

Using Highlighting to Train Attentional Expertise

BRETT ROADS AND MICHAEL C. MOZER, University of Colorado
 THOMAS A. BUSEY, Indiana University

Acquiring expertise in complex visual tasks is time consuming. To facilitate the efficient training of novices on where to look in these tasks, we propose an *attentional highlighting* paradigm. Highlighting involves dynamically modulating the saliency of a visual image to guide attention along the fixation path of a domain expert who had previously viewed the same image. We trained novices via attentional highlighting on a fingerprint-matching task. We collected fixation sequences before and after training, and automatically scored each for the degree of expertise via probabilistic generative models of novices and experts. Highlighted training causes gaze behavior to become more expert-like not only on the trained images but also on transfer images, indicating generalization of learning. As a control for the possibility that the increase in expertise was due to mere exposure, we trained subjects via highlighting using novice fixation sequences and observed no transition toward expertise. To determine the nature of the training effect, we trained subjects with fixation sequences from images other than the one being viewed, which preserves coarse statistics of expert gaze, and still observed a transition toward expertise. We discuss possible improvements to the highlighting procedure to increase the specificity of the training effect.

Categories and Subject Descriptors: **H.5.2 [Information Interfaces and Presentation]**: User Interfaces—*Training, help, and documentation*; **H.1.2 [Models and Principles]**: User/Machine Systems—*Human Information Processing*; **I.5.1 [Pattern Recognition]**: Models—*Statistical*

General Terms: Experimentation, Human Factors

Additional Key Words and Phrases: Highlighting, expertise, visual attention, fingerprint matching

ACM Reference Format:

Roads, B., Mozer, M. C., and Busey, T. A. 2015. Using Highlighting to Train Attentional Expertise. *ACM Trans. Appl. Percept.* 20 pages.

DOI=10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

1. INTRODUCTION

Individuals spend a majority of their waking hours performing complex visual tasks and many occupations specifically require operating in challenging visual environments, e.g., monitoring multiple stock exchange status displays, controlling air traffic, screening baggage, examining fingerprints, inspecting medical images, driving trucks. Yet, acquiring visual expertise in any task domain is challenging and time consuming.

Visual expertise can be decomposed into two interacting abilities: *attentional expertise*—knowing where to attend in complex, cluttered scenes—and *procedural expertise*—knowing what to do with the information gathered at the focus of attention. Acquiring these two abilities poses a circular challenge: individuals cannot learn what features and locations in the environment are task relevant until they understand how the information should be integrated and processed, but individuals cannot learn how to process information until they identify relevant feature locations.

Studying the acquisition of visual expertise is quite challenging due to the interaction between the attentional and procedural skills required. Can one type of skill be studied in isolation? It does not make sense to study procedural expertise in the absence of attentional expertise, because attentional expertise provides the visual representations on which classification and judgment procedures operate. However, the reverse is not true: in principle, attentional expertise can be acquired in the absence of procedural expertise, as in, for example, a situation where one is instructed to classify based on color but not told the classification rule. Once novices have learned where and to what to attend, expert-like attention logically supports the acquisition of procedures an expert must perform.

Indeed, past work provides encouraging indications that guiding attention can guide higher-order cognitive processing (Cook, Mitchell, & Goldin-Meadow, 2008; Grant & Spivey, 2003; Mozer et al., submitted; Peterson & Eckstein, 2012, in press; Thomas & Lleras, 2007; Velichkovsky, 1995). For example, Grant and Spivey (2003) noted that certain fixation patterns predict success on the tumor-and-lasers radiation problem, and simply cueing the critical locations increased the probability of success. Thus, although participants had no top-down guidance or goals that steered attention to a critical location, the mere act of attending to the location was sufficient to increase the likelihood of the relevant insight.

Our long-term goal is to develop procedures that improve training of visual expertise. Just as visual expertise might be decomposed into attentional and procedural skills, training of visual expertise might be decomposed similarly. Given the dependencies between attentional and procedural expertise we just discussed, it seems like a sensible first step to focus on the challenge of training attentional expertise.

This research was supported by NSF Awards SBE-0542013 and SMA-1041755 and NIJ Grant #2009-DN-BX-K226. We thank John Lynch and the Center for Research on Consumer Financial Decision Making at the University of Colorado for sharing eye tracking hardware. Corresponding author's address: B. Roads, Dept. of Computer Science, University of Colorado, Boulder, CO 80309-0430. Email: brett.roads@colorado.edu.

Thus, we wish to develop an efficient and relatively effortless means by which novices learn to deploy spatiotemporal attention in a task-appropriate manner. We explore a training paradigm that leverages the knowledge of experts, which we refer to as *attentional highlighting* (Mozer, Pashler, Lindsey, & Jones, submitted) and involves the following steps: (1) Recording gaze dynamics of experts and novices performing a particular task; (2) building a model that predicts locations experts are likely to inspect in specific images and task contexts; and (3) highlighting these locations of interest by manipulating the saliency of the image as a novice performs the task.

This paradigm addresses two challenges that arise when using experts to assist in the training of novices. First, expert knowledge is often procedural and implicit, and experts are limited in their ability to articulate their strategies (Feldon, 2007). Attentional highlighting avoids this issue by analyzing where experts are looking instead of asking experts to verbally report their strategies; expert fixations indicate the locations and features that are important for accomplishing the task. Second, verbal instruction may be unhelpful or even harmful to a trainee because it interferes with the deployment of attention and the natural pace of perceptuomotor behavior while performing a task. With attentional highlighting, verbal instructions are replaced by saliency enhancements, which leverage individuals' rapid and automatic machinery for directing gaze.

This paper explores whether the acquisition of attentional expertise can be facilitated via highlighting—guiding attention by dynamically modulating saliency during training. Because individuals are capable of learning statistical correlations in visual information for statistically structured sequences of objects (e.g. Turk-Browne, Isola, Scholl, & Treat, 2008; Turk-Browne, Jungé, & Scholl, 2005), task-irrelevant perceptual information (e.g. Seitz & Watanabe, 2009), and visuospatial context (e.g. Chun & Jiang, 1998), attending to expert fixation locations may be sufficient to train novices to deploy attention in a task-relevant manner.

Although we aren't aware of past work in which highlighting is used as a training method in visual discrimination tasks, an approach that is similar in spirit has been used to boost performance in problem solving and memory tasks. For example, Litchfield, Ball, Donovan, Manning, and Crawford (2010) showed that novice radiologists benefit from viewing another individual's scanpath. In a collaborative puzzle-solving task involving a novice and an expert, Velichkovsky (1995) observed improved performance when either the novice is cued to the expert partner's fixation sequence, or vice versa. However, highlighting in conjunction with a verbal description of a solution procedure can impair a novice's performance (van Gog, Jarodzka, Scheiter, Gerjets, & Paas, 2009). In the domain of recognition memory, yoking fixations at encoding and test improves performance, although replaying other-observer fixations at test is as effective as replaying same-observer fixations, and scanpath order does not matter (Foulsham & Kingstone, Fixation-dependent memory for natural scenes: An experimental test of scanpath theory, 2013).

As a domain, we focus on the forensic task of comparing a pair of fingerprints. We begin by describing the fingerprint-matching domain, characterizing the saccade sequences of novices and experts, and constructing a model that reliably discriminates saccade sequences of novices and experts. With this model, we can conduct experiments and evaluate the degree to which training via attentional highlighting is effective in transitioning novices toward expert-like attentional control.

2. CHARACTERIZATION OF FINGERPRINT EXAMINERS

Fingerprint analysis is part of the broader field of forensic expertise, which involves the examination of partial or distorted trace evidence left at a crime scene. Figure 1 illustrates a typical fingerprint pair used in casework and adopts the common practice of placing the *latent print* on the left and the *inked print* on the right. Latent prints are obtained from crime scenes and are often distorted, partial, and overlaid on surplus visual information. However, inked prints are made under carefully controlled conditions to get as clean and consistent an impression as possible. Automated classification techniques have been developed to reliably match pairs of inked prints, but due to the

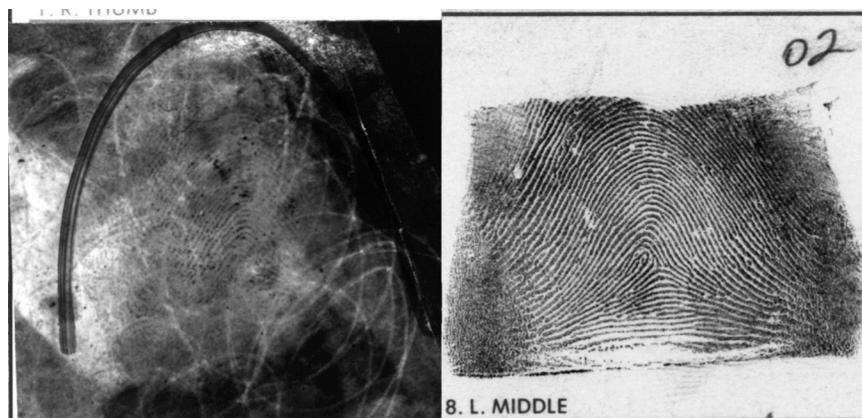


Fig. 1. An example of realistic fingerprint casework taken from the National Institutes of Standards and Technology Special Database 27. This example demonstrates the noisy and partial nature of a latent print (left) compared to the matched inked print (right).

variability and degradation of latent prints, matches involving a latent print and an inked print require expert human judgments.

Fingerprint examination expertise is acquired via a time consuming training period; it can take 1-2 years until a trainee is allowed to carry out unsupervised casework (National Academy of Sciences, 2009). However, following this training period, fingerprint examiners are exceedingly accurate compared to novices (Tangen, Thompson, & McCarthy, 2011; Thompson, Tangen, & McCarthy, 2013). Due to the substantial training required to become a fingerprint expert, even modest gains in training efficiency would be beneficial. Forensic science is well poised to benefit from novel training techniques that exploit research on perceptual expertise and cognitive science in general (Busey & Vanderkolk, 2005; Dror & Cole, 2010; Busey & Dror, 2011).

Although the expert examiner's task will always involve latent prints, in the present work we focus on the task of comparing pairs of inked prints, for two reasons. First, initial training on noise free and complete examples can benefit the learner (Doshier & Lu, 2005). Second, expert scan paths are more consistent than novice scan paths on inked prints (Busey, Yu, Wyatte, Vanderkolk, Parada, & Akavipat, 2011). In contrast, expert scan paths are actually less consistent than novices on latent prints, possibly because the partial and noisy nature of latent prints elicits more idiosyncratic strategies (Busey, Yu, Wyatte, Vanderkolk, Parada, & Akavipat, 2011). Consequently, it is easier to evaluate the expertise of an examiner using pairs of inked prints.

There is broad consensus among fingerprint experts that a fingerprint examination occurs at multiple levels of analysis (Busey & Parada, 2010). At a coarse level of analysis, experts examine the overall ridge flow of the fingerprint. Two diagnostic ridge flow patterns are the core and the delta (Figure 2). While it is possible to reject a match at this coarse level, further analysis is necessary to confirm a match. At a fine level of analysis, experts zoom in to local discriminative features called *minutiae*; the two fundamental minutiae are ridge endings and bifurcations (Figure 2). While there are technically more minutiae types, all other types can be built up from these two fundamental minutiae. At an atomic level of analysis, experts look at the shape or textures of individual ridges. Experts report that they typically rely on the intermediate level of analysis for making a judgment between two fingerprints (Busey & Parada, 2010). Although substantial variation exists among the various computational algorithms for matching fingerprints, many algorithms closely mirror some or all of these levels of analysis (Jain, Feng, & Nandakumar, 2010).

2.1 Eye Movement Data from Novice and Expert Fingerprint Examiners

To explicate differences between expert and novice fingerprint analysts we reanalyzed a dataset collected by Busey et al. (2013, dataset 2). The dataset consists of 26 matching inked fingerprint pairs and two mismatching pairs. (Matching pairs are a preferred and more informative source of data because 'same' judgments are typically slower and more deliberate than 'different' judgments.) Busey et al. (2013) collected fixation data from 12 experts and 12 novices performing the matching task. In this study, experts were recruited at forensic identification conferences and laboratories, and novices were recruited from the Bloomington, Indiana community. Experts in the dataset had a reported average of 7.6 years of unsupervised latent print work (minimum 3 years, maximum 13 years). On each trial, subjects viewed a pair of fingerprints side by side, which we'll refer to as the left and the right prints. Each trial was divided into three parts: for 5 s, the left print was presented alone, then for 5 s, the right print was presented alone, and finally for 10 s, both prints were presented simultaneously. After each trial, subjects were asked to indicate whether they believed the two prints came from the same source, different sources or 'unable to tell'.

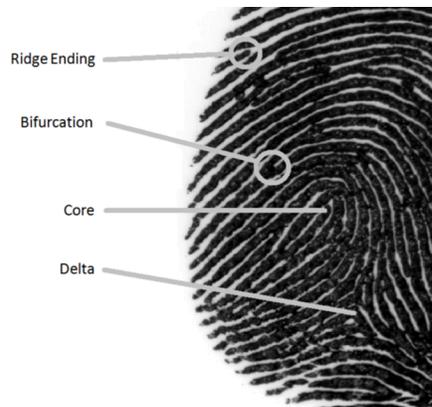


Fig. 2. Close up of an inked print from the dataset used by Busey et al. (2013) and this work. The figure shows the location of the core and delta used in level one analysis and examples of the two fundamental minutiae—ridge endings and bifurcations—that are used in level two analysis.

Statistics of the training dataset provide a window into similarities and differences between expert and novice fingerprint examiners. As reported by Busey et al. (2013), there is no statistically significant difference in the total number of fixations between experts (19.0) and novices (18.1) or in the mean duration of each fixation (experts, 176.2 ms; novices, 174.3 ms). However, Busey et al. (2013) observed that experts have smaller saccade amplitudes on both the left and right print and experts make more saccades within prints than across prints than novices. Experts also have slightly more fixations on the left print than the right, which is consistent with the standard procedure for latent print examinations. Experts' consistency largely derives from focusing the examination on the core and delta regions of the inked fingerprint pair (Busey & Parada, 2010), but variability among experts is considerable (Dror et al., 2011). A prototypical scanpath of an expert and a novice are shown in Figure 3.

Although it is possible to decouple attention from fixations (Posner, 1980), when performing complex visual tasks it is more efficient to move our eyes than to covertly shift attention (Peiyuan & Kowler, 1992; Schlingensiepen, Campbell, Legge, & Walker, 1986) and resolution is highest at the fovea. Fixations thus provide a good indication of the image regions from which information is gathered to perform a visual task, and fixation sequences provide information about the underlying cognitive operations used to compare the prints.

2.2 Discriminative Modeling of the Degree of Expertise

In the experiments we describe in this article, participants are shown fingerprint images and are instructed to inspect the images in order to determine whether or not the pair of fingerprints match. The dependent measure in the experiments is the degree of attentional expertise exhibited by a participant before and after training. In the Appendix, we explain how we compute an attentional expertise score from a participant's fixation sequence and simple probabilistic models.

Briefly, the models assign a likelihood to a fixation sequence \mathbf{F} given image \mathbf{I} for a population s (novices or experts), denoted $p(\mathbf{F}|s, \mathbf{I})$. If novices and experts systematically differ, a model of experts will assign a higher likelihood to expert sequences than to novice sequences, and a model of novices will do the opposite. By contrasting the predictions of the models, one can thus discriminate novice and expert attentional behavior. Specifically, Bayes' rule can be used to classify a fixation sequence as having been produced by a novice or by an expert:

$$p(s|\mathbf{F}, \mathbf{I}) = \frac{p(\mathbf{F}|s, \mathbf{I})p(s)}{\sum_{s'} p(\mathbf{F}|s', \mathbf{I})p(s')} \quad (1)$$

We use a measure related to Equation 1, the *log likelihood ratio (LLR)*, to characterize the degree of attentional expertise reflected in a subject's fixation sequence.

The novice and expert likelihood models used in Equation 1, $p(\mathbf{F}|s=\text{novice}, \mathbf{I})$ and $p(\mathbf{F}|s=\text{expert}, \mathbf{I})$, describe the generative process underlying fixation sequences. However, the models should not be thought of as making theoretical claims about underlying cognitive processes. Instead, the models are simply being used as a black-box tool for discriminating novice and expert fixation sequences. In the Appendix, we describe a family of models and, via previously collected data (Busey et al., 2013), identify the specific model that achieves the best inter-group discrimination.

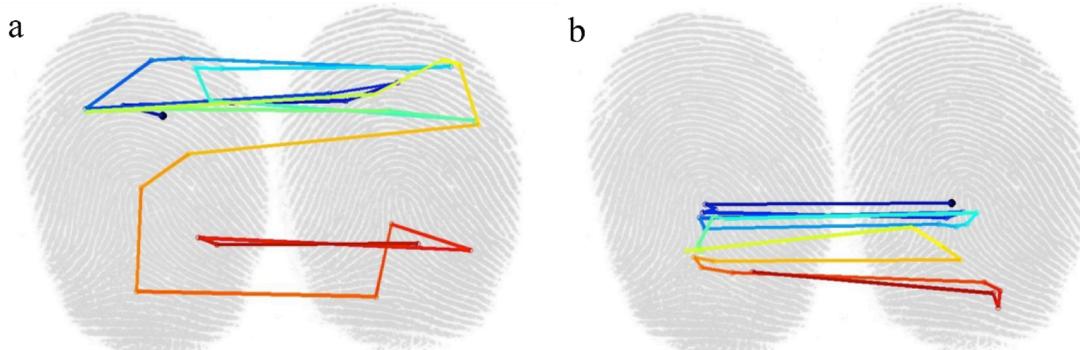


Fig. 2. A prototypical scan path of a novice (a) and an expert (b). The scan path begins at the black dot on the end of the dark blue segment. The first fixation is located at the black dot on the right and the color scheme shows the expert's fixation sequence as it progresses through time going from dark blue to yellow to dark red) This expert demonstrates a typical scanpath that concentrates on the core and delta regions of the fingerprint, diagnostic of level one analysis. In contrast, the novice exhibits less directed behavior. This figure was generated using MATLAB code modified from code provided by Busey.

3. EXPERIMENTAL STUDIES OF ATTENTIONAL HIGHLIGHTING

Assuming models exist that can be used to assess the attentional expertise exhibited by subjects, we now address the main goal of our work: to develop and evaluate training procedures that yield a rapid and relatively effortless shift toward expert-like behavior by highlighting locations where an expert is likely to attend. We describe three experiments that explore attentional highlighting to train novice fingerprint examiners.

In Experiment 1, naive subjects were trained on fingerprint pairs to replicate the fixation sequence of one randomly selected expert from Busey et al.'s dataset.¹ Each training trial began with presentation of a fingerprint pair, followed by a cue to the expert's first fixation location within the image. The cue consisted of a blinking red spot in the image background (Figure 4). When the subject made a saccade to the cue, as registered by an eye tracker, the cue shifted to the expert's next fixation, and so forth. From the subject's perspective, the training task involved following a sequence of red spots, sometimes jumping within a fingerprint and sometimes jumping to the other fingerprint. The training procedure preserves the order of fixations because there is potential information in the sequence itself (e.g., Noton & Stark, 1971), which may support learning. Experiments 2 and 3 involved a similar training procedure, but in Experiment 2, subjects were trained to follow the fixation sequence of a naive viewer, and in Experiment 3, subjects were trained to follow the fixation sequence that an expert made to a different stimulus than the current one. The rationales for these experiments will be explained shortly.

All three experiments included a pre-training and a post-training phase during which subjects performed free viewing of images. A comparison of fixation patterns pre- and post-training was used to assess the influence of training. During free viewing, the subjects' task was to perform as a fingerprint analyst and examine the pair of prints in order to determine similarities and differences between them.

3.1 Experiment 1: Expert Highlighting

3.1.1 Participants. Twelve subjects with normal or corrected to normal vision drawn from a paid subject pool at the University of Colorado, Boulder. Subjects were given \$10 in compensation for their time. One subject was dropped because the eye tracker failed to reliably detect gaze shifts. The dropped subject was rerun.

3.1.2 Apparatus and Materials. The experiment utilized a Tobii T60 XL eye tracker (24-inch 1920 x 1200 pixels widescreen monitor), an open source Talk2Tobii extension (Deligianni, 2014), the Psychophysics Toolbox extension (Brainard, 1997), and MATLAB all running on an iMac computer. The 28 stimuli (each 1680x1050 pixels) consisted of two fingerprints side by side in 256 bit gray scale.² Stimuli were displayed pixel perfect on the eye tracker screen and subtended approximately 39.2 by 25.1 degrees of visual angle. The remainder of the screen surrounding the stimulus was displayed in black. The two fingerprints in each stimulus cannot be distinguished easily. Among the 28 possible stimuli, only two stimuli contained mismatching fingerprint pairs and not every subject was assigned a mismatching case.

3.1.3 Procedure. Prior to the beginning of the experiment, subjects were calibrated to maximize performance of the eye tracker. Subjects were seated so that their eyes were approximately 25 inches (63.5 cm) from eye tracker screen, as recommended by Tobii documentation, and the eye tracker's height was adjusted so that the subject's eyes were approximately level with the center of the eye tracker screen. Prior to any experiment-specific instructions, all



Fig. 3. An example of the attentional highlighting used in Experiments 1-3. Each fixation location in a selected sequence shows attentional highlighting using a flashing red Gaussian intensity bump.

1. We chose to train each participant on a single expert to allow for the possibility of expert-specific idiosyncrasies (Dror et al., 2011).
2. Two stimuli were thrown out because some training data were corrupted on those two stimuli.

subjects were verbally informed that blinking during the experiment was acceptable and that slight movement was permitted. The eye tracker was calibrated for each subject using a script that was slightly modified from a publicly available script that accompanied the Talk2Tobii software. Halfway through the experiment the eye tracker was calibrated again to account for drift.

At the initiation of the experiment, an automated program randomly assigned each subject 16 stimuli from 28 possible stimuli in the training dataset. These 16 stimuli made up the subject's *stimulus set*. Eight of the 16 stimuli in the stimulus set were randomly assigned to the *training set* and the remaining eight were assigned to the *transfer set*. In addition to being randomly assigned stimuli, each of the 12 subjects was randomly paired with one of the 12 experts in the training dataset, in one-to-one correspondence. All attentional highlighting provided to a subject during the training phase utilized fixations from their matched expert and a given stimulus always showed the same fixation sequence.

Following calibration, subjects received on-screen instructions which asked them to imagine that it was their first day on the job as a fingerprint analyst and they should look for similarities and differences between pairs of fingerprints displayed on the screen. After subjects indicated they were ready to continue, they were immediately presented with the on-screen instructions for the pre-training phase, informing them that they would begin by examining fingerprints without assistance. (The overall trial sequence is shown in Figure 5.)

The pre-training phase consisted 16 trials each involving free viewing of a stimulus for 10 s. On each trial, one of the 16 stimuli in the subject's stimulus set was displayed, in randomized order. During the 10 s viewing period, the subject's gaze was recorded. Prior to each trial, a white fixation cross was presented on a black background for 1 s. This pre-training phase was designed to mimic the simultaneous exposure period in Busey et al. (2013) where two inked fingerprints were shown side by side. In the pre-training phase and the remainder of the experiment, subject responses were not collected because the focus of the study is on eye gaze behavior rather than on decision making. Even without overt judgments, we were confident that the prints were carefully examined by virtue of the number of saccades produced on each trial, and the observation that the vast majority of saccades landed on one of the two prints.

Following the pre-training phase, subjects were trained using attentional highlighting over 64 trials, divided into eight blocks each consisting of the eight stimuli in the training set, randomized within a block. After blocks 2, 4, and 6, subjects were informed of their progress, reminded of their task and given the opportunity to take a quick break. Prior to the training phase, subjects were presented with on-screen instructions explaining that they would be examining pairs of fingerprints aided by attentional highlighting. They were told that they should look at the locations of a series of cues and that the trial would end when all of the cues in the sequence had been visited.

Each trial began with presentation of a white fixation cross on a black background. After 1 s, the cross was replaced by a pair of fingerprints centered on the cross location. Gaze-contingent attentional highlighting sequentially

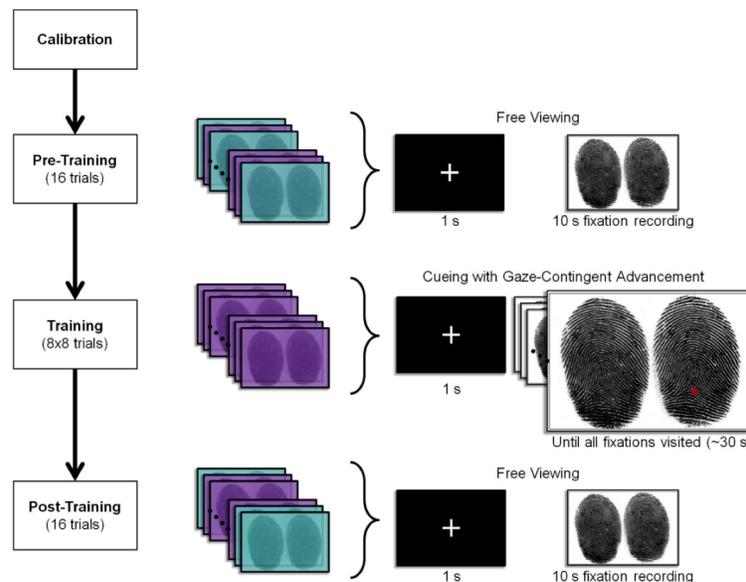


Fig. 4. The four phases of Experiments 1-3. During the pre-test and post-test phases, subjects view images for 10 s and their fixations are recorded as they examine the prints. During the training phase, attentional highlighting guides fixations to locations of interest (red dot). Once subjects saccade to the location of the cue, the next fixation in the sequence is cued. Each trial begins with a fixation cue for 1 s preceding the onset of the fingerprint image. The fingerprint pairs shaded in purple represent the trained stimulus set, while the cyan fingerprint pairs represent the transfer stimulus set.

cued subjects to the locations of their assigned fixation sequence. The particular fixation sequence that each subject saw was determined by the expert they were assigned at the beginning of the experiment. For a given subject, all training trials used fixation sequences from the same expert viewing both images simultaneously. The onset of the first highlight (i.e., the first fixation in the expert sequence) coincided with the stimulus onset. Once the subject made a correct saccade to the cued location, the highlight was immediately removed and after 180 ms, a highlight appeared at the next fixation in the sequence.³ The highlighted training proceeded until all fixations in the sequence had been visited in the correct order. The trial ended when the last fixation in the sequence was visited. By following the cues, the subjects traced out their assigned expert's scanpath.

Each cue had a Gaussian-distributed intensity and filled only the white background of the image and did not obstruct the black fingerprint ridges. The center of the cue is specified by the expert's fixation, and is fully saturated red. The cue has standard deviation of 0.25° of visual angle (10 pixels) and flashed once per second, with an onset-to-offset interval of 300 ms, to facilitate its detection. Onsets and offsets are strong attractors of attention regardless of the task (Ludwig & Gilchrist, 2002; Theeuwes, Kramer, Hahn, & Irwin, 1998).

Following the training phase, subjects were presented with on-screen instructions informing them that they would be comparing fingerprints without assistance (as they had done during the pre-training phase). The post-training phase consisted of a recorded 10 s free gaze period identical to the pre-training period. The stimulus order in the post-training was independently randomized from the pre-training stimulus order. The entire experiment took about 45 minutes.

In each phase of the experiment, each subject's gaze was recorded by the Tobii T60 XL at the rate of 60 Hz. Along with the eye coordinates at the time of sampling, the Tobii records a code for each eye indicating whether or not the eye has been detected. The data were cleaned by removing invalid gaze points. Although the Tobii is able to repair samples that have only one invalid eye, samples that have two invalid eyes must be thrown out. Any trials which have over 50% invalid gaze samples are thrown out since this suggests a systematic problem with the trial. Following the removal of invalid gaze points, the location of the gaze points is checked to make sure that all gaze points fall on the stimulus. Since the stimuli do not fill the entire screen of the eye tracker, it is possible for subjects to look off the stimulus. Any samples that are off the stimulus are removed and trials that have more than 50% of the gaze points off the stimulus are thrown out because this suggests a systematic inattention to the task.

Since the eye-tracker collects data at a uniform rate and does not directly determine fixations and saccades, it is necessary to run the gaze data through a fixation filter to extract fixation locations. The gaze filter used follows common parameters used in other research (Busey, Yu, Wyatte, & Vanderkolk, 2013; Sen & Megaw, 1984). Specifically, a median filter with a window of three consecutive points is run over the data, which serves to reduce noise in the data. The magnitude of the velocity is then calculated at every point in the smoothed data and a velocity threshold is used to segment successive gaze points into fixations. The velocity threshold used here was $10^\circ s^{-1}$. A minimum of 67 ms is established for a set of gaze points to be grouped as a fixation in order to eliminate extremely brief fixations. All subjects but one had 32 valid trials. The remaining subject produced eight trials that were over 50% invalid and was replaced by a new subject who produced 32 valid trials.

3.1.4 Results. The experiment involved four conditions in a 2-by-2 design: the time of recording (before vs. after training) crossed with whether the stimuli were used for training or not (training vs. transfer stimuli). We performed three different analyses on the cleaned data: (1) an empirically based analysis that measures the spatial proximity of subject fixations to expert fixations, (2) a model-based analysis that evaluates the degree of expertise via expert and novice image-specific models, and (3) a summary analysis of gaze statistics. We discuss each in turn.

3.1.4.1 Empirical analysis of spatial proximity to expert fixations. For each fixation of each subject on a particular image, we determined the mean Euclidean distance to all expert fixations on the same fingerprint impression. This mean distance was used to form an empirical cumulative distribution function (CDF). By comparing the CDFs before and after training, we can assess the degree to which subjects' fixations better approximated the experts'. The top row of Figure 6 shows the empirical CDF in the four conditions of Experiment 1. A two-sample Kolmogorov-Smirnov test (K-S test) allows us to test whether the probability distributions before and after training differ. If we treat each fixation as independent of the others, the K-S test produces highly significant effects for both the training stimuli ($D(2025, 1913) = 0.114, p < .001$) and the transfer stimuli ($D(2049, 1825) = 0.083, p < .001$). However, such a treatment doesn't consider the dependence among fixations in a sequence and therefore inflates the apparent degrees of freedom in the data. Consequently, we propose an approach to estimating the reliability of

3. If the subject did not saccade to the cue after 3 s, the next cue was presented. While this rarely happened, a cue timeout guaranteed the experiment would not extend beyond the allotted time. A saccade to a cue might not be registered because the subject could not locate it or the eye tracker could not get an appropriate reading.

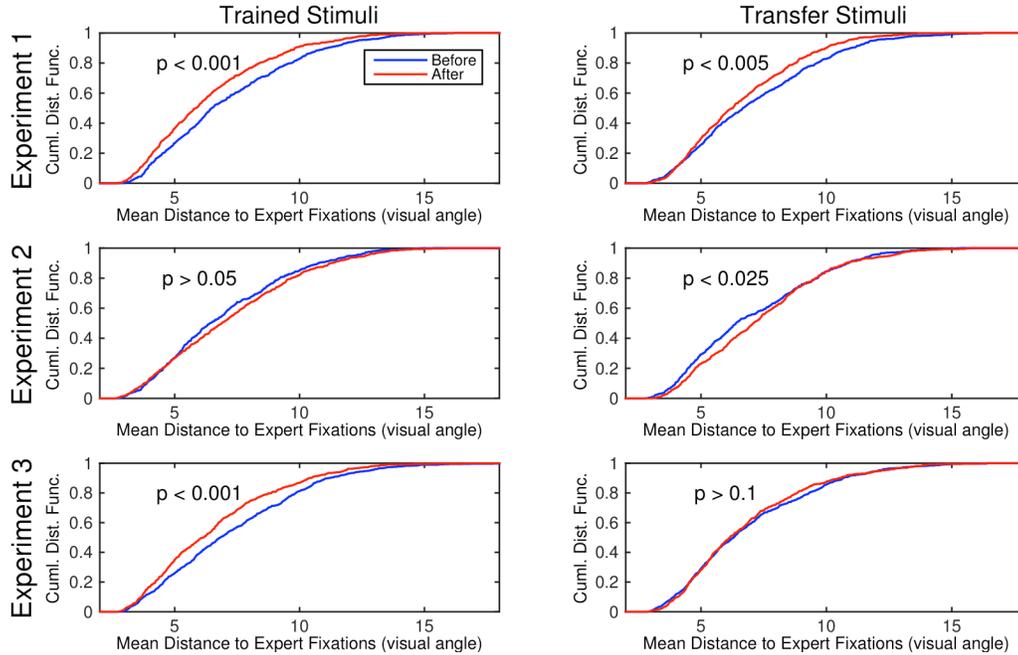


Fig. 5. Summary of spatial proximity results from Experiments 1-3. Each experiment is shown in a row of the figure. The left and right columns present results for trained and transfer stimuli, respectively. Each graph shows the empirical cumulative distribution function before (blue) and after (red) the training phase. Each sample in the cumulative distribution function derives from the mean distance to all expert fixations (one the same image and impression) for a given subject fixation. Experiment 1-3 involved highlighted training based on expert fixations, novice fixations, and incongruent expert fixations, respectively.

differences in the distributions by excluding adjacent fixations, performing the K-S test with only every k -th fixation in a sequence—i.e., fixations 1, $k+1$, $2k+1$. With a suitable k , the fixations will be independent. By intuition, we chose $k = 3$, which yields a reliable shift toward expertise for both training and transfer stimuli ($D(706,669) = 0.123$, $p < .001$; $D(718,640) = 0.096$, $p < .005$). Significance levels in Figure 6 are based on $k = 3$. Because of the arbitrariness in the choice of k , we performed a second, model-based analysis of the data.

3.1.4.2 Model-based analysis. The data were evaluated via the expert and novice image-specific models to obtain a log-likelihood ratio (the LLR, defined in Equation A5). The LLR provides a measure of *relative expertise*, allowing us to compare the nature of eye movements before and after training. The top row of Figure 7 shows relative expertise across subjects in the four conditions of Experiment 1. For both trained and transfer stimuli (left and right graphs, respectively), highlighted training shifts performance toward expertise. We conducted an analysis of variance (ANOVA) with subjects as the random variable and two within-subject factors: experiment phase (before vs. after training) and stimulus type (training, transfer). The ANOVA obtained a main effect of phase ($F(1, 11) = 24.92$, $p < .001$), no reliable effect of stimulus type ($F(1, 11) = 0.05$), and no phase-stimulus type interaction ($F(1, 11) = 0.16$). Thus, highlighted training was effective not only in altering fixation patterns on the specific images trained, but on novel images as well.

3.1.4.3 Summary gaze statistics. The data were evaluated to obtain low-level gaze statistics of various measures before and after training. This analysis included proportion of fixations on the left (vs. the right), proportion of saccades within impressions (vs. between impressions), mean fixation duration, mean saccade amplitude on the left and mean saccade amplitude on the right. Data were collapsed across training and transfer stimuli to reduce variance of the statistics. The top section of Table 1 shows the analysis of these variables for Experiment 1. There was a significant increase in the percentage of saccades that occurred within impressions; no other effect was significant.

3.1.5 Discussion. The results of Experiment 1 indicate that training via attentional highlighting leads to gaze patterns that are less like those of domain novices and more like those of domain experts. This shift toward expertise is observed both for stimuli that were used in training and for novel stimuli. Had highlighting benefited only the training stimuli, one might have explained the effect in terms of memorization of previous scanpaths (Valuch, Becker,

Table 1. Low-level statistics of subject gaze behavior for Experiments 1, 2 and 3. For each experiment, an analysis was performed on the before- and after-training data for the following statistics: the percentage of fixations that occur on the left impression versus the right (*Left Fixations*), the percentage of saccades that occur within an impression versus across impressions (*Within Saccades*), the mean fixation duration (*Fixation Duration*), the amplitude of saccades that occur on the left and right impression (*Saccade Amplitude*). Experiment 1-3 involved highlighted training based on expert fixations, novice fixations, and incongruent expert fixations, respectively.

Characteristic		Before		After		Analysis			
		%		%		X ²	df	N	p
Exp. 1	Left Fixations	57.2		55.1		3.78	1	7812	.05
	Within Saccades	68.3		71.3		7.62	1	7442	.006
		Mean	SD	Mean	SD	t	df		p
	Fixation Duration (ms)	133.3	17.5	132.9	17.9	0.15	11		.880
	Saccade Amplitude (°)								
	Left	3.04	0.55	2.92	0.57	0.59	11		.570
	Right	1.98	0.56	2.05	0.42	-0.28	11		.788
Exp. 2	Left Fixations	52.6		52.2		0.15	1	6883	.703
	Within Saccades	67.1		68.3		1.22	1	6509	.269
		Mean	SD	Mean	SD	t	df		p
	Fixation Duration (ms)	140.7	30.3	145.7	28.9	-0.71	11		.492
	Saccade Amplitude (°)								
	Left	2.94	0.61	3.08	0.82	-0.72	11		.485
	Right	2.69	0.94	2.51	0.61	0.71	11		.490
Exp. 3	Left Fixations	51.9		52.8		0.70	1	8833	.70
	Within Saccades	68.7		73.4		22.57	1	8461	<.001
		Mean	SD	Mean	SD	t	df		p
	Fixation Duration (ms)	158.0	23.4	147.8	37.6	1.33	11		.211
	Saccade Amplitude (°)								
	Left	2.92	0.55	2.99	0.75	-0.28	11		.782
	Right	2.41	0.38	2.39	0.72	0.14	11		.891

& Ansorge, 2013). However, both the empirical analysis and the model-based analysis indicate that generalization is quite robust, making it unlikely that subjects were simply memorizing the instructed fixation sequences.

The analysis of low-level gaze statistics provides some further support for the value of highlighting in facilitating the acquisition of expertise, but also points out some limitations. The increase in the percentage of saccades that occur within an impression is in agreement with the observations of Busey et al. (2013) that experts make a higher percentage of saccades within impressions than do novices. It is possible that this reflects adoption of a more efficient examination strategy that seeks to compare an arrangement of features rather than compare features one at a time. In contrast, our analysis suggests that highlighting does not increase the percentage of fixations on the left or decrease the amplitude of saccades within an impression. Although Busey et al. (2013) did find that experts fixated significantly more often on the left impression, the difference between experts and novices in their corpus was small (52.9% versus 51.4% respectively), and the reason for this difference with real-world experts is that examination strategies may focus on the latent print, which is conventionally placed on the left (Busey et al., 2013). In our experiment, no latent print was used.

Thus far, we have focused on gaze statistics and haven't reported measures of the primary discrimination task that subjects performed—comparing a pair of prints to determine if they came from the same individual. As we mentioned previously, 26 of the 28 images used in Experiment 1 contained matching prints. The predominance of matching prints ensures extended gaze sequences: with inked prints, the clarity of the images makes it easy to reject mismatching prints. Given a relatively small data set and the imbalance of matches and mismatches, we did not expect any measure of discriminability or accuracy to provide reliable indications of learning, and thus did not collect same/different judgments.

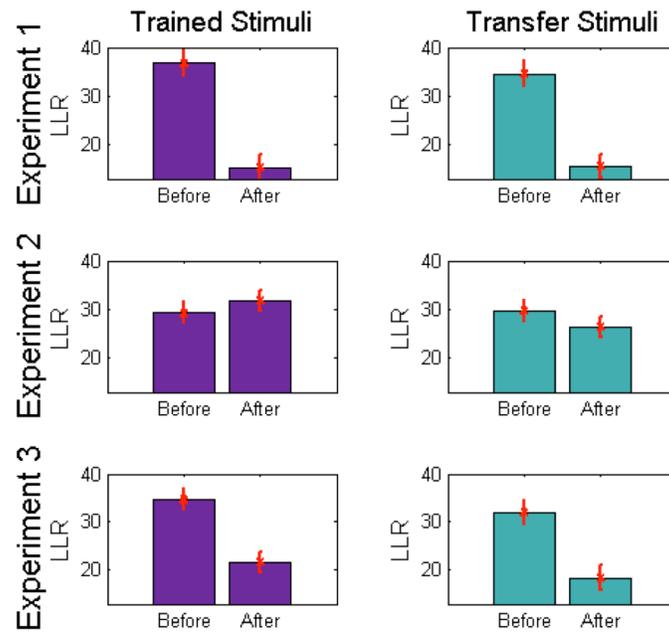


Fig. 6. Summary of results from Experiments 1-3. Each experiment is shown in a row of the figure, and the left column presents results for trained stimuli and the right column for transfer stimuli. The ordinate of each graph indicates the log likelihood ratio (LLR), a measure of *relative expertise*. An increase in the LLR from before to after training indicates acquisition of expertise. Experiment 1-3 involved highlighted training based on expert fixations, novice fixations, and incongruent expert fixations, respectively. The error bars are calculated following Masson and Lofus (2003) to account for systematic tendencies of data derived from the same subject and variability between subjects.

One concern with Experiment 1 is that subjects might not be learning from the highlighting per se; rather, their improved performance might simply be due to mere exposure to the fingerprint stimuli and the resulting increased familiarity with fingerprints in general. To test this hypothesis, Experiment 2 was identical to Experiment 1, except that instead of highlighted training based on expert fixation sequences, the training procedure was based on novice fixation sequences. Each subject in Experiment 2 was matched with a novice in the training dataset, and the training novice provided the fixation sequences for highlighting. If highlighting has no effect, then Experiment 2 should produce the same results as Experiment 1. However, if subjects are learning from the highlighted sequence, the shift towards expertise should not be observed in Experiment 2.

3.2 Experiment 2: Novice Highlighting

3.2.1 Participants. Twelve subjects were recruited from a paid subject pool at the University of Colorado, Boulder. All subjects had normal or corrected to normal vision and were given \$10 in compensation for their time.

3.2.2 Apparatus. The eye tracker setup and stimuli used were identical to Experiment 1.

3.2.3 Procedure. Experiment 2 was identical to Experiment 1, except that highlighted training was based on fixation sequences from novices not experts. The training and transfer stimuli used for the 12 subjects in Experiment 2 were matched to the training and transfer stimuli from the 12 subjects in Experiment 1, facilitating a comparison between the outcomes of the two experiments by removing one potential source of variation that differentiates the experiments. Each of the 12 subjects was paired in one-to-one correspondence with a random novice from the training dataset.

3.2.4 Results. Data collected in Experiment 2 were cleaned and filtered for fixations in the same fashion as Experiment 1. One subject had 28 valid trials and one subject had 30 valid trials, but all remaining subjects had 32 valid trials. In the case where a trial is dropped, it is necessary to drop the corresponding trial in the pre-training or post-training period to balance the analysis.

3.2.4.1 Empirical analysis of spatial proximity to expert fixations. The empirical CDF of the mean distance to expert fixations before and after training on both training and transfer stimuli is presented in the second row of Figure 6. Evaluating the difference in distributions with a K-S test that uses every third fixation, as we did in Experiment 1, we find no reliable training effect on the training stimuli ($D(599,663) = 0.069, p > .05$), and the effect on the test

stimuli ($D(564,602) = 0.093, p < .025$) is reliable, except in the opposite direction from Experiment 1: highlighted training using novice fixation sequences causes subjects to produce more novice-like fixations.

3.2.4.2 Model-based analysis. The relative expertise before and after training on both the training and transfer stimuli is presented in the second row of Figure 7. We conducted an ANOVA with subjects as the random variable and two within-subject factors: experiment phase (before vs. after training) and stimulus type (training, transfer). The ANOVA obtained no reliable effects of phase ($F(1, 11) = 0.02$) or stimulus type ($F(1, 11) = 0.22$). The phase-stimulus type interaction ($F(1, 11) = 0.71$) was also not significant.

An ANOVA was also conducted with the combined data from Experiments 1 and 2 to compare the effect of training with highlighting derived from expert versus novice fixation sequences. The ANOVA included highlighting source (expert versus novice) as a between subjects factor, and phase of evaluation (before vs. after) and stimulus type (training vs. transfer) as within subject factors. The ANOVA obtained a significant effect of phase ($F(1, 22) = 18.66, p < .001$) and an interaction of phase and source ($F(1, 22) = 17.36, p < .001$). Main effects and interactions involving stimulus type were not significant. Comparing the outcomes of Experiments 1 and 2 in Figure 7, one can see that the interaction is due to a training effect tending in opposite directions in the two experiments. Training on expert fixation sequences yields more expert-like behavior, and training on novice fixation sequences yields either no change or more novice-like behavior.

3.2.4.3 Summary gaze statistics. The analysis of low-level gaze statistics for Experiment 2 is presented in the second section of Table 3. None of the low-level gaze characteristics changed significantly as a result of highlighted training. In particular, training did not produce a significant increase in the number of saccades that occur within an impression (versus across impressions), in contrast to the effect of training in Experiment 1.

We note one caveat concerning the analysis that combines Experiments 1 and 2. The experiments were conducted consecutively in time, and thus subjects were not randomly assigned to be trained by expert vs. novice highlighting. However, the temporal lag between Experiments 1-3 was brief and a paid subject pool was used, and thus we are not concerned with systematic differences in the subject population across experiments.

3.2.5 Discussion. The results of Experiment 2 indicate that training by highlighting novice fixation sequences does not yield more expert-like gaze sequences. Consequently, the training effect observed in Experiment 1 cannot be explained by mere exposure to and familiarity with the fingerprint stimuli. If the source of the training effect was exposure and familiarity and not the highlighting procedure, Experiments 1 and 2 would have produced identical results.

Although it's conceivable that exposure might have some effect without any form of highlighting, Experiments 1 and 2 establish that subjects learn different fixation strategies as a result of expert-highlighted training versus novice-highlighted training. Nonetheless, we know little about what specifically is being learned. In our earlier characterization of expert fingerprint examiners, we discussed two spatial resolutions at which expertise is manifested. At the global scale, experts concentrate gaze on the center region of each fingerprint where the core and delta determine the gestalt flow of the ridges. At the fine scale, experts are concerned with local configurations of visual features, such as the minutiae. Highlighting might train the attentional control system to be responsive to either or both of these constraints.

To discriminate among the possibilities, Experiment 3 was designed such that highlighted training provided guidance on only the coarse scale of fingerprint analysis and not on the fine scale, in contrast to highlighted training in Experiment 1 which provided guidance at both scales. If the training effect is reduced or eliminated in Experiment 3, we obtain evidence that highlighting facilitates the learning of task-relevant, spatially local visual features of the domain.

To eliminate fine-scale information in Experiment 3, highlighted training on a stimulus X was based on the fixation sequence of an expert for a different stimulus Y . This *incongruent* expert highlighting avoids cueing the locations of minutiae, which are fingerprint and image specific. However, it does cue coarse features of the fingerprints, for the following reason. Most fingerprints have a core and delta that occupy the center region of the fingerprint, and stimulus images are normalized on the screen; consequently, the locations of the cores and deltas for pair of stimuli, X and Y , are typically aligned.

3.3 Experiment 3: Incongruent Expert Highlighting

3.3.1 Participants. Twelve subjects were recruited from a paid subject pool at the University of Colorado, Boulder. All subjects had normal or corrected to normal vision and were given \$10 in compensation for their time.

3.3.2 *Apparatus and Materials.* The eye tracker setup and stimuli used were identical to Experiment 1.

3.3.3 *Procedure.* The design of Experiment 3 was identical to that of Experiment 1. Each subject was paired with an expert whose data provided the highlighting sequence. To match Experiment 1, Experiment 3 used the same subset of stimulus items for each expert, they were divided into training and transfer sets using the same split as Experiment 1, and the trial sequences of subjects in Experiment 3 matched those of subjects in Experiment 1. The key difference from Experiment 1 is that in Experiment 3, for each stimulus X , subjects were trained not on the expert fixation sequence corresponding to X but on the sequence corresponding to a different stimulus, Y . The re-pairing of expert fixations and stimuli was performed by permuting the set of stimuli such that each training stimulus was associated with the fixation sequence of a different training stimulus. Although one might ordinarily perform a permutation at random, we wanted to ensure that the reassigned fixation sequences were not by chance a good match to the incongruent stimulus. Consequently, the fixation sequences were permuted so as to minimize the likelihood of the incongruent fixation sequence under a model based on the congruent fixation sequence of all experts.

3.3.4 *Results.* Data collected in Experiment 3 were cleaned and filtered for fixations in the same manner as Experiment 1. Only one subject had invalid trials (2 of 32 trials); therefore, no subject was rejected.

The bottom row of Figure 6 shows the empirical CDF in the four conditions of Experiment 1. The K-S test using every third fixation obtained a significant training effect for the training stimuli ($D(875, 702) = 0.126, p < .001$) but not for the transfer stimuli ($D(843, 643) = 0.048, p > .10$), suggesting that incongruent highlighted training is less effective than congruent training (Experiment 1).

Moving on to the model-based analysis, the relative expertise before and after training is presented in the third row of Figure 7. We conducted an ANOVA with subjects as the random variable and two within-subject factors: experiment phase (before vs. after training) and stimulus type (training, transfer). The ANOVA obtained a main effect of phase ($F(1, 11) = 11.32, p = .006$), no reliable effect of stimulus type ($F(1, 11) = 1.02$), and no phase-stimulus type interaction ($F(1, 11) = 0.02$).

An ANOVA was conducted with the combined data from Experiments 1 and 3 to compare the effect of highlighting training based on congruent versus incongruent expert fixation sequences.⁵ The ANOVA included congruency of highlighting (congruent in Experiment 1, incongruent in Experiment 3) as a between-subjects factor, and phase of evaluation (before vs. after) and stimulus type (training vs. transfer) as within-subject factors. The ANOVA obtained a significant effect of phase ($F(1, 22) = 35.13, p < .001$) and no significant interaction of phase and congruency ($F(1, 22) = 1.57, p = .22$). Main effects and interactions involving stimulus type were not significant.

The third section of Table 3 shows the analysis of low-level gaze statistics for Experiment 3. There was a significant increase in the percentage of saccades that occurred within impressions. There was no significant difference for proportion of fixations on the left (vs. the right), mean fixation duration, mean saccade amplitude on the left and mean saccade amplitude on the right.

3.3.5 *Discussion.* Experiment 3 was designed to examine the nature of the training effect. Via highlighted training with incongruent fixation sequences, Experiment 3 eliminated any systematic relationship between cued locations and local, fine-scale visual features of the fingerprint images (the minutiae) while preserving the relationship between cued locations and the global, coarse-scale visual features (e.g., the core and delta). A comparison of Experiments 1 and 3 suggests that subjects become sensitive to the coarse-scale information in fingerprint images as a result of highlighting training. The evidence that subjects become sensitive to fine-scale information is weaker: the K-S test failed to show a shift in the fixation distribution for transfer stimuli with incongruent highlighting, but did show a shift with congruent training.

The comparison of Experiments 1 and 3 provides a hint that our design—a between-subject comparison with only 12 subjects per condition—may not have sufficient power to detect differences in the magnitude of the training effect between congruent (Experiment 1) and incongruent (Experiment 3) fixation sequences. And with only 64 training trials, the experiments might not have had sufficient duration to instill differences either. As a final caveat concerning the interpretation of our results, our dependent measure may not be adequately sensitive to subtle differences in expert behavior across stimuli. To understand why, we need to reconsider the modeling of novice and expert behavior. Free parameters of the models are optimized to discriminate novices and experts via the log-likelihood ratio (Equation A5). Although we've treated the models as if they can be used to assess expertise in general, the model fitting procedure is aimed at maximizing between-group discrimination, and it has no explicit requirement to distinguish between fine-

5. Although subjects were not randomly assigned to Experiment 1 versus Experiment 3 (the three experiments were run in sequence) all subjects were drawn from a paid subject pool and participated early in the semester, which mitigates differences between the time each subject participated in the experiment and concerns regarding the influence of the end of the semester.

and coarse-scale expertise. Indeed, the model bandwidth parameter forces a choice along the continuum of fine- to coarse-scale spatial representation, and the best fitting parameter value (a standard deviation of .5 degrees of visual angle) essentially blurs locations within the foveal region, forcing a scale of representation that might not be sufficiently sensitive to make the precise discriminations needed to evaluate fine-scale expertise. To test the conjecture that our model was simply not well suited to discriminate coarse- and fine-scale expertise, we constructed a new model that was specifically designed to perform this discrimination. We define fine-scale expertise to be the pattern of fixations obtained from experts on a particular print. We define coarse-scale expertise to be the pattern of fixations obtained from experts combined across all prints. Since the images are fairly well aligned, the coarse-scale expertise should reflect attention on the core and delta of a print. We obtain a log likelihood ratio that distinguishes between coarse and fine scale expertise, just as previously we obtained a log likelihood ratio that distinguishes between novices and experts. Using this dependent variable, one can ask whether Experiment 3 produced a weaker shift toward fine-scale expertise than did Experiment 1. Unfortunately, the outcome of this exploration was negative: it did not further distinguish the results of Experiments 1 and 3, offering no support for the hypothesis that subjects are learning to detect critical local visual features.

One possible explanation for our failure to find support for the hypothesis is based on the possibility that expert and novice visual spans are different. In domains such as chess and medical imagery, experts do not necessarily foveate on the local features, but rather utilize parafoveal and peripheral vision to attend to local features in a larger spatial context (e.g., Reingold, Charness, Pomplun, & Stampe, 2001). Experts might fixate on the centroid of an arrangement of local features due to the learned expert unitization or chunking of co-occurring features (e.g., Reingold et al., 2001). The benefit of cueing novices to expert fixation locations presumes either that experts directly fixate minutiae or that a novice has a comparable visual span to perceive minutiae arrangements not directly fixated. Novices, with a smaller visual span, will lack the capacity to process minutiae that are not collocated at the fixation (Foulsham, et al., 2012) (Foulsham, et al., 2012)point. To test this hypothesis, we intend to conduct further experiments that will attempt to train visual span as well as fixation using local variability in expert fixations as a proxy for the appropriate visual span.

4. GENERAL DISCUSSION

Our work aims to understand the nature of attentional expertise by training subjects to fixate a series of locations via spatial cues and evaluating changes in gaze behavior as a consequence of training. Experiment 1 showed that highlighted training based on expert fixation sequences yields a shift toward expertise, not only for the trained stimuli but for novel stimuli as well. A modest number of training trials—8 trials on each of 8 images, each for about 30 s—was sufficient to yield appreciable behavioral changes. Experiment 2 provided a control to rule out the possibility that mere exposure or familiarity were responsible for the shift. Experiment 3 offered some insight into the nature of the shift: evidence was obtained for learning on a coarse spatial scale—the scale at which experts and novices are clearly discriminable—but we argued that limitations in our experimental design, dependent variable, and training procedure may explain why we failed to obtain strong evidence for learning about the fine spatial scale.

Our results clearly indicate that highlighting produces a training effect, which in itself is nontrivial in that the procedure is not framed as a training task. Subjects are instructed simply to follow the sequence of highlighted locations, and there's no explicit link between this task and the subsequent evaluation of free fixation behavior. Our results can't be dismissed as reflecting pure priming whereby subjects memorize an association between images and fixation sequences, because we observe generalization to novel images. However, one might interpret our results as a less specific form of priming whereby subjects form associations between image features and eye movements. Our perspective is that one may dismiss this as 'mere priming', but it appears to be exactly the form of learning that subserves attentional expertise. The notion of memorization in general seems misguided given that even when individuals are performing a memory task in which they are motivated to remember the details of an image during viewing, scan paths show significant variability from one viewing to the next (Foulsham et al., 2012).

Three modifications to the highlighted-training procedure may be beneficial. First, as we suggested earlier, it may be useful to highlight not only the fixation location but also the spread or span of attention. Second, one might consider enhancing the discriminative visual features themselves, e.g., by increasing the contrast at locations containing minutiae. Both of these modifications involve accentuating the features for the learner, not simply pointing the learner to a location in the visual field. Third, future research might explore a progressive training procedure in which subjects are first trained to locate and identify individual minutiae types and then training shifts to emphasize configural grouping of local features.

Our long-term objective is to design an approach that efficiently trains novices to perform complex visual tasks, leveraging the spatiotemporal dynamics of expert attention. The present work addresses two key challenges of this process: modeling attentional expertise and assessing the effectiveness of highlighting in training subjects to fixate

task-appropriate locations. However, a third key challenge remains: determining whether expert-like attention will aid in the classification or judgment task an individual must perform. As we discussed in the introduction to this article, past work provides encouraging indications that steering attention can guide higher-order cognitive processing. Expert-like attention should certainly support learning how to utilize critical visual information in a display, and therefore should facilitate the speed or ease with which expert performance is attained. Indeed, if the visual information at the attended location is sufficiently interpretable and diagnostic, knowing where to look could be sufficient to obtain expert performance.

It remains for future research to explore the relative efficiency of highlighting versus more traditional training methods such as explicit instruction. However, highlighting is an attractive training tool because it enables attentional expertise to be shared with novices efficiently and effortlessly while circumventing complications regarding the articulation of visual expertise. Since the highlighting technique employed in this study has little dependency on the domain, it would be straightforward to explore training of expertise in other domains. Many domains could benefit from such training, including those involving visual search (e.g., baggage screening, satellite and medical imagery analysis), sequential decision making (e.g., airplane piloting, air traffic control, driving), classification (e.g., detecting skin abnormalities), and highlighting might benefit not only typical individuals but also individuals with cognitive impairments (e.g., to train face and emotion recognition by autistic individuals).

4.1 Using Models to Understand the Nature of Attentional Expertise

Although our primary focus was on evaluating the benefits of attentional highlighting, the models we constructed to perform this evaluation also offer insight into the nature of attentional expertise. As we described earlier, we compared models that considered the duration of fixations to models that treated all fixations as equivalent. We had expected that the more critical locations would be fixated longer, and thus would be more predictive of expertise. However, we found that consideration of fixation duration provided no leverage in discriminating novice and expert behavior.

Fixation duration is only one means of incorporating time into models. The image-specific models we used to judge expertise predict the spatial distribution of attention by placing a Gaussian blob (kernel) of probability at the x-y location of each fixation. Such a model can be extended to the spatiotemporal domain via a 3D blob that extends in time as well as space. However, such a spatiotemporal model does no better than the pure spatial model in predicting expertise in the Busey et al. (2013) data, even when the blob is allowed to spread broadly in time. This spatiotemporal model makes the strong assumption that individuals examine images in roughly a fixed temporal order. The model's failure to perform better than the pure spatial model suggests that fixation order is relatively flexible.

Nonetheless, there are other means of incorporating sequential information into a model, such as considering the distribution of fixations with respect to the current fixation. In contrast to the spatiotemporal model that assumes locations are visited in a fixed order, this *relative-fixation* formulation claims merely that locations are visited based on their proximity to the current fixation. Proximity weighted gaze models have been constructed by incorporating an oculomotor bias favoring nearby fixations (Chukoskie, Snider, Mozer, Krauzlis, & Sejnowski, 2013; Tatler & Vincent, 2009). In order to model skill level, one might assume that the proximity bias is affected by expertise. In order to model the fingerprint matching task, one might assume an additional bias in favor of saccades to roughly the corresponding location on the fingerprint not currently fixated. Our initial simulations strongly suggest that incorporating a proximity-weighted bias and a switch-to-other-print bias leads to better prediction of subsequent fixations. Further research will determine if including such sequential constraints leads to improved discrimination of expertise.

APPENDIX

A.1 Discriminative Modeling of Novice and Expert Fixation Sequences

This appendix describes the construction of probabilistic models of novice and expert fixation sequences. A probabilistic model will assign a likelihood, $p(\mathbf{F}|s, \mathbf{I})$, to a sequence of fixations, \mathbf{F} , given an image, \mathbf{I} , and a level of performance, $s \in \{\text{novice, expert}\}$. We describe and evaluate a variety of potential models, select the best model, and use the resulting novice and expert models to obtain a discriminative measure, the *degree of expertise* exhibited by an individual based on their fixation sequences. The degree of expertise is used in experimental studies to evaluate the effect of training.

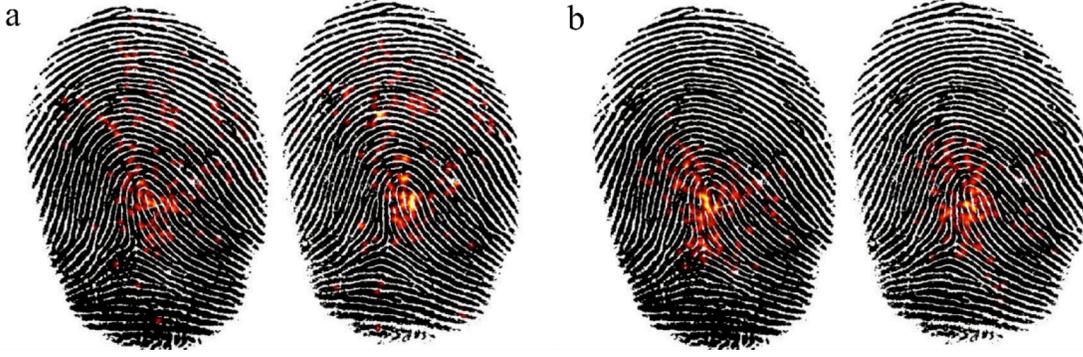


Fig. A1. A kernel density estimate of novices (a) and experts (b) on the same stimulus. Each KDE is constructed using fixations from twelve subjects of the respective population who viewed this stimulus for 10 s. Each fixation has a Gaussian bump centered on it with a bandwidth of approximately 0.5 degrees of visual angle. The color overlaid on the profile of the print corresponds to the likelihood a fixation would be generated at that location. From lowest probability to highest probability is black, dark red, red, orange, yellow and light yellow.

We start by describing two broad classes of models. One class predicts the likelihood of any fixation sequence in *any* fingerprint image. We will call such a model an *image general model (IGM)*. An IGM extracts visual features from the image and uses these features to determine the likely sequence of fixation locations. Using machine-learning tools, one could construct such a model from a set of example images and a set of example fixation sequences observed for those images. The resulting model could make predictions about fixation sequences for novel images. Although such a model is highly desirable, it demands a large dataset and further assumptions concerning feature extraction. As an alternative, in the present work, we use an *image specific model (ISM)* which ignores visual features but instead models the likely sequence of locations for a specific image.

In our modeling, we make the further assumption that the set of fixation locations in a sequence are chosen independently of one another, i.e.,

$$p(\mathbf{F}|s, I) = \prod_i p(\mathbf{f}_i|s, I) \quad (\text{A1})$$

where i is an index over the fixation sequence and \mathbf{f}_i denotes a single fixation. This assumption is certainly incorrect, but it makes modeling tractable and we've found in practice that even without sequential information, the models are sufficiently powerful to discriminate novices and experts. Perhaps this is due to the high variability in the order that fixations occur (Mannan, Ruddock, & Wooding, 1997). Under the independence assumption, a straightforward way to construct an ISM for a population, say experts, is to combine the set of fixations observed for the population on a specific image into a *kernel density estimate (KDE)*. A KDE places a bump of probability mass at the location of every fixation. We assume that the bumps are Gaussians with a spherical covariance matrix, meaning that the bandwidth along the x-dimension is equal to the bandwidth along the y-dimension, and is specified by a single free parameter, σ . The KDE is normalized such that the total area under the mixture of Gaussian components integrates to one, yielding a probability density. Examples of an expert and a novice KDE with a bandwidth of approximately 0.5 degrees of visual angle are shown in Figure A1. Given a location in the image, the KDE can be used to estimate the likelihood that an individual will fixate at that location.⁶

Because the data set used to form the KDE is finite and many locations in the image have near zero likelihoods, we regularize the KDE model (i.e., make it more robust to the variability in the training sample) by mixing it with a uniform distribution. This mixture makes the generative claim that while a fixation is likely to fall around locations where others have looked, with a small probability, c_s , viewers in population s pick an arbitrary location on the image. In total, the resulting ISM for population s has two free parameters: the bandwidth parameter σ_s and the regularization coefficient c_s . The ISM then specifies the likelihood of a fixation sequence \mathbf{F} conditioned on population s based on Equation A1:

6. The reader may be familiar with the related Gaussian mixture model (GMM), which is also a probability density formed from a mixture of Gaussian components. Typically, the number of GMM components is much smaller than the number of data points, and the mean and covariance of each component are fit to the data. Although the KDE might appear to be a more complex model, in that it has as many components as there are data points (fixations), it actually has a far smaller number of parameters to be fit, because the KDE component means are set directly from the data and the covariances are assumed to be identical for all components. We explored GMM fits to the data but found that they were not nearly as useful for our goal, which is to assess expertise of an unknown participant. We suspect that the problem with GMMs is that they lump together fixations on distinct minutiae and therefore are unable to discriminate fixation on the minutiae versus regions between minutiae.

$$p(\mathbf{F}|s, \mathbf{I}) \equiv \prod_i \left[c_s \mathcal{U} + (1 - c_s) \sum_{k=1}^{K_s} \pi_{ks} \mathcal{N}(\mathbf{f}_i; \mu_{ks}, \sigma_s) \right] \quad (\text{A2})$$

where \mathcal{U} is the uniform distribution and \mathcal{N} is the normal (Gaussian). The KDE is specified by the number of mixture components (K_s) in the Gaussian mixture, the relative weighting of the k th mixture component for image \mathbf{I} (π_{ks}), the spatial coordinates (μ_{ks}) of the k th component, and the bandwidth parameter (σ_s) in the spherical covariance matrix for population s . For the KDE model, K_s is the number of fixations.

We considered two possibilities for the mixture weights. First, we supposed that all fixations are treated as equally important, in which case the relative weighting of each mixture component is given by $\pi_k = 1/K_s$. Second, we supposed that the importance of a fixation depends on the duration at which the eyes dwell on a location. To be precise, each fixation \mathbf{f}_i is characterized by a location (f_{xi}, f_{yi}) and a duration f_{di} such that the relative weighting of the k th fixation of n could be given by

$$\pi_k = \frac{f_{dk}}{\sum_{i=1}^n f_{di}} \quad (\text{A3})$$

where i is an index over all fixations used to form the ISM.

A priori it is not clear which of the two fixation weighting schemes to utilize. Although dwell time has been shown to reflect ongoing perceptual and cognitive processing in domains such as reading (Rayner, 1998), it is less clear to what extent fixation duration indexes moment-to-moment attentional processing at the point of fixation versus contributions from other types of processing such as saccade planning or parafoveal processing. However, dwell times have been shown to be diagnostic of false-negative versus true-negative responses for a variety of medical images such as mammography (e.g. Krupinski, 1996) and bone trauma (e.g. Krupinski & Lund, 1997), indicating that dwell times reflect at least some level of task-relevant attentional processing. Because it is difficult to adequately justify one weighting scheme over the other a priori, we consider both schemes in the model selection process.

A.2 Measure of Expertise Using Image Specific Models

The purpose of the models we've described is to evaluate a fixation sequence from an unknown individual and assess the degree of expertise that the individual exhibits. The posterior probability of expertise in Equation 1 can be used to provide this assessment, but an equivalent formulation arises when the prior probabilities of novice and expert are both set to 0.5. In this case, the log likelihood ratio (LLR) conveys the same information as Equation 1:

$$LLR(\mathbf{F}|\mathbf{I}) = \sum_{i=1}^n w_i [\ln p(\mathbf{f}_i|s = \text{novice}, \mathbf{I}) - \ln p(\mathbf{f}_i|s = \text{expert}, \mathbf{I})] \quad (\text{A4})$$

where the weights w_i will be explained shortly. (The logarithm of the likelihood is used to compress the scale of the measure.) The LLR is 0 if the novice ISM and the expert ISM both score the sequence \mathbf{F} as being equally likely; the LLR is positive if \mathbf{F} is scored as being more novice-like and negative if \mathbf{F} is scored as being more expert-like. The magnitude of the LLR reflect confidence in the classification, and can be used to observe incremental changes in expertise.

To this point, we have presented calculations for a single image, but the approach is readily extended to accommodate data from a subject over multiple images via:

$$LLR = \frac{1}{S} \sum_j LLR(\mathbf{F}_j|\mathbf{I}_j) \quad (\text{A5})$$

where S is the number of images and j is an index over images and fixation sequences. Because fixation sequences will have different lengths and we wish to give equal emphasis to each trial regardless of the sequence length, n_j , we set the weight $w_i = 1/n_j$ in Equation A4.⁷

A.3 Training and Validating an Image Specific Model

The log-likelihood ratio is a useful measure only if it can reliably discriminate between expert and novice gaze behavior. In this section, we describe how we train models and then validate their quality, in order to justify the use of these models in measuring changes in expertise resulting from experimental manipulations.

Because we are using ISMs, we require $2N$ models for N images—one for each of the two (novice and expert) populations. In principle, each model has 2 free parameters (σ_s, c_s) beyond the kernel centers which are specified by a set of fixation points. Although models are image specific, the parameters reflect characteristics of an individual's

7. For duration-weighted modeling we scaled the influence of a single fixation by the duration of fixation relative to sum of the duration of fixations in the sequence.

eye-movement control system, and thus should not vary across images. Consequently, the number of parameters is independent of the number of stimuli.

The remaining parameters can be constrained in three different ways. In a *4-parameter model*, the regularization coefficient (c_s) and bandwidth parameter (σ_s) are specific to the population. Allowing the regularization coefficient to depend on expertise follows the intuition that novices will be more likely to go 'off task' and make haphazard gaze choices; and allowing the bandwidth parameter to depend on expertise captures the notion that the degree of specificity in gaze choice may be lower for novices. In a *3-parameter model*, the regularization coefficient is modulated by expertise but the bandwidth parameter is shared between experts and novices. This model is based on the intuition that gaze accuracy depends on noise in the attentional and saccade generation system that is independent of expertise. Finally, in a *2-parameter model*, both regularization coefficient and bandwidth parameter are shared between experts and novices. This 2-parameter formulation supposes that a difference arises between novices and experts only in the kernels (the fixation locations). We consider these three different model parameterizations, and for each, we consider two schemes for weighting fixations (equally weighted, or duration weighted), for a total of six distinct models.

Fitting a model means tuning parameters to maximize the likelihood of the training examples. Specifically, if a fixation sequence F_j is that of a novice, then parameters should be found that maximize $LLR(F_j | I_j)$ (Equation A4), or if F_j is that of an expert, then the parameters should maximize $-LLR(F_j | I_j)$. We therefore use a variant of Equation A5 as our objective function. Constrained nonlinear optimization is performed using the Nelder-Mead simplex algorithm (Lagarias, Reeds, Wright, & Wright, 1998) with the constraints that bandwidth parameters cannot be less than 0.5° of visual angle, which is the approximate lower bound error reported for the training data (Busey et al., 2013), and the regularization coefficient is bounded by the interval $[0, 1]$.

To validate the models, we perform cross validation by splitting the training dataset into a hold-in set and a hold-out set. The hold-in set consists of the fixation data of eleven experts and eleven novices. The hold-out set consists of the remaining expert and novice fixation data and can be used to predict model performance on unseen data via the same LLR measure as used for training.⁸ The training dataset can be split into a hold-in and hold-out set 144 different ways. For each split the hold-out data is evaluated in two ways. The first is a binary *ranking measure* that asks if the hold-out expert has a more expert-like ratio than the hold-out novice. We evaluate all 144 different hold-in/hold-out splits. The second test performs *individual classification*, labeling each hold-out subject as an expert or novice, creating 288 different test cases. The results of this hold-one-out cross validation procedure are shown in Table A1. According to the ranking measure, the two parameter and three parameter formulation perform equally well regardless of the weighting scheme (equal weighting or weighting by fixation duration). However, the individual classification measure reveals that the two parameter formulation is more appropriate. Models with more free parameters perform worse, suggesting overfitting of the training data. Therefore the most parsimonious model is the two parameter formulation. Because the weighting scheme does not seem to affect performance for the 2-parameter model, we prefer the simpler model that gives equal weighting to all fixations. This model appears to do very well in classifying unknown individuals as novices or experts, with 97.92% accuracy in ranking two individuals as more expert like or more novice like, and with a 92.01% accuracy of classifying an individual as novice or expert.

We trained the selected 2-parameter model on the entire dataset; parameters for this model are shown in Table A2. This new fully trained model is used to evaluate gaze data in the experiments that follow. We also leverage the LLR scores for held-out novices and experts in the 144 validation splits; the mean LLR scores are 24.7 and -9.6 for novices and experts, respectively, reflecting a clear differentiation. We use these mean LLR scores to scale the scores of subjects in our experiments, in order to determine where they lie on the continuum of novice to expert fixation behavior.

8. In principle validation could be done by partitioning the data into sets based on stimuli or randomly. However, validation by subjects creates a more realistic and challenging test of how the models will generalize to novel subject gaze data because subjects are likely to have their own systematic gaze traits.

Table A1. Results of a hold-one-out test procedure using a log-likelihood ratio after optimizing on two, three, and four parameter implementations of two weighting schemes. The accuracy of the ranking measure is determined by the correct ranking of the hold-out expert and novice based on their log-likelihood ratio and is comprised of 144 test cases. Individual classification is carried out by independently classifying all hold-out subjects, resulting in 288 test cases.

Weighting Scheme	Equal Weight			Weighting by Duration		
	2	3	4	2	3	4
Parameter Formulation						
Ranking Measure (%)	97.92	97.92	96.53	97.92	97.92	96.53
Classification Measure (%)						
Expert	85.42	76.39	35.42	85.42	81.25	36.11
Novice	98.61	100	100	98.61	100	100
Combined	92.01	88.19	67.71	92.01	90.62	68.06

Table A2. Optimized parameters for two, three, and four parameter formulations for the two weighting schemes. The bandwidth parameters are reported in units of visual angle and the regularization parameters are reported in terms of their exponent.

Weighting Scheme	Equal Weight			Weighting by Duration		
	2	3	4	2	3	4
Parameter Formulation						
Shared Regularization	10^{-148}	-	-	10^{-127}	-	-
Expert Regularization	-	10^{-149}	10^{-153}	-	10^{-117}	10^{-149}
Novice Regularization	-	10^{-98}	10^{-8}	-	10^{-11}	10^{-31}
Shared Bandwidth	0.500°	0.500°	-	0.504°	0.507°	-
Expert Bandwidth	-	-	0.502°	-	-	0.502°
Novice Bandwidth	-	-	2.30°	-	-	2.51°

REFERENCES

- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433-436.
- Busey, T. A., & Dror, I. E. (2011). Special abilities and vulnerabilities in forensic expertise. In *The Fingerprint Sourcebook* (pp. 15.1-15.23). Washington, DC: NIJ Press.
- Busey, T. A., & Parada, F. J. (2010). The nature of expertise in fingerprint examiners. *Psychonomic Bulletin & Review*, *17* (2), 155-160.
- Busey, T. A., & Vanderkolk, J. R. (2005). Behavioral and electrophysiological evidence for configural processing in fingerprint experts. *Vision Research*, *45* (4), 431-448.
- Busey, T. A., Yu, C., Wyatte, D., & Vanderkolk, J. (2013). Temporal Sequences Quantify the Contributions of Individual Fixations in Complex Perceptual Matching Tasks. *Cognitive Science*, *37* (4), 731-756.
- Busey, T. A., Yu, C., Wyatte, D., Vanderkolk, J., Parada, F. J., & Akavipat, R. (2011). Consistency and variability among latent print examiners as revealed by eye tracking methodologies. *Journal of Forensic Identification*, *61* (1), 60-91.
- Chukoskie, L., Snider, J., Mozer, M. C., Krauzlis, R. J., & Sejnowski, T. J. (2013). Learning where to look for a hidden target. *Proceedings of the National Academy of Sciences*, *110*, 10438-445.
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36* (1), 28-71.
- Cook, S. W., Mitchell, Z., & Goldin-Meadow, S. (2008). Gesturing makes learning last. *Cognition*, *106* (2), 1047-1058.
- Deligianni, F. (2014, 4 19). *Talk2Tobii Toolbox*. Retrieved from Birkbeck University of London Web site: <http://www.cbcd.bbk.ac.uk/people/affiliated/fani/talk2tobii>
- Dosher, B. A., & Lu, Z. L. (2005). Perceptual learning in clear displays optimizes perceptual expertise: Learning the limiting process. *Proceedings of the National Academy of Science*, *102* (14), 5286-5290.
- Dror, I. E., & Cole, S. A. (2010). The vision in “blind” justice: Expert perception, judgment, and visual cognition in forensic pattern recognition. *Psychonomic Bulletin & Review*, *17* (2), 161-167.
- Dror, I. E., Champod, C., Langenburg, G., Charlton, D., Hunt, H., & Rosenthal, R. (2011). Cognitive issues in fingerprint analysis: Inter- and intra-expert consistence and the effect of a ‘target’ comparison. *Forensic Science International*, *208*, 10-17.
- Feldon, D. F. (2007). The implications of research on expertise for curriculum and pedagogy. *Educational Psychology Review*, *19*, 91-110.
- Foulsham, T., & Kingstone, A. (2013). Fixation-dependent memory for natural scenes: An experimental test of scanpath theory. *Journal of Experimental Psychology: General*, *142* (1), 41-56.
- Foulsham, T., Dewhurst, R., Nyström, M., Jarodzka, H., Johansson, R., Underwood, G., et al. (2012). Comparing scanpaths during scene encoding and recognition: A multi-dimensional approach. *Journal of Eye Movement Research*, *5* (4), 1-14.
- Grant, E. R., & Spivey, M. J. (2003). Eye movements and problem solving: Guiding attention guides thought. *Psychological Science*, *14* (5), 462-466.
- Jain, A. K., Feng, J., & Nandakumar, K. (2010). Fingerprint matching. *Computer*, *43* (2), 36-44.
- Krupinski, E. A. (1996). Visual scanning patterns of radiologists searching mammograms. *Academic Radiology*, *3* (2), 137-144.
- Krupinski, E. A., & Lund, P. J. (1997). Differences in time to interpretation for evaluation of bone radiographs with monitor and film viewing. *Academic Radiology*, *4* (3), 177-182.
- Lagarias, J. C., Reeds, J. A., Wright, M. H., & Wright, P. E. (1998). Convergence properties of the Nelder--Mead simplex method in low dimensions. *SIAM Journal on Optimization*, *9* (1), 112-147.
- Litchfield, D., Ball, L. J., Donovan, T., Manning, D. J., & Crawford, T. (2010). Viewing another person’s eye movements improves identification of pulmonary nodules in chest x-ray inspection. *Journal of Experimental Psychology: Applied*, *16* (3), 251-262.
- Ludwig, C. J., & Gilchrist, I. D. (2002). Stimulus-driven and goal-driven control over visual selection. *Journal of Experimental Psychology: Human Perception and Performance*, *28* (4), 902-912.
- Mannan, S. K., Ruddock, K. H., & Wooding, D. S. (1997). Fixation sequences made during visual examination of briefly presented 2D images. *Spatial Vision*, *11* (2), 157-178.
- Masson, M. E., & Loftus, G. R. (2003). Using confidence intervals for graphically based data interpretation. *Canadian Journal of Experimental Psychology*, *57* (3), 203-220.
- Mozer, M. C., Pashler, H., Lindsey, R., & Jones, J. (Submitted for publication). Efficient training of visual search via attentional highlighting.

- National Academy of Sciences. (2009). *Strengthening forensic science in the United States: A path forward*. Washington, D.C.: National Academy of Sciences.
- Noton, D., & Stark, L. (1971). Scanpaths in saccadic eye movements while viewing and recognizing patterns. *Vision Research*, *11* (9), 929-942.
- Peiyuan, H., & Kowler, E. (1992). The role of saccades in the perception of texture patterns. *Vision Research*, *32* (11), 2151-2163.
- Peterson, M. F., & Eckstein, M. P. (in press). Learning optimal eye movements to unusual faces. *Vision Research*.
- Peterson, M. F., & Eckstein, M. P. (2012). Looking just below the eyes is optimal across face recognition tasks. *Proceedings of the National Academy of Sciences*, *109*, E3314-23.
- Posner, M. I. (1980). Orienting of attention. *Quarterly journal of experimental psychology*, *32* (1), 3-25.
- Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, *124* (3), 372-422.
- Reingold, E. M., Charness, N., Pomplun, M., & Stampe, D. M. (2001). Visual span in expert chess players: Evidence from eye movements. *Psychological Science*, *12* (1), 48-55.
- Schlingensiepen, K. H., Campbell, F. W., Legge, G. E., & Walker, T. D. (1986). The importance of eye movements in the analysis of simple patterns. *Vision Research*, *26* (7), 1111-1117.
- Seitz, A. R., & Watanabe, T. (2009). The phenomenon of task-irrelevant perceptual learning. *Vision Research*, *49* (21), 2604-2610.
- Sen, T., & Megaw, T. (1984). The effects of task variables and prolonged performance on saccadic eye movement parameters. *Advances in Psychology*, 103-111.
- Tangen, J. M., Thompson, M. B., & McCarthy, D. J. (2011). Identifying fingerprint expertise. *Psychological Science*, *22* (8), 995-997.
- Tatler, B. W., & Vincent, B. T. (2009). The prominence of behavioural biases in eye guidance. *Visual Cognition*, *17* (6-7), 1029-1054.
- Theeuwes, J., Kramer, A. F., Hahn, S., & Irwin, D. E. (1998). Our eyes do not always go where we want them to go: Capture of the eyes by new objects. *Psychological Science*, *9* (5), 379-385.
- Thomas, L. E., & Lleras, A. (2007). Moving eyes and moving thought: On the spatial compatibility between eye movements and cognition. *Psychonomic Bulletin & Review*, *14* (4), 663-668.
- Thompson, M. B., Tangen, J. M., & McCarthy, D. J. (2013). Expertise in Fingerprint Identification. *Journal of Forensic Sciences*, *58*, 1519-1530.
- Turk-Browne, N. B., Isola, P. J., Scholl, B. J., & Treat, T. A. (2008). Multidimensional visual statistical learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34* (2), 399-407.
- Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General*, *134* (4), 552-564.
- Valuch, C., Becker, S. I., & Ansorge, U. (2013). Priming of fixations during recognition of natural scenes. *Journal of Vision*, *13* (3), 1-22.
- van Gog, T., Jarodzka, H., Scheiter, K., Gerjets, P., & Paas, F. (2009). Attention guidance during example study via the model's eye movements. *Computers in Human Behavior*, *25* (2), 785-791.
- Velichkovsky, B. M. (1995). Communicating attention: Gaze position transfer in cooperative problem solving. *Pragmatics and Cognition*, *3* (2), 199-223.