christopher A. Dickinson, Mark B. Neider, and Gregory J. Zelinsky. 2008. Coordinating cognition: The costs and benefits of shared gaze during collaborative search. *Cognition* 106, 3 (2008), 1465–1477.
8. Zoya Bylinskii, Adrià Recasens, Ali Borji, Aude Oliva, Antonio Torralba, and Frédo Durand. 2016. Where should saliency models look next?. In *European Conference on Computer Vision*. Springer, 809–824.
9. Mary L. Cummings, Sylvain Bruni, and Paul J. Mitchell. 2010. Human supervisory control challenges in network-centric operations. *Reviews of Human Factors and Ergonomics* 6, 1 (2010), 34–78.
10. Arturo Deza and Miguel P. Eckstein. 2016. Can Peripheral Representations Improve Clutter Metrics on Complex Scenes?. In *Neural Information Processing Systems*.
11. Ivan Diaz, Miguel P. Eckstein, Anaïs Luyet, Pierre Bize, and François O. Bochud. 2012. Measurements of the

- detectability of hepatic hypovascular metastases as a function of retinal eccentricity in CT images. In *SPIE Medical Imaging*. International Society for Optics and Photonics, 83180J–83180J.
12. Birsen Donmez, Patricia E Pina, and ML Cummings. 2009. Evaluation criteria for human-automation performance metrics. In *Performance Evaluation and Benchmarking of Intelligent Systems*. Springer, 21–40.
 13. Trafton Drew, Melissa Le-Hoa Vo, Alex Olwal, Francine Jacobson, Steven E Seltzer, and Jeremy M Wolfe. 2013. Scanners and drillers: Characterizing expert visual search through volumetric images. *Journal of vision* 13, 10 (2013), 3–3.
 14. Leila H Eadie, Paul Taylor, and Adam P Gibson. 2012. A systematic review of computer-assisted diagnosis in diagnostic cancer imaging. *European journal of radiology* 81, 1 (2012), e70–e76.
 15. Miguel P. Eckstein. 2011. Visual search: A retrospective. *Journal of Vision* 11, 5 (2011), 14–14.
 16. Maryellen L Giger. 2014. Medical imaging and computers in the diagnosis of breast cancer. In *SPIE Optical Engineering+ Applications*. International Society for Optics and Photonics, 918908–918908.
 17. Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
 18. DM Green and JA Swets. 1966. Signal detection theory and psychophysics. 1966. *New York* 888 (1966), 889.
 19. Kelly S Hale, Katie Del Giudice, Jesse Flint, Darren P Wilson, Katherine Muse, and Bonnie Kudrick. 2015. Designing, developing, and validating an adaptive visual search training platform. In *International Conference on Augmented Cognition*. Springer, 735–744.
 20. Tim Halverson and Anthony J Hornof. 2011. A computational model of active vision for visual search in human–computer interaction. *Human–Computer Interaction* 26, 4 (2011), 285–314.
 21. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385* (2015).
 22. David C. Klein. 2014. *Using Adaptive Automation to Increase Operator Performance and Decrease Stress in a Satellite Operations Environment*. Ph.D. Dissertation. Colorado Technical University.
 23. Kathryn Koehler, Fei Guo, Sheng Zhang, and Miguel P Eckstein. 2014. What do saliency models predict? *Journal of vision* 14, 3 (2014), 14–14.
 24. Ian Krajbich and Antonio Rangel. 2011. Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. 108, 33 (2011), 13852–13857.
 25. Alex Krizhevsky. 2009. Learning multiple layers of features from tiny images. (2009).
 26. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
 27. Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (2015), 436–444.
 28. Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
 29. Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*. Springer, 740–755.
 30. David Lyell and Enrico Coiera. 2016. Automation bias and verification complexity: a systematic review. *Journal of the American Medical Informatics Association* (2016), ocw105.
 31. Sandra P. Marshall. 2002. The index of cognitive activity: Measuring cognitive workload. In *Human factors and power plants, 2002. proceedings of the 2002 IEEE 7th conference on*. IEEE, 7–5.
 32. David E Meyer and David E Kieras. 1997. A computational theory of executive cognitive processes and multiple-task performance: Part I. Basic mechanisms. *Psychological review* 104, 1 (1997), 3.
 33. Jiri Najemnik and Wilson S. Geisler. 2005. Optimal eye movement strategies in visual search. *Nature* 434, 7031 (2005), 387–391.
 34. Anh Nguyen, Jason Yosinski, and Jeff Clune. 2015. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 427–436.
 35. Aude Oliva, Michael L. Mack, Mochan Shrestha, and Angela Peeper. 2004. Identifying the perceptual dimensions of visual complexity of scenes.
 36. Misha Pavel, Guoping Wang, and Kehai Li. 2003. Augmented cognition: Allocation of attention. In *System Sciences, 2003. Proceedings of the 36th Annual Hawaii International Conference on*. IEEE, 6–pp.
 37. Eli Peli. 1990. Contrast in complex images. *JOSA A* 7, 10 (1990), 2032–2040.
 38. Jeffrey R. Peters and Luca F. Bertuccelli. 2016. Robust Task Scheduling for Multi-Operator Supervisory Control Missions. *AIAA Journal on Aerospace Information Systems* (2016). To Appear.
 39. Jeffrey R. Peters, Vaibhav Srivastava, Grant S. Taylor, Amit Surana, Miguel P. Eckstein, and Francesco Bullo. 2015. Human supervisory control of robotic teams: integrating cognitive modeling with engineering design. *IEEE Control Systems* 35, 6 (2015), 57–80.

40. Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*. 91–99.
41. Ruth Rosenholtz, Yuanzhen Li, and Lisa Nakano. 2007. Measuring visual clutter. *Journal of vision* 7, 2 (2007), 17–17.
42. Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, and others. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
43. Mark W. Scerbo. 2001. Adaptive Automation. In *Neuroergonomics: The Brain At Work*, Raja Parasuraman and Matthew Rizzo (Eds.). 239–252.
44. Thom Shanker and Matt Richtel. 2011. In New Military, Data Overload Can Be Deadly. *The New York Times*. (January 16, 2011).
45. Karthikeyan Shanmuga Vadivel, Thuyen Ngo, Miguel Eckstein, and B.S. Manjunath. 2015. Eye tracking assisted extraction of attentionally important objects from videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3241–3250.
46. Thomas B. Sheridan. 1992. *Telerobotics, Automation, and Human Supervisory Control*. MIT press.
47. Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
48. Yuan-Chi Tseng and Andrew Howes. 2015. The adaptation of visual search to utility, ecology and design. *International Journal of Human-Computer Studies* 80 (2015), 45–55.
49. US Air Force. 2010. *Report on technology horizons, a vision for Air Force Science And Technology during 2010–2030*. Technical Report. AF/ST-TR-10-01-PR, United States Air Force. Retrieved from <http://www.af.mil/shared/media/document/AFD-100727-053.pdf>.
50. François Vachon and Sébastien Tremblay. 2014. What Eye Tracking Can Reveal about Dynamic Decision-Making. In *Int. Conference on Applied Human Factors and Ergonomics*. Kraków, Poland, 3820–3828.
51. Marlies E. van Bochove, Lise Van der Haegen, Wim Notebaert, and Tom Verguts. 2013. Blinking predicts enhanced cognitive control. *Cognitive, Affective, & Behavioral Neuroscience* 13, 2 (2013), 346–354.
52. Carl Vondrick, Donald Patterson, and Deva Ramanan. 2013. Efficiently scaling up crowdsourced video annotation. *International Journal of Computer Vision* 101, 1 (2013), 184–204.
53. Carl Vondrick and Deva Ramanan. 2011. Video Annotation and Tracking with Active Learning. In *NIPS*.
54. Thomas D. Wickens. 2001. *Elementary signal detection theory*. Oxford university press.
55. Jeremy M. Wolfe. 2012. When do I quit? The search termination problem in visual search. In *The influence of attention, learning, and motivation on visual search*. Springer, 183–208.
56. Chen-Ping Yu, Wen-Yu Hua, Dimitris Samaras, and Greg Zelinsky. 2013. Modeling clutter perception using parametric proto-object partitioning. In *Advances in Neural Information Processing Systems*. 118–126.
57. Yanxia Zhang, Ken Pfeuffer, Ming Ki Chong, Jason Alexander, Andreas Bulling, and Hans Gellersen. 2016. Look together: using gaze for assisting co-located collaborative search. *Personal and Ubiquitous Computing* (2016), 1–14.