

PAPER

GENERAL

Itiel E. Dror,^{1,2} Ph.D.; Kasey Wertheim,³ M.B.A.; Peter Fraser-Mackenzie,^{2,4} Ph.D.; and Jeff Walajtys,³ B.A.

The Impact of Human–Technology Cooperation and Distributed Cognition in Forensic Science: Biasing Effects of AFIS Contextual Information on Human Experts*

ABSTRACT: Experts play a critical role in forensic decision making, even when cognition is offloaded and distributed between human and machine. In this paper, we investigated the impact of using Automated Fingerprint Identification Systems (AFIS) on human decision makers. We provided 3680 AFIS lists (a total of 55,200 comparisons) to 23 latent fingerprint examiners as part of their normal casework. We manipulated the position of the matching print in the AFIS list. The data showed that latent fingerprint examiners were affected by the position of the matching print in terms of false exclusions and false inconclusives. Furthermore, the data showed that false identification errors were more likely at the top of the list and that such errors occurred even when the correct match was present further down the list. These effects need to be studied and considered carefully, so as to optimize human decision making when using technologies such as AFIS.

KEYWORDS: forensic science, AFIS, contextual effects, distributed cognition, cognitive influences, technology, decision making, bias

The landscape in forensic science, as in other expert domains (e.g., medicine and policing), is changing drastically. A main force in shaping these (and future) changes is technology. Especially influential are *cognitive technologies*—that is, systems that can carry out cognitive operations that were once the sole domain of humans (1). The increased use and reliance on technology have reached a level whereby humans and technology are more and more intertwined and collaborating with one another, creating *distributed cognition* (2,3). With distributed cognition, humans “offload” some cognitive operations onto technology thereby increasing their performance abilities and capacity (4). As human–technology cooperation increases, as they become more intertwined and cognition is increasingly distributed, new opportunities and capabilities arise, as well as new challenges. These have transformed a technological evolution into a revolution. These new possibilities affect human cognition and alter how we go about our professional and personal lives (5).

Distributed cognition may take different forms and generate a variety of modes of collaboration and interaction between the human and technology. Dror and Mnookin (6) specifically distinguished between three modes: At a low level, technology merely

offers a quantitative gain in efficiency; it does not qualitatively transform what is possible. Technology at this level might include, for example, using a computer to store information rather than memorizing it, or using a calculator rather than doing the math. In these cases, the human expert is using technology to save time and cognitive resources. A higher level of distributed cognition and cooperation occurs when the human and technology work side by side as partners. In this case, the technology plays a role that the human expert is incapable of doing (and vice versa: the human expert plays a role that cannot be carried out by the technology). Such human–technology partnerships are based on critical and unique contributions from both the human and the technology, which cannot be simply “offloaded” to the other, and are not a mere matter of convenience and efficiency (e.g., a clinical diagnosis based on an interpretation of an X-ray). At the higher levels of technological use, the technology takes over the more significant and predominant role, leaving the human expert to operate in its shadow (e.g., in breath test detection for alcohol the technological instrument is making the meaningful judgments to the extent that they basically produce a result without significant human intervention). For more detail of this taxonomy, see Ref. (6).

Understanding each mode, both its potential and its limitations, is necessary to make optimal use of both the technological and the human elements in the collaboration. In other words, the success of human experts and technology working in such close collaborations depends on correctly distributing the work among them, and taking advantage of the relative strength each has to offer, and avoiding their respective weakness and vulnerabilities (see, e.g., in face recognition [7,8]).

In general, human expertise, by its very cognitive nature, encompasses a paradox: as experts acquire the cognitive architecture that

¹Institute of Cognitive Neuroscience, University College London (UCL), London, U.K.

²Cognitive Consultants International (CCI), London, U.K.

³Complete Consultants Worldwide (CCW), Bridgeport, WV.

⁴School of Management, University of Southampton, Southampton, U.K.

*Supported by a grant from the National Institute of Justice (Grant Number 2009-DN-BX-K224). Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agency.

Received 4 Oct. 2010; and in revised form 10 Jan. 2011; accepted 21 Jan. 2011.

makes them more effective and efficient, they are also more susceptible to error. This is a result, for example, of using schemas, selective attention, chunking information, automaticity, and more reliance on top-down information, all of which may make them susceptible to missing and ignoring information, and to suffer from tunnel vision and bias (9). This paradox is a result of how the brain processes information and characterizes experts in medicine, policing, aviation, as well as in the forensic domain (10).

One of the vulnerabilities of experts, across domains, is that with their superior performance, they are also susceptible to bias and other contextual and emotional influences. This holds true across expert domains, including forensic fingerprinting (11–20). While interpretations of the research findings do vary, the studies do clearly and consistently show that biasing effects exist, but may or may not change decision outcomes depending on a variety of factors and circumstances. As stated in Langenburg et al. (20), “There is strong evidence that some fingerprint specialists *can* be biased by contextual information. The decision made by a specialist is not necessarily based solely on the ridge detail when comparing images” (p. 577).

It is important to recognize that finding a bias within a decision-making process does not necessarily mean that the conclusions reached are incorrect, nor that they would necessarily have been different in the absence of the biasing information or process. As stated in Dror (14, p. 19).

Bias affects the decision-making process, but not necessarily the decision outcome. Decision-making models clearly illustrate how bias can shift the objective ‘evidentiary weights,’ but that does not mean that every time this shift moves the overall decision outcome past the decision threshold (for details, see Decision Field Theory [21] and Sequential Sampling Models [22]). Bias may just shift the decision in the same direction as the objective ‘evidentiary weights’ which have already surpassed the decision threshold. Furthermore, even when the bias is in the opposite direction of the actual objective ‘evidentiary weights,’ if this shift does not cause a movement past the decision threshold, then it will not result in a change in decision outcome. Therefore, the existence of bias does not necessarily mean that it affects the decision outcome every time, it plays a role in the decision making process.

However, it is equally important to realize that if a decision process is shown to be biased in some way, this generates the *potential* for that bias to affect decision outcomes in some circumstances, depending on the extent of the bias and the proximity to the decision threshold.

All the previous studies that have examined bias in forensic science have focused solely on human examiners and general contextual influences. They have not studied or examined the potential bias introduced by the use of technology. Thus, they do not examine a critical contextual influence that may affect human decision makers: that which is introduced by technology. The growing use and role that technology is (and will be) playing in forensic science requires careful attention and consideration (6). Our study aims to address this gap in the scientific literature.

A relatively new feature of the latent print identification landscape is the increasing use of Automated Fingerprint Identification Systems (AFIS). AFIS is a computerized system that extracts and stores individual characteristics of digitized fingerprints and can be used to search unknown fingerprints or partial latent marks against the stored known fingerprints in the database. AFIS has the ability to store tens of millions of known fingerprints and to perform an automated search against that repository in seconds. In criminal cases, an AFIS is often used to generate possible matches to a latent mark found at a crime scene. The AFIS presents the human

examiner with a set of candidate prints from those contained within the database that, according to the system’s algorithms, are deemed most likely to share a common source with the latent mark in question. AFIS, with its ability to find potential matches among millions of exemplars, is an extremely powerful forensic tool.

There has been significant discussion about the impact that AFIS has had on management, performance and the organization (e.g., [23,24]), but there has been very little discussion of the effects of AFIS systems on the decision processes of human examiners (apart from [6]). Although it was widely recognized and understood that in the latent fingerprint identification context, AFIS is a tool used by examiners rather than a technique for making matches—that “Latent examiners make identifications, not AFIS” (25, p. 4)—the examination of AFIS’s potential influences on those comparisons made by the human latent examiners has been neglected.

It has been suggested that although AFIS is very powerful and has been effective in solving crimes, it also has introduced potential problems that can (and have) lead to erroneous identification (see [6,15–17]). For example, a contributing factor to the Madrid bomber erroneous identification was the remarkable similarity of the two prints (26,27), but the ability to locate these similar prints was a direct result of the great power of the AFIS to search tens of millions of prints. One potential problem is that with the introduction of this new powerful technology, fingerprint comparison by humans has not changed, in the sense that examiners generally continue to use the same thresholds for “sufficient similarity” for determining a match without considering the increased likelihood of seeing two prints from different sources with a high degree of similarity as a result of searching a large AFIS database (6). Davis and Hufnagel (28) seem to suggest that latent fingerprint analysts might have very different views about the matches suggested by the AFIS, but not because of contextual biasing information, which is the focus of this research.

Bias introduced by AFIS has not been studied. However, in a 2005 study that examines bias in forensic science, Dror et al. (17) note that “With the growing use of technology in fingerprint identification, some claim that such human biases and weakness will be reduced, if not eliminated altogether. Although technology is an important ally in fingerprint matching, the issues addressed in this study [i.e., biases], as well as other psychological/cognitive issues, will continue to exist and even increase” (p. 807). A recent theoretical paper on bias and technology in forensic science makes the point that

...the appropriate effects of AFIS on the process of latent fingerprint identification warrant significant further inquiry and consideration. We believe that AFIS does change in important ways the cognitive tasks in which latent fingerprint experts are engaged. Our key point is that there has not yet been sufficient consideration of either the meaning or consequences of the new distributed cognition that AFIS offers, either the new potentialities, or the new risks for error. Such understanding will enable better deployment and utilization of this technology (6, pp. 8–9).

Until now, however, the potential bias introduced by AFIS has not been empirically studied.

Although AFIS makes the scrutiny of large fingerprint databases manageable, thereby offering advances in both process and productivity, its current design also introduces metadata that may impact the subsequent human expert comparison. These data are subject to being controlled by process and workflow design. With the growing use of AFIS, it becomes important to understand how AFIS may affect human fingerprint examiners. Our study scientifically and empirically examines how contextual information provided by AFIS may affect the human examiner. Specifically, AFIS systems return

results to examiners in a ranked order. The primary focus of this study is to examine whether there are decision effects or bias introduced by this ranking. This paper takes a step toward the empirical investigation of the effects of distributing cognition between the human and the AFIS within latent fingerprint examination.

The study reported here is specifically aimed at examining these issues scientifically and empirically, and examining whether the data support any recommendations to minimize any potential biases. This experimental study is designed to detect and quantify the potential influence that the AFIS ranking has on the human experts' perception and cognition that may bias their decision-making process. The conclusions from such empirical studies should be used to develop suggestions for best practices and proper ways to use AFIS technology so as to maintain its benefits while reducing any potential vulnerability that it may introduce.

Our experimental design entails changing the ranking provided by AFIS, that is, for example, taking the "top candidate" in the list, and placing it at the bottom of the list, then observing if the print's position in the list affected how it was judged by the human examiner.

Psychological and cognitive research has demonstrated a general bias toward the first choice, even when the order of choices is random (29–37). AFIS's ranking of potential matches, along with the general already existing bias for the first choice, may work together to influence and affect the perception and judgments of human examiners when they conduct a comparison generated by AFIS. This question—how the rank ordering generated by AFIS affects examiner's cognition—is the topic of this research study.

Method

Participants

We used only latent print experts, all of whom do casework as examiners in forensic laboratories. The 23 examiners used in this study were all experienced examiners that were qualified to do latent fingerprint comparison and to appear as experts in court (mean years of experience was 19.17, $SD = 11.51$). Nearly half of them ($n = 11$) were Certified Latent Print Examiners by the International Association for Identification (IAI). Fourteen were males and nine were females.

Materials

Prints of high quality are less likely to create interpretive problems for examiners; indeed, contextual influences and bias have more pronounced affect with degraded images (e.g., [16,17]). Therefore, we used only latent marks of medium and low quality. We used 160 latent marks and their corresponding prints that matched. We obtained the AFIS list using the National Institute of Standards and Technology (NIST) data set that contains over 3000 known tenprint files. For half of the latents, we generated a list of 10 candidates and for the other half a 20-candidate list. All fingerprint images were captured in standard formats and resolutions; 500 ppi wsq tenprints and 1000 ppi bmp latent mark images.

We then inserted the correct matching print into some of the AFIS lists but not others. When we inserted the match print into the list, we inserted it at or near the top of the list (either as candidate number 1 or near but not at the top (number 2 in the 10-candidate list or number 3 in the 20-candidate list) or at or near the bottom (either as the last candidate or low on the list, as number 8 in 10-candidate list and number 15 in the 20-candidate list). We thereby produced five different AFIS lists for each of the 160 latent marks:

- With no matching print.
- Matching print at the very top of the list (candidate number 1).
- Matching print high on the list (candidate number 2 or 3, for lists of 10 and 20 candidates, respectively).
- Matching print low on the list (candidate number 8 or 15, for lists of 10 and 20 candidates, respectively).
- Matching print at the very bottom of the list (candidate number 10 or 20, for lists of 10 and 20 candidates, respectively).

Procedure

Participants conducted comparisons in this study as part of their normal routine work, not knowing they were taking part in a study. This is critically important, as participants' awareness that they are taking part in a study affects their behavior, especially in studies that examine bias (14–16,18). All comparisons were conducted using a Web-based Remote Examination (WebRex; Lakota Software Solutions, Inc., Fairmont, WV) software that allows the examiners to remotely log in and securely conduct comparisons. This environment is especially suited for our study, as it enables us to manipulate information sent to the practitioners within their normal everyday work and without their knowledge. It is critical to test examiners in their day to day routine work rather than using a contrived experimental setup. This experimental design means that we can draw conclusions about real casework, as the examiners participating in the study are engaging in what they believe to be casework. This is especially essential for studies that examine the effects of context; if the participants know they are taking part in a study, then they do not actually believe the context and therefore its effect is diminished (if not eliminated altogether).

Participants were randomly assigned different lists associated to the same latent mark. That is, for each of the 160 latent marks, some examiners received an AFIS list that contained the matching print as the first print (number 1) on the AFIS list, other examiners got an AFIS list in which the matching print was second in the list (number 2 or 3, depending on the length of the list), other examiners got an AFIS list in which it was located just before the very bottom of the list (number 8 or 15, depending on the length of the list), and for other examiners, it was located at the very bottom of the list (number 10 or 20, again, depending on the length of the list). Finally, most examiners got an AFIS list that did not contain the correct matching print anywhere on the list (to maintain ecological validity we needed to make sure most lists did not have a match, as examiners searching AFIS in real casework do not find matches most of the time). The AFIS lists were distributed and counterbalanced across examiners.

This experimental design enabled us to collect data regarding how latent print experts examine AFIS lists, and we could examine whether the position of the print in the list affected their decision making, and if so, in what ways. Our data permitted us to compare differences in performance, if any, on the same latent mark when it was presented at different positions on the AFIS list. We were interested both in the effect of position on the decision-making process, if any, as reflected by the time it took examiners to reach a conclusion, as well as the decision outcome itself, as reflected by their actual conclusion (see the Introduction). For each of the AFIS lists, the examiners were required to make comparisons on each and every print in the list, and for each one to reach a conclusion of identification, exclusion, or inconclusive.

Each examiner conducted 2400 separate comparisons: 160 latents, each with an AFIS list (half of which were 10-candidate lists and half were 20-candidate lists). Overall, this study includes

data from 55,200 separate comparisons across the 23 latent fingerprint examiners that took part in the study.

Results

The data we obtained in this study are rich in information, and we subjected them systematically to statistical analysis, based on the research questions and experimental design we employed. Our analysis had to distinguish between AFIS lists that did not contain the matching print and those which did, as our experimental design and manipulation included placing the correct matching print at different positions in the AFIS lists for those lists that included the matched print. On the correct print, an error of false positive was not possible; either the examiner would reach the correct conclusion or might erroneously judge the print inconclusive or a false exclusion. On other prints, nonmatches, judgments could either be correct exclusions or incorrect determinations of a match (or inconclusive). Our analysis examined not only the response but also the response time, that is, how long the examiner took to reach a conclusion. Response times (comparison time) were analyzed to gain insight into the decision-making process, whereas errors were analyzed to examine the outcome of the decision-making process.

Our statistical analysis distinguishes between different types of error: false identifications, false exclusions, and false inconclusives (i.e., an inconclusive determination when an identification could, and should, have been made). Although we clearly distinguish between these types of error in our analysis, as noted above, our main experimental design and question of focus was examiner decisions on the matching prints (as a function of their position on the AFIS list). For these matched prints, no erroneous identification was possible. For the other comparisons, while we can compute the rate at which erroneous identifications were made, we would urge caution in taking this as reflecting an error rate (both because our study design was not focused on this question, and because our stimuli, which came from a modestly sized AFIS database may not reflect the range of prints found within casework). Thus, although this research may provide some information regarding error rates, its focus is an analysis of the effect of AFIS ranking on examiner decision making and processes.

Our first set of analyses focused on the false inconclusives and false exclusions made on the matching prints, statistically examining the errors and then comparison times. Our next set of analysis statistically examined false identifications. Overall, the descriptive data are provided in Table 1. Of particular interest is that there were 27.40% missed identifications. A central question is whether these false decisions were affected by our experimental manipulation of position of the matching print in the AFIS list.

Errors on the Matching Prints

This analysis examined the errors made by examiners when the latent was being compared to the matching print (the “target”).

TABLE 1—Overall descriptive statistics.

Description	Statistic (%)
Total number of comparisons	55,200
(number of match comparisons)	1832
Total number of errors (all types)	1516 (2.74)
Number false identifications	49 (0.09)
(excluding “clerical errors”)	12 (0.02)
Number of false “inconclusive” errors	1001 (1.81)
Number of missed identifications	502 (27.40)

TABLE 2—Errors as a function of the position of the matching print.

Set Size	Error	Candidate Position of the Matching Target Print			
		Bottom (%)	Low (%)	High (%)	Top (%)
10	False “Inconclusive”	9.87	14.47	13.82	10.48
10	False “Exclusion”	13.82	15.79	17.11	15.28
20	False “Inconclusive”	12.50	16.45	10.53	13.64
20	False “Exclusion”	16.45	15.79	15.13	12.99
All	Total errors	26.32	31.25	28.30	26.20

The results are shown in Table 2. A three-way repeated-measures analysis of variance (ANOVA) revealed that there was no main effect of the position of the target matching print on overall errors, $F_{3,66} = 1.150$, $p = 0.336$, or the length of AFIS list (i.e., 10 or 20 candidates), $F_{1,22} = 0.011$, $p = 0.919$; nor any statistical difference in the number of false “Inconclusive” or false “Exclusion” decisions, $F_{1,22} = 1.250$, $p = 0.276$. There were also no statistical significant two-way or three-way interactions between these factors. The results at this stage show that there is no main effect of the position of the target matching print as far as the final false exclusion and false inconclusive conclusions are concerned.

Response Times on the Matching Prints

Response time is a critical measure for examining cognitive processing. It is more sensitive and provides more cognitive insights than overall error. We started off with an analysis to examine whether the time spent on a comparison determined the likelihood of error. Specifically, we statistically examined whether a longer time spent by a given examiner comparing the latent mark to a matching print reduces the likelihood of making an error. A logistic regression showed that the log of the comparison time significantly predicted the likelihood of making an error, $\chi^2(1, N = 1832) = 101.28$, $p < 0.001$ (the log of the comparison time was used to reduce the skewed distribution of the response time (RT) data, normalizing the data).

The analysis provides a significantly negative coefficient ($\log RT = -0.477$, $SE = 0.049$, $z = -9.69$, $p < 0.001$), meaning that as the (log of) the comparison time decreased, the likelihood of an error rate is increased. This reflects a speed-accuracy trade-off. While it is not necessarily true that spending extra time than usual would reduce errors, if an individual makes a quicker decision than they need to perform at their optimum, extra errors may occur. Of course, our data simply show a correlation between comparison time and error rates, and we cannot confidently conclude causality without further research.

The second analysis of the comparison times was to determine whether the position of the target matching print (which was the main manipulation in this study) influenced the length of the comparison time. If position did affect the examiners' expectations and biased their comparison, then that will be reflected in the data as a statistically different decision time as a function of the position of the print in the AFIS list (i.e., Top, High, Low, or Bottom of the list).

A two-way repeated-measures ANOVA of the position of the target matching print (Top, High, Low, or Bottom) and AFIS list length (10 or 20 candidates) was performed on the log of the time spent on each comparison (see Fig. 1). The response time for the Bottom position was artificially higher than it should have been if it only included the comparison time, because examiners at the end had to press additional buttons to finalize and submit the batch. The study design did not permit us to distinguish “examination”

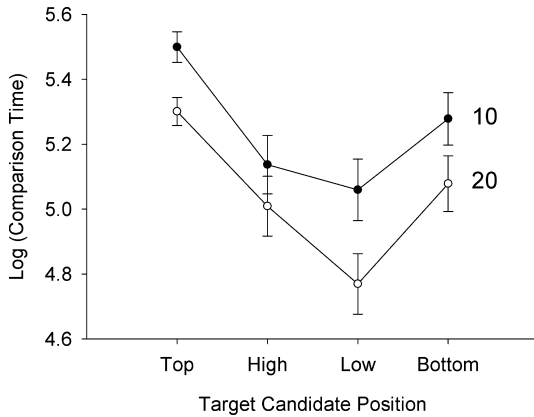


FIG. 1—The effect of the position of the target matching print on average comparison times for the 10- and 20-length Automated Fingerprint Identification Systems (AFIS) lists. The times for the Bottom position were artificially higher, as examiners had to finalize the batch.

time from this finalization time, so the response time for the Bottom position is inclusive of both. Nevertheless, even with this artificial confound, the results showed that there was a significant statistical effect of target position on the comparison time of the target matching print, $F_{3,63} = 18.59, p < 0.001$.

As illustrated in Fig. 1, the top position comparison times are much greater than the comparison times of all the other positions. Indeed, Bonferroni *t*-tests revealed this to be true; when the target was present in the top position (i.e. number 1), examiners spent significantly longer comparing the print than when it appeared lower in the list ($p < 0.001$ for the Top vs. High comparison, $p < 0.001$ for the Top vs. Low comparison, and $p < 0.05$ for the Top vs. Bottom comparison which included the additional artificial time, which made the Bottom vs. Low comparison significant $p < 0.05$). Furthermore, there was a significant effect of length of list on decision times, $F_{1,16} = 14.81, p = 0.001$, reflecting that examiners took on average, per comparison, more time to consider the target matching candidates when the AFIS list length was 10 compared to when it was 20. There was no interaction between the list length and the candidate position on comparison times, $F_{3,66} = 0.291, p = 0.832$.

The results of this analysis are important in light of the data linking comparison time and error rates. Although the result of the analysis of the position of the target matching print candidate on error rates was not significant, the comparison time showed an effect. Accordingly, it is necessary to consider whether the effects of the position on error rate may depend on the time spent on the comparison.

To test for this potential effect, a logistic regression was performed. This regression calculated the likelihood of an error being made as a function of the matching print candidate position and the time spent on the comparison. Because the main effect of position on comparison time (Fig. 1) was the top position (number 1), we collapsed the data into two groups: Top (the target was in position “Top”) and Lower (the target was in position “High,” “Low,” or “Bottom”).

These statistics confirmed our earlier analysis: there was no main effect of target matching print position on error, $\chi^2(1, N = 1831) = 1.352, p = 0.245$, but there was a statistically significant effect of the (log of) the comparison time, $\chi^2(1, N = 1831) = 100.214, p < 0.001$, on the likelihood of error. Most important, this statistical analysis showed a critical interaction between the position

TABLE 3—The parameter estimates for the logistic regression model.

	Estimate	SE	<i>z</i> Value	<i>p</i> (> <i>z</i>)
Intercept	1.9745	0.3270	6.039	<0.001*
Position (top)	-1.3498	0.5381	-2.508	0.012*
Log (comparison time)	-0.5884	0.0660	-8.915	<0.001*
Log (comparison time) × position (top)	0.2771	0.1036	2.676	0.007*

(Top or Lower) and the (log of) comparison time, $\chi^2(1, N = 1831) = 7.187, p = 0.007$.

This is very crucial because it means that the interaction parameter estimates from the model (Table 3) indicate that as comparison time decreases, there is a stronger effect of position on error rates. Specifically, when the target is in a position other than the top position, the examiners were more likely to make an error if they have decreased their comparison time. By contrast, when the examiners took a longer time for the comparison, the effect of the position of the candidate had less of an effect on error rates. In the Discussion section, we consider whether such shorter comparison times reflect reduced effort and motivation in the decision-making process or reflect the use of different cognitive mechanisms and processes of attention (32,33,37).

Figure 2 shows the results of the logistic regression analysis by plotting the model predictions of the log of the comparison time for trials in which the target matching print is at the top of the list compared to when it is elsewhere in the list of candidates.

As illustrated in Fig. 2, the greatest effect of the position of the target is found when the examiner takes a short time for the comparison. When the examiner takes a longer time to do a comparison, the effect of the position is diminished. Accordingly, in contrast to the preliminary analyses regarding error rates and target position, these results demonstrate that the position of the target does influence error rates but that this effect is dependent on the comparison time. Specifically, the regression estimates that the strongest effect is when the comparison time is quicker.

False Identifications

Erroneous identifications are very much a function of the relative similarity between the latent mark and the compared nonmatching print; the greater the similarity is between the two, the more they “look alike” and the greater the chances of an erroneous




FIG. 2—The logistic regression model estimates of the effect of comparison times on error probabilities when the target appears at the top of the list versus at lower positions.

identification (6,15,16). Because we obtained the nonmatching prints from a relatively very small database of just over 3000 ten-prints, there may have been a very small number of such “look alike” nonmatching prints in our study.

Furthermore, we expected to find a low number of false identifications because when examining a number of prints from a list provided by AFIS (on average, 15 comparisons per latent in our study), the potential rate of erroneous identifications is drastically reduced: The maximum number of identifications per list is one. Therefore, the worse scenario—of an examiner making an erroneous identification on each and every list—would give a maximum false identification rate of 6.7% (one out of each 15 comparisons). We elaborate on these issues in the Discussion section, but remind the reader that the research questions and experimental design of this study were to examine the contextual biasing effect of the ranking that the AFIS technology provides, and our analysis (below) on false identification maintains this research focus.

Figure 3 shows that the distributions of false identifications are centered at the upper positions of the AFIS list (numbers 1 and 2). Indeed, logistic regression analyses revealed that candidate position was significantly predictive of whether a choice was a false identification or not for both lengths of AFIS lists (for 10-candidate lists, $\chi^2(1, N = 18,380) = 35.41, p < 0.001$, and for 20-candidate lists, $\chi^2(1, N = 36,840) = 44.17, p < 0.001$).




FIG. 3—False identifications, as a function of position in the Automated Fingerprint Identification Systems (AFIS) list, for both “match present” and “match absent” lists, in AFIS lists with (a) 10- and (b) 20-candidate prints.

TABLE 4—Regression parameter calculated for 10 and 20 Automated Fingerprint Identification Systems (AFIS) candidate lists.

	Estimate	SE	z Value	p (> z)
Intercept (length = 10)	-4.6022	0.2965	-15.521	<0.001*
Candidate position (length = 10)	-0.4378	0.0888	-4.932	<0.001*
Intercept (length = 20)	-5.0609	0.4138	-12.231	<0.001*
Candidate position (length = 20)	-0.5471	0.1386	-3.947	<0.001*

Table 4 shows that the coefficients for both length lists were negative, indicating that there was a greater likelihood of false identifications when the candidate was at the top of the list.

The data above relate to false identifications, regardless of whether they occurred within a list that contained the actual matching print, or within a list that did not include the matching print (some jurisdictions allow examiners to stop doing comparisons after they find an identification in an AFIS list, whereas others require them to continue to check every candidate on the list; in our study, examiners were required to conduct comparisons with all the candidates in the list, i.e., even if they found a match, they were obligated to finish all comparisons on the list).

In a list that does not contain the matching print, the print identified (as false as it is) may still be the most similar to the latent from the entire set of prints in the list. In this circumstance, the examiner has incorrectly concluded that the prints matched, but could have nevertheless select the print with the highest degree of similarity to the latent mark. The contribution to such an error may be a result of the position bias, but also from the mere fact the print chosen was the most similar form all the candidates in the list.

However, if a false identification occurs in a list where the actual matching print is present, the implication is that the examiner has missed a more similar print in the set—the matching print. Therefore, in this circumstance, an error reflects a bias introduced by the position (not because it is the most similar candidate in the AFIS list), as the examiner has selected a wrong print higher in the list, while missing the correct identification print present in a lower position.

We therefore conducted further analysis to focus on false identifications that occurred only within lists that included the actual matching print. This focus also tests a critical issue of this research, that is, whether the matching print might be more likely to be missed if it is positioned lower down the AFIS list. The results of this analysis, illustrated in Fig. 4, show that even in these lists, false

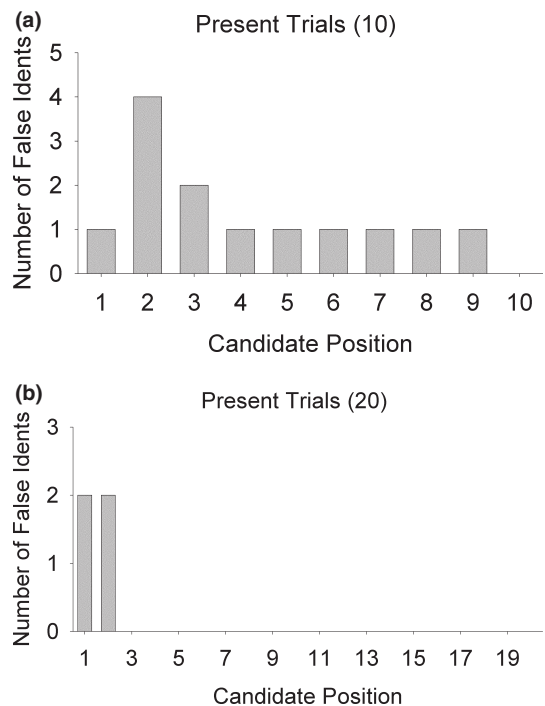


FIG. 4—The number of false identifications in Automated Fingerprint Identification Systems (AFIS) lists that contained the actual matching print, for lists of (a) 10 and (b) 20 candidates.

TABLE 5—Regression parameter calculated only for lists that contained the actual matching print, for 10 and 20 Automated Fingerprint Identification Systems (AFIS) candidate lists.

	Estimate	SE	z Value	p (> z)
Intercept (length = 10)	-5.7297	0.5127	11.175	<0.001*
Candidate position (length = 10)	-0.1713	0.1037	-1.652	0.099
Intercept (length = 20)	-4.7389	1.0013	4.733	<0.001*
Candidate position (length = 20)	-1.0997	0.5778	-1.903	<0.001*

identifications were concentrated at higher candidate positions (but this effect was more pronounced in the longer, 20-candidate, lists, where false identifications were only made in the first and second candidate positions).

A logistic regression further confirmed these findings. Table 5 shows the parameter of the logistic regression model. The results showed that for a list length of 10, the candidate position with a trend for statistical significance, $\chi^2(1, N = 9140) = 2.93, p = 0.087$, and for list length of 20, the candidate position was significant, $\chi^2(1, N = 18,360) = 16.34, p < 0.001$. These analyses show that false identifications tend to be made in high candidate positions. This clearly demonstrates that examiners making use of AFIS technology can result in false identifications even when the actual matching print is present in the same list, as long as it is in a lower position in the list. Indeed, in all four of the false identifications made in the 20-candidate lists, the actual matching print was lower down the list than the candidate that was falsely identified as matching.

A potential concern is that some of these 49 false identifications may have included errors that may appropriately be classified as “clerical errors.” Arguably, errors that are the result of “clerical” mistakes should be excluded from the data set. However, such a step is problematic and can be questioned and criticized. For example, Wertheim et al. (38) reported an error rate of 1.041% but classified almost all the errors (1.007%) as “clerical” mistakes, and only 0.034% of errors were classified and attributed as actual erroneous identification judgments. In an examination of that study, Cole (39) questions their exclusion of errors as “clerical” mistakes.

There are several issues surrounding exclusion of data as “clerical” mistakes. To begin with, how should a clerical error be defined or identified? Even if some false identifications are indeed “clerical errors,” does this automatically justify exclusion? It may well be that such clerical errors are a result of not paying attention; however, this reduced attention may itself be caused by bias, such as, for example, a general cognitive bias to prefer the first choice (32,33,37), or a specific bias relating to AFIS-induced expectation and motivation associated with comparisons of prints at different positions on the AFIS list, or possibly a combination of both: the exact topic and focus of this research. Therefore, one must exercise extreme caution in excluding data. Nevertheless, including data that do not reflect the cognitive processes one is investigating may taint the results and interpretation of the study.

Given this quandary, we decided to include two analyses. The first, reported above, included all the false identifications. The second, reported below, excluded the false identification that were potentially “clerical” in nature, resulting in excluding 37 of the 49 false identifications, giving a false identification error rate of 0.02%. As our criterion for “clerical error,” we excluded any false positive errors on lists where the examiner reported more than one identification within the same list. Examiners knew that—at least in principle—any list could not contain more than one matching




FIG. 5—False identifications (excluding potential “clerical errors”), as a function of position in the Automated Fingerprint Identification Systems (AFIS) list, for both “match present” and “match absent” lists, in AFIS list with (a) 10- and (b) 20- candidate prints.

print; so in those instances when they indicated two (or more) matching prints, this reveals a process mistake that can, at least arguably, be deemed likely to be “clerical” in nature.

With the new data set, now only with 12 false identification rather than 49, we were interested to examine, as we did before, their positions in the AFIS lists. Figure 5 clearly illustrates that false identifications were still centered at the upper positions of the lists (1 and 2).

Logistic regression analyses (shown in Table 6) further confirmed our findings: it revealed that candidate position was significantly predictive of whether a choice was a false identification or not for the 10-candidate lists, and a trend for statistical significance in the 20-candidate lists. The coefficients for both lengths of AFIS lists (10 and 20) were negative, indicating that there was a greater likelihood of false identifications when the candidate was at the top of the list.

As before, we wanted to conduct an analysis when false identifications were made in lists that contained the actual matching print (see Fig. 6).

Although Table 7 shows that the candidate position was not quite significant, this may be the result of the low number of false identifications and therefore the lack of statistical power. Nevertheless, the fact remains that false identifications were only made in

TABLE 6—Regression parameter calculated for 10 and 20 Automated Fingerprint Identification Systems (AFIS) candidate lists (excluding potential “clerical errors”).

	Estimate	SE	z value	p (> z)
Intercept (length = 10)	-4.737	0.708	-6.687	<0.001*
Candidate position (length = 10)	-1.010	0.409	-2.691	0.007*
Intercept (length = 20)	-7.576	0.793	-9.557	<0.001*
Candidate position (length = 20)	-0.209	0.123	-1.697	0.090




FIG. 6—The number of false identifications made in lists that contained the actual matching print (excluding potential “clerical errors”). There were no such errors in the 20-candidate lists, and the data in the figure are from the 10-candidate lists.

TABLE 7—Regression parameter calculated for the 10 Automated Fingerprint Identification Systems (AFIS) candidate lists that included the actual matching print (excluding potential “clerical errors”).

	Estimate	SE	z Value	P (> z)
Intercept (length = 10)	-5.313	1.084	-4.804	<0.001*
Candidate position (length = 10)	-0.915	0.550	-1.665	0.100

the upper candidate positions when the matching print was present lower in the list. This, combined with our earlier analysis of all of the false positive errors, suggests that AFIS and its design may influence the human examiner to make a false identification, even when the actual matching print is present in the list, as long as it is in a lower position.

The important and consistent result is that in both analyses, with and without the potential “clerical errors,” the position in the AFIS list played a critical contributing role in the way examiners conduct their comparisons and conclusions. It can be argued that those candidates in the more upper positions are more similar to the latent and that drives the error, not the position in the list, per se. However, our analysis of the lists that contained the actual matching print dismisses such a claim: if the actual similarity was driving the false identifications and there was no effect based on the position within AFIS, then errors of identification would only be apparent in the lists that did not contain the actual matching print.

Discussion

Technology, and especially cognitive technology that performs cognitive operations that were once only possible by humans, gives rise to new forms of distributed cognition. A forensic technology such as AFIS is a clear and excellent example of new achievements and possibilities that can arise from new human–technology collaboration. It also demonstrates the complexity of such endeavors and illustrates that one should fully consider and understand their impact. Too often technology is deployed with only technical training in its use, without optimizing the human–technology collaboration, adapting human cognition to the new technological environment, or taking steps to minimize potential harm or new risks generated by the technology.

Automated Fingerprint Identification Systems (AFIS) is a major technological apparatus which is widely used. Nevertheless, there has not been a single empirical scientific study that examined its cognitive impact on the world of fingerprinting. Indeed, fingerprint examiners continue to consider similarities for identifications in

exactly the same way as they did in the pre-AFIS era. That is, they used the same decision-making processes whether a print is provided by a suspect or from an AFIS search. SWGFAST (the Scientific Working Group that established guidelines for this profession), the IAI (their professional body), as well as dozens of laboratories that we have examined, do not provide any guidelines or stipulations concerning how examiners may need to change their decision-making processes and threshold when using AFIS.

AFIS has in many ways been a great success, but nevertheless it may introduce problems that have never been studied or researched. AFIS changes the way that comparisons are presented to an examiner. In non-AFIS searches, an examiner is frequently presented with a limited set of prints for comparison. By contrast, in an AFIS setting, an examiner is presented with a ranked set of prints, beginning with the most probable match. Indeed, most AFIS hits are found at the top position of the candidate list. While this ranking may therefore be very useful information for the examiner, it may also create a bias. The research reported here examines the potential biasing effects of the ranking of prints in an AFIS list.

Do human examiners take this ranking information provided by AFIS and use it in their decision-making process? And if they do, is that a problem? The fact that AFIS views certain candidates as more likely to be a match than others may constitute valid and important information for the human examiner to take on board. If AFIS rankings tend to be accurate, human examiners may experience efficiency gains by utilizing that information, and focusing their cognitive resources on the highest-ranking exemplars. However, such influences—warranted and helpful as they may be—need to be carefully researched, considered and formalized (6). Furthermore, such influences may be biasing in a negative way. For example, they may affect expectations, leading to lower attention levels and motivation for comparisons of candidates not at the top of the list, and thereby missing identifications provided by AFIS. Moreover, the less accurate the AFIS rankings, the more problematic it may be if the examiner’s own processes internalize them. There may also be too much of an examiner focus on the top prints in a ranked list, especially given the general psychological and cognitive bias to prefer the first choice (29).

Providing a ranked list is not a necessary feature of a database search process. It would be simple to modify AFIS’s output to eliminate the examiners’ knowledge of AFIS’s ranking, by providing lists to examiners with prints in a random order. However, it is far from clear if that is warranted. To determine whether that change would be beneficial, research would be needed on at least two distinct questions: first, research must examine whether the AFIS ranking has an impact on the human examiner. Our study examined this specific issue and found that AFIS ranking does impact examiner decision-making processes, both by decreasing the time spent by some examiners on lower-ranked exemplars and as a result of this decreased time, increasing the number of missed identifications and also by revealing that false identifications were more likely to occur at the top of an AFIS list. But before we could conclude that such information ought therefore not be provided to the examiner, we would also need to examine whether the ranking information was sufficiently valid that its presentation aided accurate decision making (or more efficient) notwithstanding its potential biasing effect.

In this study, we inserted the matching print into different positions in AFIS lists. We compared how human examiners approached the comparison as a function of where we inserted the matching prints. We examined whether the position in the list would affect their conclusions (i.e., whether they were more likely

to make an error, either a false inconclusive or a false exclusion) and whether it biased their decision-making process (i.e., are they likely to go about the comparison differently, based on the position of the print in the list, as reflected by their response time). Our study design recognized that bias could affect the decision-making process, but not necessarily change the decision outcome every time. The change in decision outcomes is determined by the direction of the bias and its magnitude, and the complexity of the decision itself (14). We therefore were interested in both whether the position on an AFIS list affected decisions, but also whether it affected the cognition and the decision processes itself.

Our empirical study clearly shows that the ranked position of a print in an AFIS list affects the human examiners. This effect may well be without conscious awareness, but the data demonstrate that examiners change their decision-making process as a function of print position. We are able to ascertain this with confidence as our experimental setup and design enabled us to statistically compare performance on the exact same comparison (same pair of latent and print), which only differed in the position of the print in the AFIS list, and all else being equal. Furthermore, given that our expert participants conducted the comparisons as routine casework and hence did not know they were taking part in a study, we can confidently attribute our findings to the real world of fingerprinting (14–16,18).

Our findings show that examiners are more likely to miss an identification (false exclusion) or erroneously judge a matched print inconclusive (false inconclusive) when comparison time is lower. We also found that examiners take less time to compare items when they are presented at a lower position on the list. These findings are not a function of the print itself; the same print is considered differently when presented at a lower position on the ranked list. Furthermore, these two factors appear to interact such that there is a greater biasing effect of the position in the list when the comparison is made quickly than when comparison time is longer.

Thus, the position of an actual matching print on an AFIS list biases the human examiners: if it is at the top of the list, then examiners are less likely not to identify it (i.e., concluded it is an exclusion or an inconclusive), in contrast, if it is lower on the list, then the likelihood for wrongly concluding that the matching print is an exclusion or an inconclusive is greater. This fits with current sequential sample models of perception and judgment which argue that biases are often more influential in low-threshold decision making (21,22).

Our findings also show that when false identifications occur, they are closely centered at the top of the list, further showing the biasing effects of position. Such false identifications occurred even when a more similar print (the actual matching one) was present in a lower position on the same list. We are confident in our finding as we found the exact same results when we included all the false identifications (49 errors, 0.09%) and also when we excluded all those that were deemed potentially as “clerical errors” (keeping only 12 errors, 0.02%).

Although our study does not directly address error rates, it is important to note that false identifications are more likely as the comparison print is more similar to the latent. All of our prints were generated from an AFIS search of a database that contained only a small number of tenprints (just over 3000). We cannot ascertain if the rate of false identification would have been higher if we obtained prints from an AFIS search of tens of millions of prints rather than a few thousands.

Similarly, our finding of 27.40% of missed identifications (false inconclusive and false exclusions) needs to be taken within the

scope of this study. The position of the matching print in the AFIS list (which was the topic of research of this study) contributed to this error rate, and we cannot ascertain what the rate of error would have been without this biasing factor.

Our research demonstrates the effects of the prints’ position in an AFIS list but does not explain its cognitive implications. Do examiners utilize different cognitive schemas and processing when comparing prints in different positions (see [40] for details of different cognitive identification schemas and processing), or do their expectations of finding a match influence their motivation and drive the time, effort, and attention they dedicate to the comparison (32,33,37)? Further research can address these questions, which will help guide recommendations of how best to use AFIS and optimize the human–technological collaboration. Our findings also indicated the potential relevance of list size on examiners’ decision-making processes, that is, whether AFIS provided a 10- or 20-candidate list to the human examiner. While the length of the list mediates the effect of position, more research and data are needed before we can confidently draw conclusions and best practice recommendations on this issue.

Other cognitive aspects regarding the use of AFIS also require further research and data, such as whether there are accuracy or efficiency benefits to provide the exemplar list to an examiner sequentially or simultaneously. Such questions have been heavily researched in other areas of the criminal justice system (e.g., eye witness identifications; see [41–43]) and even in some forensic domains (e.g., face recognition technology; see [7,8]). However, no such research has investigated these issues with regard to AFIS, and therefore, it is impossible to make scientifically informed decisions on better or worse ways to make use of this important and prevalent technology.

Forensic science requires such cognitive research, especially in the domains that rely heavily on human perception and judgment, such as pattern and impression evidence. To make such research scientifically sound as well as applicable to the real world of forensic science, collaborative research projects—like the one reported here—significantly benefit from the involvement of both cognitive research scientists and practicing forensic examiners. It would be premature to recommend best practices regarding AFIS from this study standing alone, but the clear effects of ranking on examiners’ time to decision and false identification decisions demonstrate that AFIS design does have an effect on examiner’s cognitive processes. This study illustrates the importance of continued further research regarding how cognition is, and should be, distributed between humans and cognitive technologies.

Acknowledgments

We would like to thank Phillip Merritt and Aaron Wilson of Lakota Software Solutions for their help in setting up the study. We further owe thanks to Deborah Boehm-Davis, Bruce Budowle, John Butler, Gerald Clough, Jennifer Mnookin, and Brad Ulery for comments on an earlier version of this paper.

References

1. Dror IE, editor. *Cognitive technologies and the pragmatics of cognition*. Amsterdam, The Netherlands: John Benjamins Publishing, 2007.
2. Baber C, Smith P, Cross J, Hunter J, McMaster R. Crime scene investigation as distributed cognition. In: Dror IE, Harnad S, editors. *Cognition distributed: how cognitive technology extends our minds*. Amsterdam, The Netherlands: John Benjamins Publishing, 2008;159–84.
3. Dror IE, Harnad S, editors. *Cognition distributed: how cognitive technology extends our minds*. Amsterdam, The Netherlands: John Benjamins Publishing, 2008.

4. Dror IE, Harnad S. Offloading cognition onto cognitive technology. In: Dror IE, Harnad S, editors. *Cognition distributed: how cognitive technology extends our minds*. Amsterdam, The Netherlands: John Benjamins Publishing, 2008;1–23.
5. Dror IE. Land mines and gold mines in cognitive technologies. In: Dror IE, editor. *Cognitive technologies and the pragmatics of cognition*. Amsterdam, The Netherlands: John Benjamins Publishing, 2007;1–7.
6. Dror IE, Mnookin J. The use of technology in human expert domains: challenges and risks arising from the use of automated fingerprint identification systems in forensics. *Law Prob Risk* 2010;9:47–67.
7. Dror IE, Shaikh A. *Face recognition technology: cognitive considerations in system design*. London, UK: United Kingdom Passport Services (UKPS), 2005.
8. Dror IE, Shaikh A. *Training for expertise in face recognition and working with face recognition technology (TNA)*. London, UK: United Kingdom Passport Services (UKPS), 2005.
9. Dror IE. The paradox of human expertise: why experts can get it wrong. In: Kapur N, editor. *The paradoxical brain*. Cambridge, UK: Cambridge University Press, 2011;177–88.
10. Busey T, Dror IE. Special abilities and vulnerabilities in forensic expertise. In: *The fingerprint sourcebook*. Washington, DC: NIJ Press, 2011; Chapter 15, pp. 1–23.
11. Dror IE, Champod C, Langenburg G, Charlton D, Hunt H, Rosenthal R. Cognitive issues in fingerprint analysis: inter-and intra-expert consistency and the effect of a 'target' comparison. *Forensic Sci Int* 2011; 208:10–17.
12. Schiffer B, Champod C. The potential (negative) influence of observational biases at the analysis stage of fingermark individualisation. *Forensic Sci Int* 2007;167:116–20.
13. Charlton D, Fraser-Mackenzie P, Dror IE. Emotional experiences and motivating factors associated with fingerprint analysis. *J Forensic Sci* 2010;55(2):385–93.
14. Dror IE. On proper research and understanding of the interplay between bias and decision outcomes. *J Forensic Sci* 2009;191:17–8.
15. Dror IE, Charlton D. Why experts make errors. *J Forensic Ident* 2006;56(4):600–16.
16. Dror IE, Charlton D, Péron AE. Contextual information renders experts vulnerable to make erroneous identifications. *Forensic Sci Int* 2006;156(1):74–8.
17. Dror IE, Peron A, Hind S, Charlton D. When emotions get the better of us: the effect of contextual top-down processing on matching fingerprints. *Appl Cogn Psychol* 2005;19(6):799–809.
18. Dror IE, Rosenthal R. Meta-analytically quantifying the reliability and biasability of fingerprint experts' decision making. *J Forensic Sci* 2008; 53(4):900–3.
19. Hall LJ, Player E. Will the introduction of an emotional context affect fingerprint analysis and decision making? *Forensic Sci Int* 2008;181:36–9.
20. Langenburg G, Champod C, Wertheim P. Testing for potential contextual bias effects during the verification stage of the ACE-V methodology when conducting fingerprint comparisons. *J Forensic Sci* 2009;54(3): 571–82.
21. Busemeyer JR, Townsend JT. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol Rev* 1993;100:432–59.
22. Dror IE, Busemeyer JR, Basola B. Decision making under time pressure: an independent test of sequential sampling models. *Mem Cognit* 1999;27:713–25.
23. Klug D, Petersen J, Stoney D. *Automated fingerprint identification systems: their acquisition management performance and organizational impact*. Washington, DC: National Institute of Justice, 1992. Report No.: NCI 13749.
24. Peterson JL, Moore J. The status of AFIS systems worldwide: issues of organization, performance and impact. In: Almog J, Spinger E, editors. *Proceedings of the International Symposium on Fingerprint Detection and Identification*; 1995 June 26–30; Néurim, Israel. Jerusalem, Israel: Israel National Police, 1996.
25. Komarinski PD. Considerations for improving latent print processing. Presentation at the National Institute of Standards and Technology (NIST), 2009. http://biometrics.nist.gov/cs_links/latent/workshop09/proc/Komarinski.pdf (accessed 11/28/2011).
26. Stacey RB. Report on the erroneous fingerprint individualization bombing case. *J Forensic Ident* 2004;54(6):706–18.
27. Office of the Inspector General (OIG), Oversight and Review Division U.S. Department of Justice. A review of the FBI's handling of the Brandon Mayfield case. Washington, DC: OIG, 2006.
28. Davis CJ, Hufnagel EM. Through the eyes of experts: a sociocognitive perspective on the automation of fingerprint work. *MIS Quart* 2007;31(4):681–703.
29. Mantonakis A, Rodero P, Lesschaeve I, Hastie R. Order in choice: effects of serial position on preferences. *Psychol Sci* 2009;20:1309–12.
30. MacFie HJ, Bratchell N, Greenhoff K, Vallis LV. Design to balance the effect of order of presentation and first-order carry-over effect in hall tests. *J Sens Stud* 1989;4:129–48.
31. Becker SL. Why an order effect? *Public Opin Q* 1954;18:271–8.
32. Carney DR, Banaji MR. First is the best in rapid social judgment and consumer decision. Cambridge, MA: Harvard University, 2008.
33. Miller JM, Krosnick JA. The impact of candidate name order on election outcomes. *Public Opin Q* 1998;62:291–330.
34. Berg HW, Filippello FE, Hinreiner E, Sawyer FM. Consumer wine-preference methodology studies at California fairs. *Food Technol* 1955;9: 90–3.
35. Coney KA. Order bias; the special case of letter preference. *Public Opin Q* 1977;41:385–8.
36. Dean ML. Presentation order effects in product tasting tests. *J Psychol* 1980;105:107–10.
37. Sulmont-Rosse C, Chabanet C, Issanchou S, Köster E. Impact of the arousal potential of uncommon drinks on the repeated exposure effect. *Food Qual Prefer* 2008;19:412–20.
38. Wertheim K, Langenburg G, Moenssens A. A report of latent print examiner accuracy during comparison training exercises. *J Forensic Ident* 2006;56(1):55–93.
39. Cole SA. The prevalence and potential causes of wrongful conviction by fingerprint. *Golden Gate UL Rev* 2006;37:39–105.
40. Ashworth ARS, Dror IE. Object identification as a function of discriminability and learning presentations: the effect of stimulus similarity and canonical frame alignment on aircraft identification. *J Exp Psychol Appl* 2000;6(2):148–57.
41. Charman SD, Wells GL. Applied lineup theory. In: Lindsay RCL, Ross DF, Read JD, Togli MP, editors. *Handbook of eyewitness psychology: memory for people*. Mahwah, NJ: Lawrence Erlbaum Associates, 2006; 219–54.
42. Wells GL, Olson E. Eyewitness identification. *Ann Rev Psychol* 2003;54:277–95.
43. Turtle JW, Lindsay RCL, Wells GL. Best practice recommendations for eyewitness evidence procedures: new ideas for the oldest way to solve a case. *Canadian J Pol Sec Serv* 2003;1:5–18.

Additional information and reprint requests:

Itiel E. Dror, Ph.D.
 Institute of Cognitive Neuroscience
 University College London
 London
 U.K.
 E-mail: i.dror@ucl.ac.uk
 URL: www.cci-hq.com