The Effect of Image Presentation Rate on Person Identification

Carlos Ibaviosa

The University of Adelaide

This thesis is submitted in partial fulfillment of the Honours degree of Bachelor of Psychological Science (Honours)

Submitted on 1 October, 2019

Total word count: 13,647

Assessable word count: 9,442

Table of Contents

List of Figures

List of Tables

Abstract

Our ability to recognise complex images across contexts depends on our exposure to similar instances. For example, despite much natural variation, it is easier to recognise a new instance of a familiar face than an unfamiliar face. As we encounter similar images, we automatically notice structural commonalities and form a representation of how the image generally looks, even when each image is presented rapidly (i.e., several milliseconds each). However, it is not clear whether this process allows us to better identify new instances of an image compared to assessing single images for a longer duration. Across two experiments, I tested observers’ person recognition ability when presented with rapid image streams at varying rates compared to a single image. Experiment 1 compares performance between upright and inverted faces. Experiment 2 compares performance between fingerprints from the same finger and from the same person more generally. My results suggest that viewing images rapidly is better than single images when identifying faces, but not fingerprints; and that people better recognise upright compared to inverted faces, but are similar in both fingerprint conditions. I discuss the theoretical implications of these results, as well as some practical implications in security and forensic contexts.

Keywords: Visual cognition, recognition, gist perception, ensemble coding, face processing, fingerprint analysis

Declaration

This thesis contains no material which has been accepted for the award of any other degree of diploma in any University, and, to the best of my knowledge, this thesis contains no material previously published except where due reference is made. I give permission for the digital version of this thesis to be made available on the web, via the University of Adelaide's digital thesis repository, the Library Search and through web search engines, unless permission has been granted by the School to restrict access for a period of time.

Oct 2019

Acknowledgements

This thesis would not have been possible without the support and guidance of those around me.

Special thanks to Dr. Rachel Searston for sharing valuable research skills and advice, being an amazing sounding board for ideas, and providing much-needed support throughout the year. From answering questions after work hours to bringing cakes at lab meetings, the extra steps you take to support your students has not gone unnoticed. I wouldn't pick anyone else as my supervisor.

I would also like to thank my lab family and fellow Honours students for creating such a positive environment throughout the year. You were all incredibly comforting during times of stress, whether I needed support or we were all stressed together. There are few things more beneficial to one's mental health than being surrounded by psychology students.

Last but not least, I would like to thank my friends and family. Despite the importance of the Honours year, thank you for reminding me of the wonderful life outside of academia. You'll probably see me a little more now.

Contribution Statement

In writing this thesis, my supervisor and I collaborated to generate research questions of interest and design the appropriate methodology. I conducted the literature search, completed the ethics application, preregistered the project and uploaded all relevant information and materials onto the Open Science Framework. My supervisor and I collaborated in selecting the image databases for both experiments and generating the appropriate subsets: I manually filtered through and formatted over 9,600 face images, and my supervisor selected and formatted the appropriate fingerprint images. I was responsible for all participant recruitment and testing, and my supervisor provided all participation incentives. My supervisor and I collaborated to code all analyses in R Markdown. I wrote up all aspects of the thesis, except for the Materials section relevant to the fingerprint subset.

CHAPTER 1 - INTRODUCTION

People have a remarkable ability to recognise familiar faces. When seeing a friend, for example, we may immediately recognise their face despite having never seen them in that particular context. Their facial expression may change, they may style their hair differently, wear a hat or sunglasses, or appear under different lighting conditions (see Figure 1). Still, we can somehow see past this visual noise and recognise that they are the same person (e.g., Bruce & Young, 1986; Bruce, Henderson, Newman, & Burton, 2001). Recognising a stranger, however, even under optimal viewing conditions, is more difficult because natural variations in hair, expression or pose can be mistaken for genuine differences between people (e.g., Bruce, Henderson, Greenwood, Hancock, Burton, & Miller, 1999; Clutterbuck & Johnston, 2002; Megreya & Burton, 2006). That is, our ability to recognise more familiar faces seems different from our ability to recognise less familiar faces (Megreya & Burton, 2006). In this thesis, I focus on the cognitive processes that may underlie our ability to rapidly identify people. In particular, I explore the effect of varying image presentation rate on face recognition to determine what kind of exposure best familiarises people with new faces—carefully assessing the details of a face, or being rapidly exposed to different examples of the same face.

Figure 1. Natural variation among 64 images of the same face.

1.1 The Role of Prior Experience in Face Recognition

One explanation for our improved recognition for familiar faces is that we develop superior visual memory for objects we have great exposure to (e.g., Sunday, Donelly, & Gauthier, 2018; Tanaka, Curran, & Sheinberg, 2005; Wong, Palmeri, & Gauthier, 2009). Previous research suggests that we become more familiar with a person when exposed to their face for longer compared to shorter durations (e.g., Memon, Hope, & Bull, 2003); and other research suggests that we become more familiar with someone when exposed to several varying instances of their face in different contexts (e.g., Bruce & Young, 1986; Murphy, Ipser, Gaigg, & Cook, 2015). In real life, we can easily accumulate hours of direct exposure to someone we spend a lot time with or see regularly in the media, and this allows us to encode thousands of instances into our visual long-term memory of particular people (see Standing, 1973; Standing, Cornezio, & Haber, 1970). As these instances accumulate, we seem to remember the facial features that are most consistent across contexts to understand what a person looks like (Bruce, 1994; Burton, Jenkins, Hancock, & White, 2005). For example, we may learn to rely more on stable features such as the eyes or nose, and learn to rely less on the more variable features such as hairstyle when making an identification (Bruce et al., 1999, 2001; Young, Hay, McWeeny, Flude, & Ellis, 1985). When viewing a friend's face in a novel context, therefore, we draw upon this expansive repository of similar experiences, automatically assessing whether the stable features of their face align with our visual representation of that face in long-term memory (Bruce & Young, 1986; Burton & Bruce, 1993).

1.2 Ensemble Coding: Understanding the 'Average' of a Face

The influence of these stable representations in long-term memory may work in tandem with a process called 'ensemble coding'—where we make sense of the regularities in our environment by computing the average representation of a set of similar instances (e.g., Ariely, 2001; Chong & Treisman, 2003). Indeed, to make sense of multiple instances of the same object, instead of expending much of our limited cognitive resources remembering specific details of each instance, we perceive all the instances' 'average' properties, or their 'ensemble representation' (see Alvarez, 2011). This process of ensemble coding is incredibly flexible and seems to apply automatically to many basic visual properties, from simple stimuli (e.g., average circle size, Chong & Treisman, 2003; average orientation of patches, Parkes, Lund, Angelucci, Solomon, & Morgan, 2001), to more complex stimuli, such a person's average facial expression and identity (e.g., Haberman & Whitney, 2009).

In the context of face recognition, as we become more familiar with a face, each encounter adds another new instance to visual memory, and as each new memory competes for cognitive resources, new instances are thought to be represented with less precision, generating a more abstract face representation (Dunn, Ritchie, Kemp, & White, 2019; see Perrett, Oram, & Ashbridge, 1998). These ensemble representations allow us to rapidly glean the nature of a given face with a high degree of accuracy. Indeed, when presented with images of unfamiliar faces, the average of a face has demonstrated superior recognition performance even when compared to more analytic judgments of individual instances (Burton et al., 2005; White, Burton, Jenkins, & Kemp, 2014; Whitney & Leib, 2018). These findings suggest that an important aspect of face recognition may lie in our ability to generate ensemble representations of faces from exposure to that face across different contexts.

While research on ensemble coding in face recognition suggests that we create average representations of faces as we store new instances in memory, some evidence suggests that it may not even be necessary to commit particular instances to visual memory. It has been suggested that we need approximately 300 milliseconds of uninterrupted processing to encode an image into visual memory (Potter, 1976); however, research using Mary Potter's rapid serial visual presentation (RSVP) manipulation, where a series of images is presented sequentially at a fixed location very briefly (i.e., several milliseconds per image), has shown that participants can rapidly process and understand each image even when they have no time to encode the images into memory (Potter & Levy, 1969). That is, when instructed to search for a particular image (e.g., a beach) among a rapid sequence of different images (e.g., a forest, followed by a mountain range, followed by a house, etc.), participants could successfully identify the prompted image despite being unable to recall any details (see Potter, 1975, 1976; see also Intraub, 1979, 1980, 1981). Although these RSVP studies initially involved different scene images, studies of ensemble coding have now suggested that when similar images are presented rapidly, we can extract an average representation of the images in the stream (e.g., Im & Chong, 2009; Morgan, Watamaniuk, & McKee, 2000). Indeed, even when faces are presented in our peripheral vision and are not the main focus of the task at hand, we can still discriminate the average emotion or identity presented in the streams (Haberman, Harp, & Whitney, 2009; Ying & Xu, 2017). Despite some research suggesting that our ability to recognise people depends on our representation of that person in long-term memory, it seems that we can also rapidly understand what a person generally looks like even when we cannot commit any individual instance to memory.

1.3 Gist Perception: How Rapid Face Identification Occurs

The ability to instantly process images presented in this RSVP format is an example of ‘gist perception’. Gist perception refers to our ability to accurately identify the nature of an image or object without having to perceive every detail (e.g., Larson & Loschky, 2009; Oliva, 2005; Oliva & Torralba, 2001). Indeed, many studies have shown that we can accurately identify and categorise images even when they are blurred (e.g., see Ruiz-Soler & Beltran, 2006; Torralba, 2009), presented at a reduced resolution (e.g., Searston, Thompson, Vokey, French, & Tangen, 2018), presented for a few milliseconds (e.g., Evans, Georgian-Smith, Tambouret, Birdwell, & Wolfe, 2013; Potter, 1975, 1976), and presented in peripheral vision (e.g., Larson & Loschky, 2009). Studies have suggested that instead of relying on particular features to categorise an image, observers can rapidly identify the image by noticing broad commonalities among similar instances, such as the spatial relations between features (e.g., the distance between the eyes and nose for a particular face; Ruiz-Soler & Beltran, 2006; Schyns, 1998), or other information dispersed throughout the image diagnostic of the object’s category, such as colour or texture (e.g., skin tone or texture; see Oliva, 2005; Oliva & Schyns, 2000). It may be the case, therefore, that much of our immediate recognition of familiar faces, or faces more broadly, occurs within the first few moments of seeing the person, based largely on the general characteristics, rather than particular details, of their face. In an RSVP stream of images, this mechanism may allow us to broadly categorise each image despite having insufficient time to memorise every detail.

1.4 The Face Inversion Effect: Not All Faces Are Perceived Equally

While ensemble coding and gist perception allow us to rapidly process faces, these processes may only operate so efficiently due to our extensive experience with this image category generally. Much of the literature surrounding visual memory, ensemble coding and gist perception incorporate either very basic stimulus classes (e.g., circle size, Chong & Treisman, 2003) or very familiar natural stimulus classes (e.g., facial expression, Haberman & Whitney, 2009)—and so it is less clear how these processes may operate when viewing natural stimulus classes that are less familiar. Exploring how we process faces presented in an unfamiliar manner may help to reveal important aspects of the recognition process, and how it may change under different conditions and with experience. Several studies have investigated, for example, the processing differences between upright and inverted faces (see Tanaka & Simonyi, 2016, and Valentine, 1988 for reviews). Despite generally being considered experts in facial recognition—whether due to specialised processing regions in the brain (e.g., Yin, 1969) or our extensive experience relative to other stimuli (Tanaka & Gauthier, 1997), or a combination—this expertise does not extend to inverted faces. When presented with inverted faces, observers typically perform significantly worse in identification tasks (Hochberg & Galper, 1967), old/new recognition tasks, and familiar and unfamiliar face processing tasks (see Rossion, 2008). These results suggest that the mechanisms we rely on for efficient face processing may not operate equally for all faces—and are most strongly recruited for upright face recognition, given their prominence in everyday life.

A prominent explanation for the face inversion effect is that inversion disrupts ‘holistic processing’. That is, when we perceive the gist of a complex image such as a face, we typically integrate all the features into one single image, rather than processing the features separately (Tanaka & Simonyi, 2016). However, holistic processing seems to depend somewhat on our familiarity with the object (e.g., Campbell & Tanaka, 2018; Tanaka & Simonyi, 2016; Wong et al., 2009; but see McKone & Robbins, 2007 for a critique); and so, compared to upright face processing, where separate facial features are perceived together as a single face (Farah, Wilson, Drain, & Tanaka, 1998; Ingvalson & Wenger, 2005), inverted faces are instead perceived as two eyes, a nose, and a mouth separately (see Rossion, 2008 for a review). This disruption reduces our ability to accurately identify and remember inverted faces (Rossion, 2008). Although the face inversion effect is incredibly robust to task demands and has been widely studied in the context of visual memory (see Tanaka & Simonyi, 2016, and Valentine, 1988 for reviews), no studies have examined the effect of face inversion on gist perception and ensemble coding in an RSVP stream (but see Haberman et al., 2009; Haberman, Lee, & Whitney, 2015; and Leib, Puri, Fischer, Bentin, Whitney, & Robertson, 2012, in relation to ensemble coding generally); and so it is unknown how the perceptual processes that typically allow rapid face recognition in these contexts may operate with inverted faces.

1.5 The Current Study

Although previous research has focused on our ability to glean the average of an image category from an RSVP stream, less is known about how this RSVP paradigm might be used to improve recognition of new instances. In this thesis, I investigate the extent to which the rapid serial presentation of face images affects recognition of new images of the same identity compared with a single face image presented for the same duration. While no research has directly tested this, previous studies have suggested that, if given a limited amount of time to recognise people, exposure to more varying images may be more beneficial than viewing single images (e.g., Murphy et al., 2015; Ritchie & Burton, 2017). Research on gist perception also suggests that we can rapidly process new images of a face within a few milliseconds. Therefore, it is plausible that presenting more images in a rapid stream would boost face recognition compared to viewing a single image for the same duration, given that these conditions provide participants with more exposure to a face in naturally varying contexts, despite the images being presented more briefly.

However, it is also possible that identification may be impeded at certain image rates due to the flashed-face distortion effect (FFDE), where faces presented serially for 200-250 milliseconds have been reported to look distorted as the relative differences between facial features from one image bleed into the next (Tangen, Murphy, & Thompson, 2011). Given that this contrast effect seems to require holistic processing and is less prominent in inverted faces (Bowden, Whitaker, & Dunn, 2019; Tangen et al., 2011), it is possible that it may distract from our ability to correctly identify upright faces when presented in a rapid stream.

To test these possibilities, my first experiment uses the RSVP methodology to vary the number of face images presented per second and attempts to reveal an optimal presentation rate for face recognition given a set amount of exposure time. My guiding hypothesis is that rapidly presenting images will allow participants to become more familiar with the face in different contexts, thereby facilitating recognition of new instances. Seeing a single image of a face for an equal amount of time, on the other hand, may enable better memory for specific details of the face but at the cost of exposure to within-face variability. By keeping the overall duration of the streams constant across conditions, my experiment examines whether it is better to glean the general gist of how a person varies in the more rapid presentation streams, or to assess the details that can be gleaned from any particular image in the less rapid presentation streams. I will also explore whether recognition in these different image presentation conditions is affected by familiarity with general properties of the stimulus class, such as orientation, by presenting the faces as upright and inverted images.

Predictions:

1. In accordance with upright and inverted face-matching literature (see McKone & Yovel, 2009), I predict that my participants will perform better than chance in all conditions
2. I predict that face recognition will increase as image rate increases (i.e., a single image, 2, 4, or 8 images per second)
3. I predict that this benefit will be more pronounced for inverted compared to upright faces, given the existing advantage for upright face-matching and possible interference from the flashed-face distortion effect
4. In the absence of similar prior studies, I predict that confidence will be highest when viewing single images, as it provides the most amount of time for conscious encoding of image details

CHAPTER 2 - EXPERIMENT 1

In this experiment, participants viewed streams of face images from the same person and made a judgement about whether a new face presented immediately afterwards depicted the same or different identity to the person in the stream. I varied the orientation of the faces (upright, inverted) and the image presentation rate for the streams within-subjects, so each participant rated face images after 8-second streams of 1 image, 16 images, 32 images, or 64 images.

2.1 Method

I preregistered my research plan for this experiment on the Open Science Framework (OSF), available here. The wiki page includes a full description of my predictions and hypotheses, methodology, power analysis, analysis plan, and links to all available materials, software, raw data files, and R markdown scripts.

2.1.1 Participants. 30 participants took part in this experiment (19 male, 11 female, mean age of 25) consisting of students from the University of Adelaide and members of the general Adelaide population. All participants were required to be at least 18 years of age, fluent in English, and have normal or corrected-to-normal vision. Participants were incentivised by receiving a $20 Coles/Myer gift card in exchange for their time (see Appendix A). All participants provided informed consent prior to commencing the experiment (see Appendix B).

Participants' responses were to be excluded if they failed to complete the experiment due to illness, fatigue or excessive response delays (i.e., longer than the session allows). Participants who responded in less than 500ms, or consecutively provided the same response, for over 30 percent of trials were also to be excluded. In these cases, another participant was to be recruited and given the same stimulus set according to the previous participant's experiment number. None of the 30 participants met any of these pre-specified exclusion criteria.

2.1.2 Power Analysis. To my knowledge, no previous research has analysed the effect of image presentation rate in a face recognition task. The sample size was determined based on a power analysis assuming a Smallest Effect Size of Interest (SESOI; Lakens, Scheel, Isagar, 2018) of $d = 0.45$ for all effects. Previous studies on face recognition typically show face inversion effect sizes ranging between 0.96 and 1.29 (e.g., Civile, Elchlepp, McLaren, Galang, Lavric, & McLaren, 2018), and so this SESOI was a conservative estimate. With a sample of 30 participants and 96 observations per participant (12 trials × 4 different image presentation rates × 2 levels of image orientation = 96 trials), the experiment had an estimated power of 83.2% to detect a main effect of image presentation rate, and an estimated power of 98.2% to detect an interaction between image presentation rate and orientation. I used Jake Westfall's PANGEA R Shiny App to calculate power given these design parameters.

2.1.3 Design. This experiment had a 4 (presentation rate: single image, 2, 4, 8 images per second) × 2 (orientation: upright vs. inverted) fully within-subjects design. In Experiment 1, participants were presented with a series of 96 face streams for eight seconds. Presentation rate varied across the streams, with participants viewing streams of 64 face images for 125 milliseconds each (8 images per second), streams of 32 face images for 250 milliseconds each (4 images per second), streams of 16 images for 500 milliseconds each (2 images per second), and single images of faces for eight seconds. After a brief 500 millisecond delay, a new 'target' face image from either the same or different person was displayed and participants indicated on a scale whether they believed this new face was the same or different person as the face in the stream, and their confidence in their decision (see Figure 2).

The faces were presented upright for one half of the trials and inverted on the other half. Both orientation blocks were counterbalanced across participants. The four presentation rate blocks were also randomly presented to each participant within the two orientation blocks. Within each presentation rate block, half of the trials depicted the same person as the target image, and the other half depicted a different person to the target image. These trials were randomly presented for each participant.

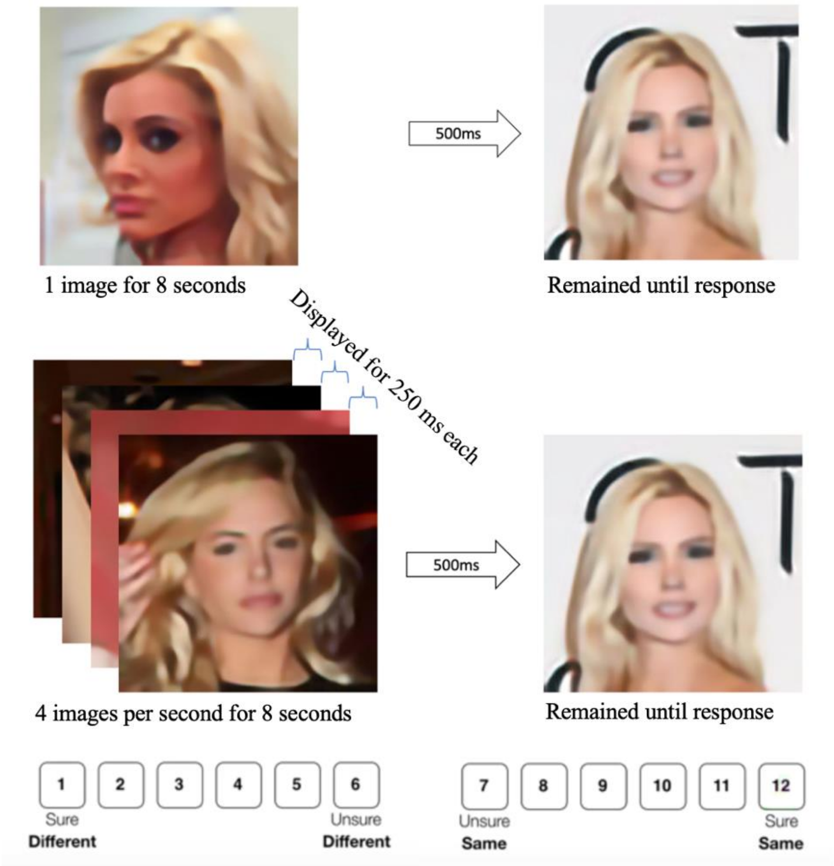

Figure 2. An illustration of the single image condition (top) compared to an example stream (4 images per second; bottom), with a target depicting the same person. The confidence rating scale that appeared with the target is provided underneath.

2.1.4 Measures. Participants indicated their judgments on a 12-point forced choice confidence rating scale: 1 to 6 indicates a “Different” response and 7 to 12 a “Same” response, with ratings closer to 1 and 12 indicating higher confidence than ratings closer to 6 or 7 (see Figure 2). This scale allows me to compute participants’ accuracy (mean proportion correct), and mean confidence (between 1 and 6), and has been used in previous research to compute individuals' discriminability as indicated by the area under their proper Receiver Operating Characteristic (ROC) curve ('AUC'; Vokey, Tangen, & Cole, 2009).

To measure discriminability, I computed each participant's AUC for each condition from their cumulative confidence ratings on same and different trials (see Hanley & McNeil, 1982; Vokey, 2016). An AUC of 1 indicates perfect discriminability, and an AUC of .5 indicates chance performance. A large number of 'hits' (i.e., participant correctly says "Same") and a small number of 'false alarms' (i.e., participant incorrectly says "Same") indicates high discriminability and would produce an AUC score closer to 1, whereas an equal number of hits and false alarms would indicate chance discriminability, resulting in lower AUC scores closer to .5. Participants' confidence is also taken into account in computing AUC, such that lower confidence judgments reflect lower discriminability.

Confidence was computed by collapsing the 12-point rating scale to a 6-point scale. The original scale provided six degrees of confidence for both "Different" (1-6) and "Same" (7-12) responses; and so the collapsed scale isolates confidence by coding all "unsure" responses (6 or 7) to 1, all "moderately unsure" responses (5 or 8) to 2, all "slightly unsure" responses (4 or 9) to 3, and so on—until all "sure" responses (1 or 12) are coded to 6.

2.1.5 Materials. The faces were sourced from the VGGFace 2 dataset (Cao, Shen, Xie, Parkhi, & Zisserman, 2018). The original set contains 3.31 million images of 9,131 identities collected from Google Image searches. I used a subset of 9,600 images of 48 identities (200 images per identity; see the Materials component of the preregistered study). I preserved all natural variation across the images of each identity to increase the difficulty of the target trials (i.e., dissimilar matching identities are more challenging to tell together). The original dataset also contains a large number of blonde, Caucasian, female identities. I constrained my subset to this demographic to increase the difficulty of the distractor trials in the experiment (i.e., similar mismatching identities are more challenging to tell apart). My supervisor and I further increased similarity by computing the distributional characteristics (mean, min, max of image) of each identity and pairing similar identities side-by-side to increase target-distractor resemblance (see Appendix C).

I reduced the original set of images for each identity down to 200 by manually excluding any images with dimensions under 100 $\times$ 100 pixels, drawings, illustrations or animations of faces, significantly occluded faces, faces with distracting watermarks, duplicates or images that clearly depicted a different identity. All other original details were left intact, including natural variation in pose, age, illumination, etc. I then cropped each face to a square using a script in Adobe Photoshop CC (version 20.0.4) and centred the images around the eyes as close as possible. To increase task difficulty, my supervisor and I initially reduced all the images to 64 $\times$ 64 pixels, then upsized them to 400 $\times$ 400 pixels in MATLAB. However, after pilot testing (N = 2) revealed that the task was still too easy for upright faces (mean proportion correct = .92), we further reduced the images to 32 $\times$ 32 pixels. A second pilot (N = 5) then revealed near-chance performance with the inverted faces (mean proportion correct = .59), and so we generated a fresh batch of images reduced to 48 $\times$ 48 pixels to avoid ceiling or chance performance in either condition (see Figure 2).

2.1.6 Software. The video instructions and face recognition task were presented to participants on a 13-inch MacBook Pro, with over-ear headphones. My supervisor developed the software used to generate the trial sequences, present stimuli to participants, and record their responses in LiveCode (version 9.0.2; the open source 'community edition'). The LiveCode source files and experiment code are available in the Software component of the OSF project. The data analytic scripts and plots for this project were produced in RStudio with the R Markdown package. A list of other package dependencies needed to reproduce my plots and analyses are listed in the data visualisation and analysis html file found in the Analyses component of the OSF project.

2.1.7 Procedure. Participants commenced the task after reading an information sheet, signing a consent form, and watching an instructional video. Participants rated a total of 96 faces as the same or different identity to the faces in the stream. In each case, they indicated their judgments on the 12-point confidence rating scale. The response buttons remained on screen until participants selected their rating; however, a prompt to respond within 4 seconds was displayed between trials if participants took longer to decide. Corrective feedback in the form of an audio (correct or incorrect tone) and visual (the selected response button turns green or red) cue is presented to participants after every trial. The whole face recognition task took about 25 minutes to complete.

2.2 Results

I repeated all reported analyses with participants' AUC (discriminability) and raw proportion correct (accuracy) data as planned in my pre-registration. As the pattern of results was the same using both of these performance indicators, for brevity I will only report the analyses I conducted on participants' discriminability. My analyses on participants' accuracy can be found in Appendix D and Appendix E.

2.2.1 Checking Assumptions. My statistical tests involve the following assumptions: normality of differences between paired observations, no extreme outliers, and a continuous dependent variable. Given that the same participants completed each condition, all observations were paired, and the data did not appear to have any severe skewness or deviations from normality (see histograms in Appendix F). Although there was one outlying observation on the upright face trials (see orientation boxplot in Figure 3), other observations from this participant were not outliers and no participants displayed response patterns consistent with my exclusion criteria. It was impossible to determine whether the outlying observation was a genuine outlier or merely due to an experimental artefact (e.g., interruption or distraction), and so removing it may have artificially inflated my orientation effect size. To err on the side of caution, I did not remove the outlier from my dataset. Both discriminability and confidence are continuous measures of performance.

2.2.2 Comparison to Chance. To examine whether participants' performance was better than chance, I calculated participants' discriminability scores for each condition and compared them to randomly generated data. To generate these responses, my supervisor programmed a complementary "sim" participant to respond randomly (i.e., a random match/no-match response at a random response time between 0 and 4000 milliseconds, and a random 1-12 confidence rating) to an identical stimulus set as completed by each human participant. A paired-samples *t*-test comparing participants to their simulated counterparts suggests that participants' discriminability is significantly better than chance ($t(239) = -6.689$, $p < .001$, $d = 0.121$), supporting my prediction that participants should identify faces reasonably well, despite the complexity of the current task and the low resolution (48×48).

2.2.3 Presentation Rate and Orientation. I conducted repeated measures ANOVAs on participants' AUC scores to test whether their ability to distinguish faces of the same versus different identities significantly increased as presentation rate increased, and whether these effects varied as a function of familiarity with the stimulus orientation.

As shown in Table 1, my results suggest that participants are better at recognising faces when viewing rapid streams of the same face compared to single images for both upright and inverted conditions, despite discriminability decreasing overall with inverted faces. A repeated measures ANOVA yielded a significant, medium-to-large (see Cohen, 1988 for conventions) main effect of orientation ($F(1, 29) = 68.258$, $p < .001$, $\eta_G^2 = .148$) and a significant, small-to-medium main effect of image rate ($F(3, 87) = 3.788$, $p = .013$, $\eta_G^2 = .041$) on participants' discriminability scores (see Figure 3). No significant interaction was found

($F(3, 87) = 1.952$, $p = .127$, $\eta_G^2 = .019$). Mauchly's test for sphericity suggested that the assumption of sphericity was met (image rate: $W = .756$, $p = .17$; orientation-image rate interaction: $W = .957$, $p = .942$); and so no corrections were applied to the reported *p*-values. A treatment-control contrast suggested that when compared to viewing a single image, participants' discriminability scores significantly improved under all rapid presentation rate conditions (2 images: $t = 2.192$, $p = .029$; 4 images: $t = 2.468$, $p = .014$; 8 images: $t = 2.431$, $p = .016$). A subsequent trend analysis also revealed a significant linear trend over presentation rate conditions ($t = 2.394$, $p = .018$). That is, discriminability increased in a linear fashion as a function of increasing presentation rate for both upright and inverted faces, despite inverted faces being harder to recognise.

Table 1.
Descriptive statistics for recognition performance in Experiment 1.

	Mean Discriminability (AUC)		Mean Proportion Correct	
Image rate	Upright (*SD*)	Inverted (*SD*)	Upright (*SD*)	Inverted (*SD*)
Single image	.548 (.216)	.462 (.163)	.619 (.138)	.542 (.117)
2 images	.715 (.242)	.473 (.202)	.733 (.190)	.547 (.145)
4 images	.698 (.208)	.513 (.218)	.733 (.151)	.625 (.143)
8 images	.684 (.176)	.524 (.201)	.733 (.139)	.603 (.145)

Note: Discriminability estimates participants' ability to discriminate faces depicting the same versus different people; proportion correct depicts accuracy.

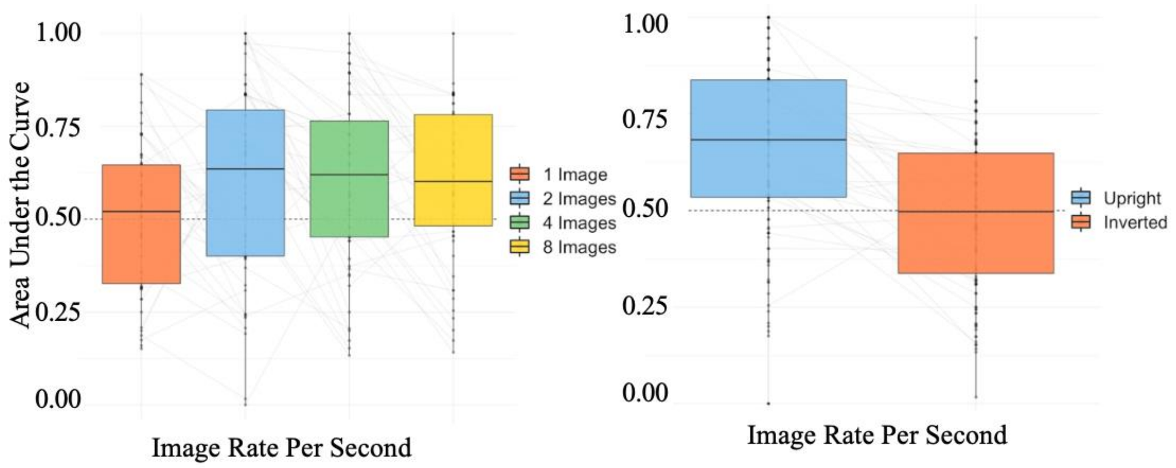

Figure 3. Box plots of participants' discriminability as measured by their AUC scores for the main effects of image rate (left) and image orientation (right).

To address my prediction that confidence will be highest when viewing single images, I analysed participants' confidence ratings for each condition. As shown in Table 2, participants were more confident at identifying upright compared to inverted faces, though confidence seems similar across different presentation rates. A repeated measures ANOVA revealed a significant, medium-to-large main effect of orientation ($F(1, 29) = 8.655$, $p = .006$, $\eta^2_G = .020$), but no significant main effect of image rate ($F(3, 87) = 0.785$, $p = .505$, $\eta^2_G = .002$), and no significant interaction ($F(3, 87) = 0.365$, $p = .779$, $\eta^2_G = .001$; see Figure 4). Mauchly's test for sphericity suggests that the assumption of sphericity was met (image rate: $W = .923$, $p = .818$; orientation-image rate interaction: $W = .885$, $p = .643$); and so no correction was applied to the reported *p*-values. Given that confidence did not significantly differ across image rate conditions, my data did not support the third hypothesis.

Table 2.
Descriptive statistics for confidence in Experiment 1.

	Mean Confidence Scores	
Image Rate	Upright (*SD*)	Inverted (*SD*)
Single image	3.644 (1.394)	3.347 (1.539)
2 images	3.664 (1.538)	3.322 (1.395)
4 images	3.739 (1.446)	3.292 (1.465)
8 images	3.925 (1.419)	3.392 (1.448)

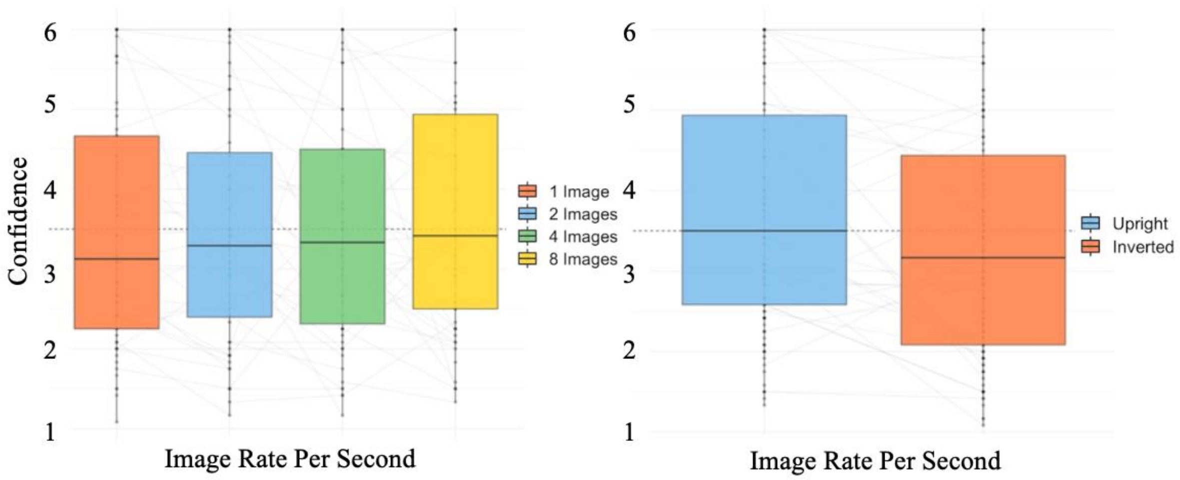

Figure 4. Confidence data for the main effects of image rate (left) and orientation (right).

2.3 Discussion.

2.3.1 Addressing Predictions. This experiment aimed to assess what kind of exposure leads to better face recognition when presented with upright and inverted faces. In line with previous face-matching literature (e.g., Murphy et al., 2015; Ritchie & Burton, 2017), my analyses suggest that overall recognition performance increases as participants view more examples of naturally varying face images. This finding builds upon our previous understanding of the ensemble coding literature. While previous research suggests that RSVP streams allow observers to recognise the average representation easier than individual instances in the stream (e.g., Ariely, 2001; De Fockert & Wolfenstein, 2009), the current study suggests that this ensemble can also facilitate the recognition of new instances of the same category. This is not surprising, given that previous face recognition research suggests that we compare new instances of a familiar face to the average representation of that face in our long-term memory (e.g., Bruce & Young, 1986; Burton & Bruce, 1993).

My results also suggest that the benefit associated with increasing image rate occurred in a similar manner for both upright and inverted faces, despite inverted faces being harder to recognise overall. While lower performance when recognising inverted faces was expected (see Tanaka & Simonyi, 2016, and Valentine, 1988), it is surprising that the RSVP paradigm influenced both upright and inverted faces equally. Given that we already process upright faces more successfully than inverted faces, possibly due to experience (Tanaka & Simonyi, 2016), I expected that image streams may only provide a slight benefit over single images, compared to inverted faces, which may show a larger benefit as image rate increased. The fact that the two orientation conditions increased in a similar manner may be a product of presenting the images at a reduced resolution. During pilot testing, my supervisor and I blurred the images to increase difficulty with upright faces and prevent ceiling effects (e.g., Balas, Gable, & Pearson, 2019). It is possible, therefore, that while an advantage for upright face processing is still evident, it may be less prominent at low resolutions, allowing the image streams to demonstrate a similar advantage for both orientation conditions. However, no studies seem to have tested the face inversion effect at reduced resolutions, and so future research may wish to confirm this conclusion.

I also suspected a lesser advantage for upright faces due to the flashed-face distortion effect (FFDE). The FFDE refers to the apparent distortion of upright (but not inverted) faces presented in an RSVP stream of different random faces, and is thought to emerge due to the relative differences between facial features contrasting from one identity to the next (Tangen et al., 2011). The lack of interaction between orientation conditions, however, suggests that the FFDE had no detrimental effect on either condition. Given that each face in the streams belonged to the same person in the current experiment, rather than different people as is typically the case with FFDE studies (e.g., Balas & Pearson, 2019; Bowden et al., 2019), it may be that the *commonalities* across each face image were exaggerated, rather than the differences, thereby increasing performance when viewing rapid streams. However, given that I did not directly manipulate the FFDE, future experiments may wish to explicitly measure the potential influence of this effect in similar face recognition tasks, to investigate whether it aids encoding of an unfamiliar face.

2.3.2 Considerations. One minor consideration regarding the current methodology is that, given that the selected database sampled faces from Google Images, several of the identities depicted celebrities. Although this provided a suitably large sample of naturally varying face images that could not be found in other databases, this may have increased participants' performance in some trials and inflated my effect sizes, as familiar faces are easier to recognise than unfamiliar faces (Megreya & Burton, 2006). Although an informal post-experiment assessment of each participant's prior familiarity with each face demonstrated that most participants were unfamiliar with the faces (although eight participants reported being previously familiar with 8-13 faces out of 48, and one reported 25), future research may wish to use a dataset containing exclusively unfamiliar faces if possible.

CHAPTER 3 – EXPERIMENT 2

Experiment 1 suggests that presenting similar images in an RSVP stream can facilitate the identification of new instances even when viewing less familiar stimuli (e.g., inverted faces). This method of rapidly presenting multiple similar instances may also be useful in improving performance in other disciplines that rely on identifying naturally varying images—such as fingerprint examination (see Figure 5).

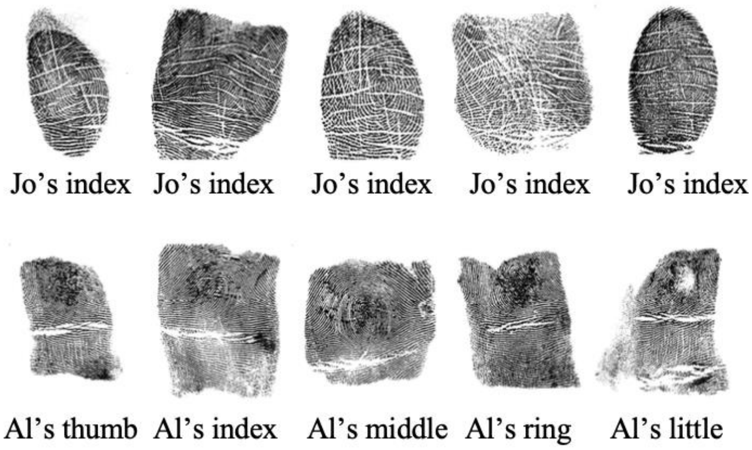

Figure 5. Natural variation among plain or ‘arrest’ fingerprints left by the same finger (top) and same person (bottom).

3.1 Current Practices and Research on Fingerprint Analysis

Fingerprint identifications are made by human examiners who judge the similarity of a crime scene (‘latent’) print against either a single print or a set of known (‘arrest’) prints returned from a large national computer database (e.g., Emerick, Vanderkolk, & Busey, 2015; PCAST, 2016). Fingerprint examiners have been trained to carefully mark up and classify distinguishing features of a crime scene print before comparing them to the known prints since the turn of the 20th century (see Henry, 1900). It is commonly claimed that deliberate analysis is required for accurate fingerprint identification (Busey & Parada, 2010; Cole, 2005); however, recent evidence suggests that fingerprint examiners’ superior identification skills also derive from non-analytic processes that emerge from vast exposure to many different fingerprints over time (Searston & Tangen, 2017a, 2017b; Thompson & Tangen, 2014). Previous research highlights the possibility that fingerprint training could be streamlined by incorporating exposure to varying instances, rather than just carefully comparing individual prints (Thompson & Tangen, 2014). It is possible, therefore, that presenting novices with rapid streams of varying fingerprints will help to simulate experts' experience, therefore improving their fingerprint recognition ability and developing 'expertise' more efficiently.

3.2 The Current Study

In Experiment 2, I explore whether the RSVP method of increasing image exposure, while keeping exposure duration constant, improves person recognition with fingerprints. Experiment 2 employed a similar design to Experiment 1; however, to more closely resemble fingerprint identification procedure, participants were shown the target image of a crime scene print first, before viewing the RSVP stream or single comparison print. Additionally, given that no studies have directly compared performance with fingerprints belonging to same finger versus belonging to the same person more generally (but see Searston & Tangen, 2017c, Tangen, Thompson, & McCarthy, 2011, and Thompson, Tangen, & McCarthy, 2014 for indirect comparisons), I presented participants with image streams of prints from the same finger half the time, and from different prints from the same person the other half the time, rather than in upright and inverted orientations. While evidence suggests that novices can notice general similarities among prints from different fingers of the same person (Searston & Tangen, 2017c), streams from the same finger may contain less variation and therefore may generate a more stable average representation of the finger to compare with the latent print (see Whitney & Leib, 2018).

Predictions:

1. In accordance with previous literature (e.g., Searston & Tangen, 2017c; Tangen et al., 2011; Thompson et al., 2014), I predict that participants would perform better than chance in all conditions.
2. I predict that fingerprint identification would increase as image presentation rate increases, and that this improvement would be more pronounced with streams of the same finger compared to streams of the same person.
3. I predict that confidence would be higher when viewing single images in both conditions.

3.3 Method

In this experiment, participants viewed single images of a latent crime scene fingerprint *before* viewing a stream of fingerprint images. They then determined whether the fingerprints in the stream belonged to the same or different finger, or the same or different person more broadly, to the latent fingerprint (see Figure 5 and Figure 6). As in Experiment 1, presentation rate varied for each stream, and participants' confidence and discriminability were the main performance measures of interest. This experiment was preregistered along with Experiment 1.

3.3.1 Participants. Both experiments were conducted concurrently with the same participants.

3.3.2 Design. Experiment 2 had a 4 (image presentation rate: single image, 2, 4, 8 images per 8-second stream) $\times$ 2 (image specificity: prints from the same finger vs. prints from the same person) fully within-subjects design. Participants judged if a latent fingerprint belonged to the same or different *finger* or *person* as the fingerprint images in a rapidly presented stream of images. In this experiment, participants viewed the latent fingerprint (single image) *before* viewing the image stream. Due to the limited number of fingerprint images in the selected dataset, streams consisted of one-second fingerprint streams presented ‘on loop’ for eight seconds. Participants viewed streams of eight images per second for 125 milliseconds each, streams of four images per second for 250 milliseconds each, streams of two images per second for 500 milliseconds each, and single fingerprint images for eight seconds. Fingerprint streams remained on-screen until a response was made, though participants were prompted to respond within eight seconds (see Figure 6). Participants received corrective feedback for every decision.

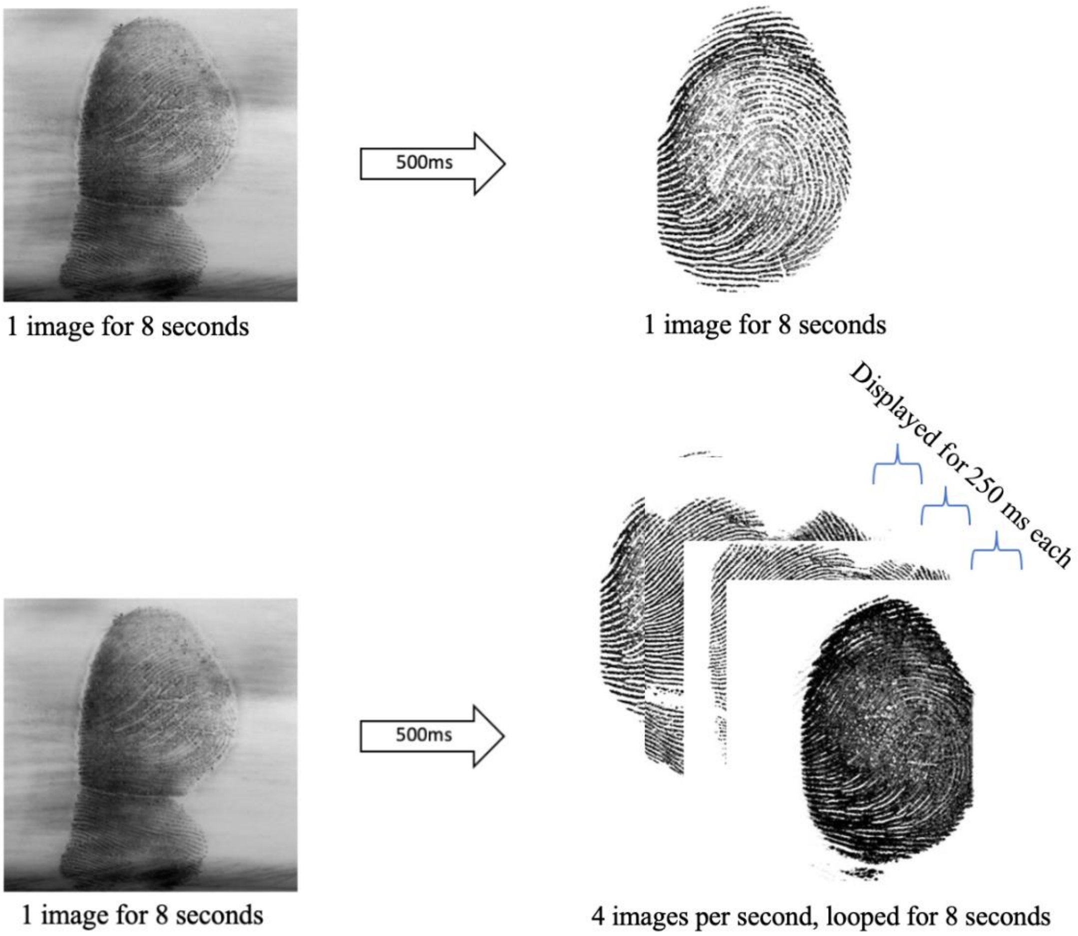

Figure 6. An illustration of the single image condition compared to an example stream (4 images per second), with the left thumb from the same person as the test stimulus.

3.3.3 Materials. The fingerprints were generated from a subset of the Forensic Informatics Biometric Repository (Tear, Thompson, & Tangen, 2010). For the person recognition component of the task, there are ten fully-rolled prints, one from each finger, from 48 different individuals. These served as the rolled prints presented in the rapid streams. For each individual there is also one ‘target’ latent print from the same person, and a ‘distractor’ latent print from another person. The targets and distractors were always taken from the left thumb, as previous research suggests that novices can distinguish prints based on hand type (less so based on finger type; Searston & Tangen, 2017a, 2017b; Thompson & Tangen, 2014). For the finger recognition component of the task, there are eight different fully-rolled impressions from the left thumb of the same 48 individuals. The target and distractor latent prints are the same as those used in the person component of the task

All natural variation in the latent prints was preserved, while the rolled prints presented in the streams were centred on a white background, grey-scaled, level balanced, and cropped to 400 $\times$ 400 pixels (as with the faces). Any distracting borders and text from the arrest cards were removed to isolate the prints.

3.3.4 Software. The software for Experiment 2 was identical to that in Experiment 1. The relevant files are similarly available under the same pre-registration link.

3.3.5 Procedure. Participants were randomly assigned to complete Experiment 2 either immediately before or after Experiment 1. The procedure for Experiment 2 was identical to that in Experiment 1, except for the necessary design changes, and participants were prompted to respond within eight seconds.

3.4 Results

As planned in my pre-registration, I repeated all reported analyses with participants’ AUC (discriminability) and raw proportion correct (accuracy) data. While proportion correct analyses revealed mostly similar results, it suggested no main effect of image rate (F(3, 87) =

2.149, p = .100), contrary to my discriminability analyses. This is likely due to more the difficult nature of the task compared to Experiment 1 (see Comparison to Chance analyses), and that proportion correct analyses are less sensitive than discriminability analyses. I will therefore only report discriminability as it better represents participants' recognition differences between conditions. My analyses on participants' accuracy can be found in Appendix G and Appendix H.

3.4.1 Checking Assumptions. My statistical tests involve the same assumptions as in Experiment 1, and have been addressed in the same way (see histograms in Appendix I for distributional properties of the data).

3.4.2 Comparison to Chance. Similarly to Experiment 1, I compared human performance to chance performance using "sim" data. A paired-samples t-test suggests that my participants are significantly more accurate than chance ($t(239) = -3.318$, $p = .001$, d = 0.058), supporting my prediction that participants should identify fingerprints with reasonable discriminability.

3.4.3 Presentation Rate and Image Specificity. I conducted repeated measures ANOVAs on participants' AUC scores to test whether their ability to distinguish related and non-related fingerprints significantly increased as presentation rate increased, and whether these effects varied as a function of stimulus specificity level. As shown in Table 3, my results show that participants' fingerprint recognition performance generally decreased as image rate increased for both "same finger" and "same person" conditions. My results suggest no significant main effect of specificity ($F(1, 29) = 0.108$, $p = .744$, $\eta_G^2 < .001$), a significant, small-to-moderate main effect of image rate ($F(3, 87) = 3.367$, $p = .022$, η_G^2 = .035) on participants' discriminability, and no significant interaction ($F(3, 87) = 2.053$, p = .112, $\eta_G^2 = .018$; see Figure 7). Mauchly's test for sphericity suggests that the assumption of sphericity was met (image rate: $W = .934$, $p = .862$; specificity-image rate interaction: W =

.827, *p* = .386); and so no corrections were applied to the reported *p*-values. A treatment-control contrast suggested that compared to viewing a single image, participants' discriminability scores significantly decreased when presented with 4 and 8 images per second (2 images: t = -0.897, p = .371; 4 images: t = -2.016, p = .045; 8 images: t = -2.663, p = .008). A subsequent trend analysis also revealed a significant linear trend over presentation rate (t = -2.880; p = .004). That is, discriminability decreased in a linear fashion as presentation rate increased for both same finger and same person conditions—contrary to my predictions.

Table 3.
Descriptive statistics for recognition performance in Experiment 2.

	Mean Discriminability (AUC)		Mean Proportion Correct	
Image rate	Person (*SD*)	Print (*SD*)	Person (*SD*)	Print (*SD*)
Single image	.595 (.242)	.531 (.214)	.656 (.166)	.594 (.176)
2 images	.535 (.218)	.521 (.197)	.619 (.619)	.586 (.141)
4 images	.439 (.182)	.532 (.185)	.542 (.134)	.606 (.126)
8 images	.456 (.262)	.464 (.174)	.575 (.197)	.547 (.143)

Note: Discriminability estimates participants' ability to discriminate faces depicting the same versus different people; proportion correct depicts accuracy.

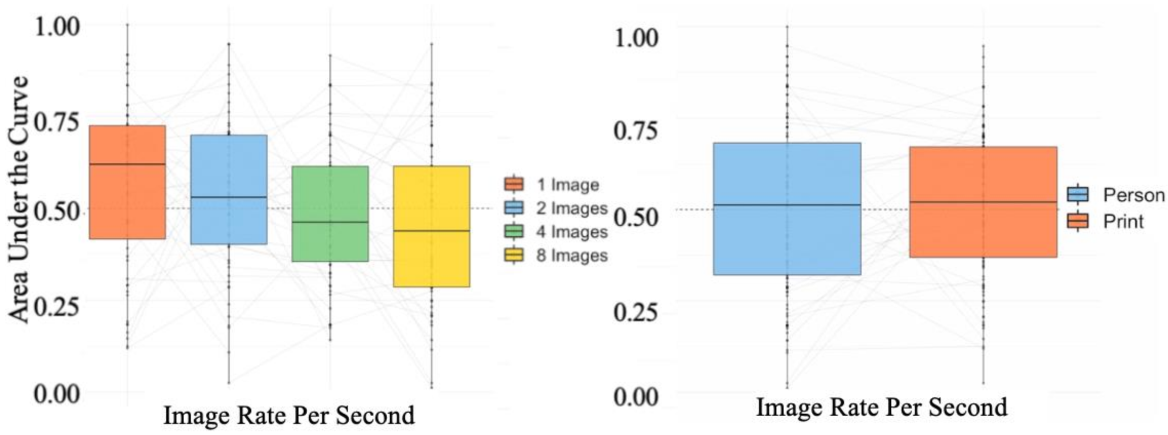

Figure 7. Box plots of participants' discriminability as measured by their AUC scores for the main effects of image rate (left) and image orientation (right).

To investigate my prediction that confidence will be highest when viewing single images, I also examined participants' confidence ratings for each condition. As demonstrated in Table 4, participants were consistently confident across all presentation rates when viewing streams of prints from the same person and prints from the same finger. A repeated measures ANOVA revealed no significant main effect of specificity (F(1,29) = 3.994, p = .055, η_G^2 = .006) or image rate (F(3,87) = 0.763, p = .518, η_G^2 = .002), and no significant interaction (F(3,87) = 0.486, p = .693, η_G^2 < .001; see Figure 8). Mauchly's test for sphericity suggests that the assumption of sphericity was met (image rate: W = .743, p = .144; specificity-image rate interaction: W=.676, p = .054); and so no corrections were applied to the reported p-values. Given that confidence did not significantly differ across image rate conditions, my data does not support my initial prediction.

Table 4.
Descriptive statistics for confidence in Experiment 2.

	Mean Confidence Scores	
Image Rate	Person (*SD*)	Print (*SD*)
Single image	3.097 (1.544)	3.292 (1.367)
2 images	3.147 (1.631)	3.475 (1.381)
4 images	3.086 (1.591)	3.233 (1.509)
8 images	3.008 (1.662)	3.286 (1.552)

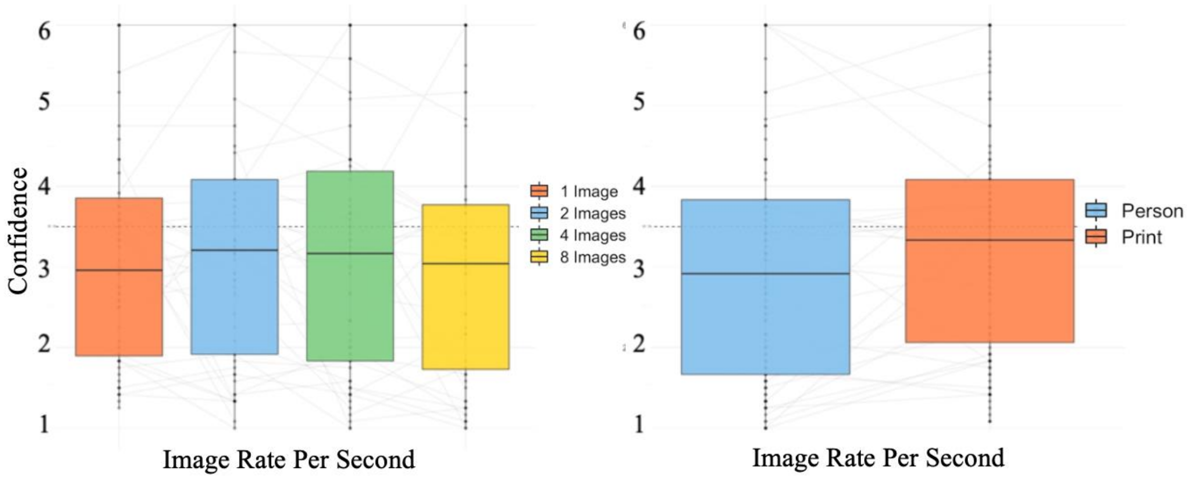

Figure 8. Confidence data for the main effects of image rate (left) and orientation (right).

4.1 Discussion

4.1.1 Addressing Predictions. This experiment aimed to assess whether viewing several impressions of similar fingerprints, either from the same finger or the same person, would better assist novices in making an identification compared to viewing a single fingerprint for a longer duration. My results suggest that this is not the case for either condition. Since novices have no experience in fingerprint matching, it is possible that recognition may benefit from carefully assessing fingerprints, as is currently standard practice (e.g., Busey & Parada, 2010), during the early stages of training. Indeed, given that understanding the images in an RSVP stream seems to rely on holistically processing each image (i.e., perceiving a complex image as a whole, rather than a collection of features; see Oliva, 2005), which may depend on image familiarity (e.g., Tanaka & Simonyi, 2016), it may be that the completely novel nature of the stimulus class required longer exposure to compensate for a lack of holistic processing. If this is true, it is plausible that rapidly presenting fingerprints may have introduced a floor effect in participants' performance—obscuring any positive effect that viewing multiple exemplars may have otherwise exerted. This explanation seems likely, as discrimination performance significantly decreased as presentation rate dropped below 300 milliseconds per image—the approximated minimum duration required to process visual stimuli (Potter, 1976).

The fact that there was no significant difference or interaction between the same person and same finger conditions was also surprising. I suspected that performance would be higher when participants viewed streams from the same finger, to the extent that these streams contain less variation compared those in the 'same person' condition, thus providing a more stable ensemble representation with which to compare the latent print (see Whitney & Leib, 2018). However, while no studies have directly compared the two conditions as in the present experiment, evidence suggests that novices may not perform very differently when asked to match a print to either the same person or same finger (see Searston & Tangen, 2017c, Tangen et al., 2011, and Thompson et al., 2014). It seems likely, therefore, that because novices have no specific fingerprint matching experience like experts, the RSVP methodology allows them to notice general similarities between related prints, regardless of how precisely the prints are related.

4.1.2 Future Directions. While the current results suggest that the RSVP paradigm does not improve fingerprint novice performance, this does not necessarily mean that exposure to various naturally varying fingerprints will not benefit novices. Previous research suggests that images presented in streams of at least one second per image can be efficiently remembered for long periods (e.g., Potter & Levy, 1969; Standing, 1973); and additionally, Thompson and Tangen (2014, Experiment 3) suggested that viewing a print for two seconds only incurred a 6.8 percent decrease in accuracy for novices compared to viewing prints for one minute. It is possible, therefore, that if each fingerprint in the stream was presented for several seconds, rather than several milliseconds, this may optimally balance the advantages of both viewing the detail in a single image and being exposed to variability within images. Future research may wish to either decrease the presentation rate, or allow participants themselves to control presentation rate and view each fingerprint for as long as they deem necessary for familiarisation. The latter manipulation would preserve individual differences in evidence accumulation styles (i.e., some people may prefer more image variation, while others may prefer more viewing time), providing a less intrusive method of investigating how presentation rate might predict identification performance.

Additionally, future research may wish to administer the current experiment to participants with varying degrees of fingerprint-matching experience. Given that novices did not benefit from the RSVP stream (and were no better than chance in some conditions), it is possible that more experienced fingerprint examiners may derive greater benefits from the

RSVP paradigm, as they may process the fingerprints more holistically (Busey & Vanderkolk, 2005; but see Vogelsang, Palmeri, & Busey, 2017 for a competing study). Given that previous research suggests that the majority of learning among novices occurs within the first three months of training (Searston & Tangen, 2017b), it is possible that increasing exposure to varying prints may be most beneficial after the initial learning phase.

CHAPTER 4 - GENERAL DISCUSSION

This thesis examined whether rapidly viewing several instances of complex stimuli, across varying levels of familiarity (Experiment 1) and specificity (Experiment 2), would better facilitate recognition of a new instance compared to viewing a single image for a longer duration. Previous literature suggests that we can recognise new instances of an object based on our prior experience with similar instances (Brooks, 1987; Medin & Ross, 1989). Research on ensemble coding also suggests that we can rapidly understand the general nature of an object as we view several similar, varying instances (e.g., Im & Chong, 2009; Morgan et al., 2000). However, no research has examined how an RSVP-generated ensemble representation may assist in identifying new instances.

Experiment 1 suggests that ensemble coding may indeed facilitate recognition when viewing upright and inverted faces. Given that upright and inverted faces differ only in observers' decreased familiarity with inverted faces (Valentine, 1988), these results suggest that ensemble coding may assist recognition even when exposed to less familiar stimuli. Experiment 2, however, suggests the opposite pattern of results, as fingerprints—a completely unfamiliar stimulus class—showed worse discrimination when participants were presented with RSVP streams from either the same finger or same person as the crime scene print.

5.1 Addressing Predictions

Contrary to my predictions in both experiments, participants' confidence showed no significant differences across image rate conditions, despite single images allowing the greatest encoding time. It may be that the task demands were too difficult in each condition for participants to feel confident. Indeed, identifying different instances of unfamiliar faces has been reported to be a challenging task (e.g., Bruce et al., 1999), which would undoubtedly be harder when the faces are blurred (e.g., Balas et al., 2019; Sanford, Sarker, &

Bernier, 2018); and novice performance in fingerprint matching appears equally challenging (Searston & Tangen, 2017c; Tangen et al., 2011; Thompson et al., 2014). It seems likely that the relative disadvantages in either condition (i.e., less variation with single images compared to less processing time with several images) may have undermined confidence equally across all conditions.

5.2 Discrepancies Between Discriminability Patterns

Although my contradicting discriminability results between the two experiments were unexpected, several explanations are possible. Firstly, the fact that I presented the test stimulus in Experiment 2 before, rather than after the image streams, may have placed greater demands on working memory—especially as the 'more familiar' faces in Experiment 1 (approximated from rapid stream conditions) may have already demanded less from working memory compared to recognising 'less familiar' faces (approximated by single image conditions; Jackson & Raymond, 2008). As opposed to Experiment 1, where the test stimulus remained onscreen until the response, participants in Experiment 2 had to hold a complex, unfamiliar, noisy latent fingerprint in working memory while viewing the subsequent print streams. This working memory demand may have made Experiment 2 more difficult than Experiment 1, particularly as the images became more difficult to process at faster image rates. The fact that ensemble coding seems more beneficial during the encoding stage of learning an identity, rather than on retrieval, seems concurrent with previous research on categorisation. Such research typically suggests that we can identify a new image by comparing its similarity to previously encountered images or representations (e.g., Brooks, 1987; Dopkins & Gleason, 1997). If participants can only view similar instances after being exposed to the test stimulus, as in Experiment 2, then they are not *previously* encountering similar instances to create a representation; they view these images after the fact.

A second possible explanation is that compared to upright and inverted faces, fingerprints may be too difficult for novices to process using the current methodology. Although Experiment 1 suggests that RSVP streams may familiarise observers with less familiar stimuli, fingerprints may simply be too unfamiliar for a similar benefit to occur. The RSVP methodology seems to depend on holistic processing (see Oliva, 2005), and while previous research suggests that we process unfamiliar stimuli less holistically than familiar stimuli (e.g., Campbell & Tanaka, 2018; Wong et al., 2009), holistic and analytic processing seem to be opposing ends of a spectrum, rather than a dichotomy (see Farah, 1992, and Tanaka & Simonyi, 2016). That is, while inverted faces are not processed as holistically as upright faces (Tanaka & Simonyi, 2016), fingerprints may be processed even *less* so, and therefore benefit less from the RSVP paradigm as presentation rate increases. Future research may wish to confirm these suspicions, assessing and comparing our holistic processing abilities with a range of less familiar stimuli (e.g., fingerprints, paintings, bird species) with various recognition or categorisation tasks.

5.3 Discrepancies Between Chance Comparisons

While participants in both experiments displayed better performance than chance, participants in Experiment 1 displayed a higher difference (d = 0.121) than those in Experiment 2 (d = 0.058). In addition to the changes listed above, this difference in overall discriminability may be due to the fact that Experiment 1 had a higher degree of image variation than Experiment 2. In Experiment 1, all images were coloured and blurred and consisted of people in different contexts, including the subsequent test images; however, in Experiment 2 the stream images were somewhat controlled and artificial (i.e., fully-rolled prints, all on a white background) compared to the latent crime scene prints, which may vary in different ways to the prints used in the stream (e.g., contact surface or print pressure). That is, the streams in Experiment 1 were a closer match to the test images than in Experiment 2.

Previous research in face recognition suggests that exposure to more variable images better facilitates recognition in a new context compared to less variable images (Menon, White, & Kemp, 2015; Ritchie & Burton, 2017), and so it is possible that the more controlled nature of the stream images in Experiment 2 may have hindered participants' ability to recognise the test images compared to the more variable stream images in Experiment 1. However, Ritchie and Burton (2017) suggest that reduced variability should nevertheless increase rather than decrease recognition compared to viewing single images. As such, while reduced variability may explain why participants did not benefit from the print streams in Experiment 2, it does not account for the significant decrease in discriminability observed with increasing presentation rates. Of course, it is possible that a combination of the aforementioned design factors may have produced the opposite trends observed across the two experiments.

Another possible factor that may have contributed to the different pattern of results across the two experiments is that Experiment 2 contained fewer unique image exemplars in the streams compared to those in Experiment 1. Given the differences in the selected databases, participants viewed fewer unique fingerprints in each stream compared to the faces in Experiment 1. Indeed, even the highest presentation rate condition in Experiment 2 only showed participants eight unique prints, compared to the slowest stream condition in Experiment 1, which contained 16 unique faces. Given that previous research suggests that viewing fewer different exemplars may decrease recognition of new instances compared to viewing more (Murphy et al., 2015), it is possible that there were not enough fingerprints to produce a similar benefit of presentation rate in Experiment 2. However, it is also important to note that, in real-world fingerprint examination settings, examiners are unlikely to always have access to many varying exemplars of a suspects' fingerprints—in some cases, fingerprint databases may only contain a single comparison print, or a ten-print card consisting of fully-rolled prints and 'slapped' prints from the same *person*, and not the same finger (Jain, Nandakumar, & Ross, in press; PCAST, 2016). While Experiment 2 aimed to use prints that fingerprint analysts are likely to encounter in their daily work (e.g., latent crime scene prints presented with fully rolled suspect prints), and the aforementioned task constraints are an important limitation with respect to the experiment's theoretical implications, they also highlight real constraints in attempting to generalise these findings to more applied contexts.

5.4 Broader Implications

Despite the different pattern of results observed with faces and fingerprints, my findings nevertheless help reveal important information about how observers may best familiarise themselves with novel images under different conditions. If these findings were to be replicated or extended in different contexts, they may reveal benefits of image presentation rate beyond face recognition for other domains of perceptual expertise. Given that prior exposure to variation seems to increase recognition performance when controlling for time, the identification decisions of counterfeit investigators, passport officers, various medical practitioners, and other professionals who rely on their perceptual expertise, may benefit from accumulating as much exposure as possible to varying examples within their domain. Future research may look to improve expert identification decisions by optimising the advantages of viewing time and exposure to variation in a range of given fields.

5.5 Conclusion

This thesis is the first to explore how to best familiarise observers with complex, unfamiliar images given a fixed amount of time: should we assess the finer details, or glean the general gist of several similar images? Across two experiments, I establish a new relationship between the RSVP-based ensemble coding literature and the image recognition literature, with the caveat that this relationship may change when presented under different conditions and in other expert domains not explored in this thesis. In Experiment 2, I attempted to boost novices' fingerprint identification performance by increasing their exposure to fingerprint variation in each case, and I found tentative support for current analytical practices, as reported by analysts, during the early stages of their training. My thesis highlights the need to further investigate how to optimally balance the potential advantages of both assessing the details of individual instances, and gaining experience with natural variation, when tasked with recognising familiar or unfamiliar identities and visual categories. As it stands, this thesis provides foundational evidence for the effect of presentation rate that may inform future research on improving the training and identification decisions of professionals in medicine, security, and law enforcement—who are faced with the task of diagnosing or classifying new complex cases based on their previous experience.

References

Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Science, 15*(3), 122-131.

Ariely, D. (2001). Seeing sets: representation by statistical properties. *Psychological Science, 12*(2), 157–162.

Balas, B., & Pearson, H. (2019). The Flashed Face Distortion Effect Does Not Depend on Face-Specific Mechanisms. *Scientific Reports, 9*(1), 1-11.

Balas, B., Gable, J., & Pearson, H. (2019). The Effects of Blur and Inversion on the Recognition of Ambient Face Images. *Perception, 48*(1), 58-71.

Bowden, J., Whitaker, D., & Dunn, M. J. (2019). The role of Peripheral Vision in the Flashed Face Distortion Effect. *Perception, 48*(1), 93-101.

Brooks, L. R. (1987). Decentralized control of categorization: The role of prior processing episodes. In: U. Neisser (Ed.), *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. Cambridge: Cambridge University Press. pp.141–174.

Bruce, V. (1994). Stability from variation: The case of face recognition. The M. D. Vernon Memorial Lecture. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 47*(1)*,* 5–28.

Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology, 77,* 305–327.

Bruce, V., Henderson, Z., Greenwood, K., Hancock, P. J. B., Burton, M. A., & Miller, P. (1999). Verification of face identities from images captured on video. *Journal of Experimental Psychology: Applied, 5*(4), 339–360.

Bruce, V., Henderson, Z., Newman, C., & Burton, A. M. (2001). Matching identities of familiar and unfamiliar faces caught on CCTV images. *Journal of Experimental Psychology: Applied, 7*(3)*,* 207–218.

Burton, A. M., & Bruce, V. (1993). Naming faces and naming names: Exploring an interactive activation model of person recognition. *Memory, 1*(4)*,* 457–480.

Burton, A. M., Jenkins, R., Hancock, P. J., & White, D. (2005). Robust representations for face recognition: The power of averages. *Cognitive Psychology, 51*(3)*,* 256 –284.

Busey T. A., & Parada F. J. (2010). The nature of expertise in fingerprint examiners. *Psychonomic Bulletin and 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.

Campbell, A., & Tanaka, J. W. (2018). Inversion Impairs Expert Budgerigar Identity Recognition: A Face-Like Effect for a Nonface Object of Expertise. *Perception, 47*(6), 647-659.

Cao, Q., Shen, L., Xie, W., Parkhi, O. M., & Zisserman, A. (2018). *VGGFace2: A dataset for recognising faces across pose and age.* Paper presented at the Proceedings - 13th IEEE International Conference on Automatic Face and Gesture Recognition, FG 2018.

Chong, S. C., & Treisman, A. (2003). Representation of statistical properties. *Vision Research, 43*(4), 393-404.

Civile, C., Elchlepp, H., McLaren, R., Galang, C. M., Lavric, A., & McLaren, I. P. L. (2018). The effect of scrambling upright and inverted faces on the N170. *Quarterly Journal of Experimental Psychology, 71*(11), 2464-2476.

Clutterbuck, R., & Johnston, R.A. (2002). Exploring levels of face familiarity by using an indirect face-matching measure. *Perception, 31*(8), 985–994.

Cohen, J. (1988). Statistical Power Analysis for the Behavioral Sciences. New York, NY: Routledge Academic.

Cole, S.A. (2005). More than zero: Accounting for error in latent fingerprint identification. *Journal of Criminal Law and Criminology, 95*(3), 985–1078.

De Fockert, J., & Wolfenstein, C. (2009). Rapid extraction of mean identity from sets of faces. *The Quarterly Journal of Experimental Psychology, 62*(9), 1716-1722.

Dopkins, S., & Gleason, T. (1997). Comparing exemplar and prototype models of categorization. *Canadian Journal of Experimental Psychology, 51*(3), 212-230.

Dunn, J. D., Ritchie, K. L., Kemp, R. I., & White, D. (2019). Familiarity does not inhibit image-specific encoding of faces. *Journal of Experimental Psychology: Human Perception and Performance, 45*(7), 841-854.

Emerick, B., Vanderkolk, J., & Busey, T. (2015). The Policy Implications of Research on Fingerprint Examination Tasks. *Policy Insights from the Behavioral and Brain Sciences, 2*(1), 166-174.

Evans, K. K., Georgian-Smith, D., Tambouret, R., Birdwell, R. L., Wolfe, J. M. (2013). The gist of the abnormal: Above-chance medical decision making in the blink of an eye. *Psychonomic Bulletin & Review, 20*(6), 1-10.

Farah, M. J. (1992). Is an object an object an object? 0 Cognitive and neuropsychological investigations of domain-specificity in visual object recognition. *Current Directions in Psychological Science, 1*(5), 164–169.

Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. (1998). What is ‘‘special’’ about face perception? *Psychological Review, 105*(3), 482–498.

Haberman, J., & Whitney, D. (2009). Seeing the Mean: Ensemble Coding for Sets of Faces. *Journal of Experimental Psychology: Human Perception and Performance, 35*(3), 718-734.

Haberman, J., Harp, T., & Whitney, D. (2009). Averaging facial expression over time. *Journal of Vision, 9*(11), 1-13.

Haberman, J., Lee, P., & Whitney, D. (2015). Mixed emotions: Sensitivity to facial variance in a crowd of faces. *Journal of Vision, 15*(4).

Hanley, J. A., & McNeil, B. J. (1982). The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology, 143*(1), 29-36.

Henry, E. R. (1900). *Classification and Uses of Fingerprints*. London: George Routledge and Sons.

Hochberg, J., & Galper, R. R. (1967). Recognition of faces: An explanatory study. *Psychological Science, 9*(12), 619–620.

Im, H. Y., & Chong, S. C. (2009). Computation of mean size is based on perceived size. *Attention, Perception, & Psychophysics*, *71*(2), 375–384.

Ingvalson, E. M., & Wenger, M. J. (2005). A strong test of the dual-mode hypothesis. *Perception and Psychophysics, 67*(1), 14–35.

Intraub, H. (1979). The role of implicit naming in pictorial encoding. *Journal of Experimental Psychology: Human Learning and Memory, 5*(2), 1–12.

Intraub, H. (1980). Presentation rate and the representation of briefly glimpsed pictures in memory. *Journal of Experimental Psychology: Human Learning and Memory, 6*(1), 1–12.

Intraub, H. (1981). Rapid conceptual identification of sequentially presented pictures. *Journal of Experimental Psychology: Human Perception and Performance, 7*(3), 604–610.

Jackson, M. C., & Raymond, J. E. (2008). Familiarity Enhances Visual Working Memory for Faces. *Journal of Experimental Psychology: Human Perception and Performance, 34*(3), 556-568.

Jain, A. K., Nandakumar, K., & Ross, A. (2016). 50 years of biometric research: Accomplishments, challenges, and opportunities. *Pattern Recognition Letters, 79*, 80-105.

Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence Testing for Psychological Research: A Tutorial. *Advances in Methods and Practices in Psychological Science*, *1*(2), 259–269.

Larson, A. M., & Loschky, L. C. (2009). The contributions of central versus peripheral vision to scene gist recognition. *Journal of Vision, 9*(10), 6, 1-16.

Leib, A. Y., Puri, A. M., Fischer, J., Bentin, S., Whitney, D., & Robertson, L. (2012). Crowd perception in prosopagnosia. *Neuropsychologia, 50*(7), 1698-1707.

McKone, E., & Robbins, R. (2007). The evidence rejects the expertise hypothesis: Reply to Gauthier & Bukach. *Cognition, 103*(2), 331-336.

McKone, E., & Yovel, G. (2009). Why does picture-plane inversion sometimes dissociate perception of features and spacing in faces, and sometimes not? Toward a new theory of holistic processing. *Psychonomic Bulletin and Review, 16*(5), 778-797.

Medin, D.L., & Ross, B.H. (1989). The specific character of abstract thought: Categorization, problem solving, and induction. In: R. Sternberg (Ed.), *Advances in the Psychology of Human Intelligence*. San Diego, CA: Academic Press. pp.189–223.

Megreya, A. M., & Burton, A.M. (2006). Unfamiliar faces are not faces: Evidence from a matching task. *Memory & Cognition, 34*(4), 865-876.

Memon, A., Hope, L., & Bull, R. (2003). Exposure durations: Effects on eyewitness accuracy and confidence. *British Journal of Psychology, 94*(3), 339-354.

Menon, N., White, D., & Kemp, R. I. (2015). Variation in Photos of the Same Face Drives Improvements in Identity Verification. *Perception, 44*(11), 1332-1341.

Morgan, M., Watamaniuk, S., & McKee, S. (2000). The use of an implicit standard for measuring discrimination thresholds. *Vision Research*, *40*(17), 2341–2349.

Murphy, J., Ipser, A., Gaigg, S., & Cook, R. (2015). Exemplar variance supports robust learning of facial identity. *Journal of Experimental Psychology: Human Perception and Performance, 41*(3), 577–581.

Oliva, A. (2005). Gist of the Scene. In L. Itti, Rees, G., & Tsotsos, J. K. (Ed.), *Neurobiology of Attention* (pp. 251-256). San Diego, CA: Elsevier Academic Press.

Oliva, A., & Schyns, P. G. (2000). Diagnostic Colors Mediate Scene Recognition. *Cognitive Psychology, 41*(2), 176-210.

Oliva, A., & Torralba, A. (2001). Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision, 42*(3), 145-175.

Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nature Neuroscience, 4*(7), 739-744.

Perrett, D. I., Oram, M. W., & Ashbridge, E. (1998). Evidence accumulation in cell populations responsive to faces: An account of generalisation of recognition without mental transformations. *Cognition, 67*(1-2), 111-145.

Potter, M. C. (1975). Meaning in visual search. *Science, 187*(4180), 965–966.

Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory, 2*(5), 509–522.

Potter, M. C., & Levy, E. I. (1969). Recognition memory for a rapid sequence of pictures. *Journal of Experimental Psychology, 81*(1), 10-15.

President's Council of Advisors on Science and Technology. (2016). *Forensic Science in Criminal Courts: Ensuring Scientific Validity of Feature-Comparison Methods*. Retrieved from https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/PCAST/pcast_forensic_science_report_final.pdf.

Ritchie, K. L., & Burton, A. M. (2017). Learning faces from variability. *Quarterly Journal of Experimental Psychology, 70*(5), 897-905.

Rossion, B. (2008). Picture-plane inversion leads to qualitative changes of face perception. *Acta Psychologica, 128*(2), 274-289.

Ruiz-Soler, M., & Beltran, F. S. (2006). Face perception: An integrative review of the role of spatial frequencies. *Psychological Research, 70*(4), 273-292.

Sandford, A., Sarker, T., Bernier, T. (2018). Effects of geometric distortions, Gaussian blur, and contrast negation on the recognition of familiar faces. *Visual Cognition, 26*(3), 207–222.

Schyns, P. G. (1998). Diagnostic recognition: Task constraints, object information and their interactions. *Cognition, 67*(1), 147– 179.

Searston, R. A., & Tangen, J. M. (2017a). Expertise with unfamiliar objects is flexible to changes in task but not changes in class. *PLoS ONE, 12*(6), 1-14.

Searston, R. A., & Tangen, J. M. (2017b). The Emergence of Perceptual Expertise with Fingerprints Over Time. *Journal of Applied Research in Memory and Cognition, 6*(4), 442-451.

Searston, R. A., & Tangen, J. M. (2017c). The style of a stranger: Identification expertise generalizes to coarser level categories. *Psychonomic Bulletin and Review, 24*(4), 1324-1329.

Searston, R. A., Thompson, M. B., Vokey, J. R., & French, L. A. (2018). *How low can you go? Detecting style in extremely low resolution images*. Unpublished manuscript.

Standing, L. (1973). Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology, 25*(2), 207-222.

Standing, L., Conezio, J., & Haber, R. N. (1970). Perception and memory for pictures: Single-trial learning of 2500 visual stimuli. *Psychonomic Science, 19*(2), 73-74.

Sunday, M. A., Donnelly, E., & Gauthier, I. (2018). Both fluid intelligence and visual object recognition ability relate to nodule detection in chest radiographs. *Applied Cognitive Psychology, 32*(6), 755-762.

Tanaka, J. W., & Simonyi, D. (2016). The "parts and wholes" of face recognition: A review of the literature. *Quarterly Journal of Experimental Psychology, 69*(10), 1876-1889.

Tanaka, J. W., Curran, T., & Sheinberg, D. L. (2005). The training and transfer of real-world perceptual expertise. *Psychological Science, 16*(2), 145-151.

Tanaka, J., & Gauthier, I. (1997). Expertise in object and face recognition. In R. L. Goldstone, D. L. Medin, & P. G. Schyns (Eds.), *The Psychology of Learning and Motivation, Vol. 36. Perceptual Learning* (pp. 83-125). San Diego, CA, US: Academic Press.

Tangen, J. M., Murphy, S. C., & Thompson, M. B. (2011). Flashed face distortion effect: Grotesque faces from relative spaces. *Perception, 40*(5), 628-630.

Tangen, J. M., Thompson, M. B., & McCarthy, D. J. (2011). Identifying fingerprint expertise. *Psychological Science, 22*(8), 995-997.

Tear, M. J., Thompson, M. B., & Tangen, J. M. (2010). *The importance of ground truth: An open-source biometric repository.* Paper presented at the Proceedings of the Human Factors and Ergonomics Society.

Thompson, M. B., & Tangen, J. M. (2014). The nature of expertise in fingerprint matching: experts can do a lot with a little. *PLoS ONE, 9*(12), 1-23.

Thompson, M. B., Tangen, J. M., & McCarthy, D. J. (2014). Human matching performance of genuine crime scene latent fingerprints. *Law and Human Behavior, 38*(1), 84-93.

Torralba, A. (2009). How many pixels make an image? *Visual Neuroscience, 26*(1), 123-131.

Valentine, T. (1988). Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology, 79*(4), 471-491.

Vogelsang, M. D., Palmeri, & Busey, (2017). Holistic processing of fingerprints by expert forensic examiners. *Cognitive Research: Principles and Implications, 2*(15), 1-12.

Vokey, J. R. (2016). Single-step simple ROC curve fitting via PCA. *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale, 70*(4), 301-305.

Vokey, J. R., Tangen, J. M., & Cole, S. A. (2009). On the preliminary psychophysics of fingerprint identification. *The Quarterly Journal of Experimental Psychology, 62*(5), 1023-1040.

White, D., Burton, A. M., Jenkins, R., & Kemp, R. I. (2014). Redesigning photo-ID to improve unfamiliar face matching performance. *Journal of Experimental Psychology: Applied, 20*(2)*,* 166–173.

Whitney, D., & Leib, A. Y. (2018). Ensemble Perception. *Annual Review of Psychology, 69*, 105-129.

Wong, A. C. N., Palmeri, T. J., & Gauthier, I. (2009). Conditions for face-like expertise with objects: Becoming a ziggerin expert - but which type? *Psychological Science, 20*(9), 1108-1117.

Yin, R.K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology, 81*(1), 141-145.

Ying, H., & Xu, H. (2017). Adaptation reveals that facial expression averaging occurs during rapid serial presentation. *Journal of Vision, 17*(1).

Young, A. W., Hay, D. C., McWeeny, K. H., Flude, B. M., & Ellis, H. D. (1985). Matching familiar and unfamiliar faces on internal an external features. *Perception, 14*(6), 737–746.

Appendix A

Online participant recruitment poster

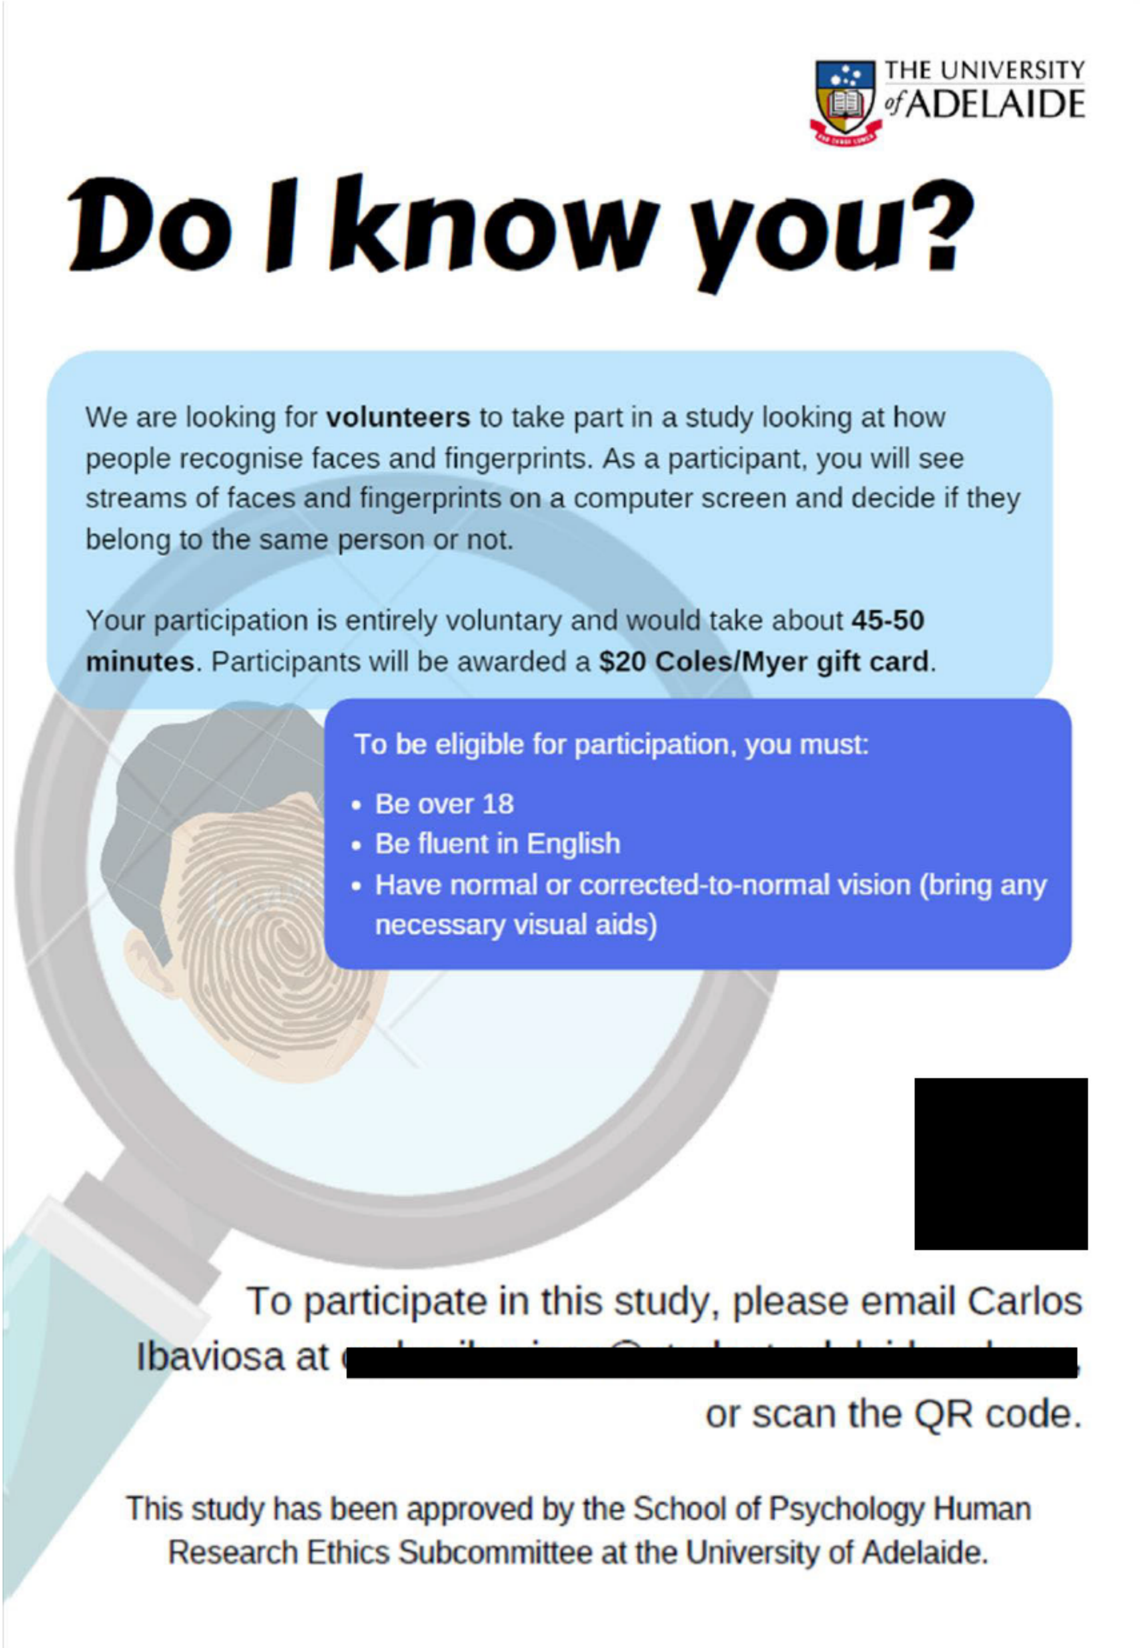

Physical participant recruitment poster

THE UNIVERSITY of ADELAIDE

Do I know you?

We are looking for **volunteers** to take part in a study looking at how people recognise faces and fingerprints. As a participant, you will see streams of faces and fingerprints on a computer screen and decide if they belong to the same person or not. Your participation is entirely voluntary and would take about **45-50 minutes**. Participants will be awarded a **$20 Coles/Myer gift card**.

To be eligible for participation, you must:

- Be over 18
- Be fluent in English
- Have normal or corrected-to-normal vision (bring any necessary visual aids)

To participate in this study, please email Carlos Ibaviosa, at [redacted] or scan the QR code.

This study has been approved by the School of Psychology Human Research Ethics Subcommittee at the University of Adelaide.

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Carlos Ibaviosa
Email: carlos.ibaviosa@student.adelaide.edu.au

Appendix B

Participant information sheet

Participant Information

THE UNIVERSITY of ADELAIDE

PROJECT: The effect of presentation rate on face- and fingerprint-identification
RESEARCHER: Carlos Ibaviosa, School of Psychology, University of Adelaide
SUPERVISOR: Dr Rachel Searston, School of Psychology, University of Adelaide
LOCATION: Hughes Building, North Terrace

PARTICIPATION: You are invited to participate in a study of human perception and cognition. We hope to understand whether varying the rate of image presentation will affect participants' ability to identify whether a new face or fingerprint belongs to the same person or finger as previously seen. This study involves making identification decisions on upright face images, inverted face images, whether similar fingerprints (e.g., all index fingers) belong to the *same person's finger*, and whether different fingerprints (e.g., thumb, index, middle fingers) belong to the *same person* more generally. While potentially useful for security or forensic disciplines, the results from this experiment will not affect any ongoing real-world cases. There are no foreseeable risks in this study that are any different to those of daily life. This study will take approximately one hour, and the researcher will be available nearby throughout the session. After you have completed the experiment, the researcher will discuss the study with you and explain the methodology of the experiment, the variables of interest, and answer any questions you have.

CONSENT: You will receive $20 in the form of a Coles/Myer gift card upon completion of the experiment and will be asked to sign an acknowledgment of payment form. Your participation is completely voluntary. You are free to withdraw from the study at any time and will not be penalised in any way, nor have a negative impact on the participation of any other participants. If, for any reason, you do not want to continue with the experiment, simply let the researcher know. In this event you will still be awarded full payment.

DATA MANAGEMENT: Any information that is obtained from this study will remain entirely confidential and will be kept on a password protected computer with multiple redundant backups. The data from this experiment will be identified by a unique code upon completion. You will not be identifiable by this code, but your performance in this experiment will be recorded. We plan to discuss the results at academic conferences both here and overseas, publish the data in international scientific journals, and store the data in an online open access repository, such as the Open Science Framework, for future meta-analyses and so that other researchers can easily reproduce our work. In any publication, presentation or online record, you cannot be identified.

ETHICS: The study has been approved by the School of Psychology Human Research Ethics Subcommittee at the University of Adelaide (approval number [redacted]. This research will be conducted according to the NHMRC National Statement on Ethical Conduct in Human Research (2007). For any questions about the ethical conduct of the research, please contact Professor Paul Delfabbro, Chair of the Human Research Ethics Subcommittee in the School of Psychology (paul.delfabbro@adelaide.edu.au).

Rachel Searston, PhD
School of Psychology
The University of Adelaide
rachel.searston@adelaide.edu.au

[redacted]

Participant consent form

Consent Form

THE UNIVERSITY of ADELAIDE

PROJECT: The effect of presentation rate on face- and fingerprint-identification
RESEARCHER: Carlos Ibaviosa, School of Psychology, University of Adelaide
SUPERVISOR: Dr Rachel Searston, School of Psychology, University of Adelaide
LOCATION: Room 253, Hughes Building, North Terrace

1. I agree to participate in the project named above, which is for research purposes. The particulars of the project, including details of the tasks, have been explained to me and provided to me via the participant information sheet.

2. I consent to any data gathered from this participation to be used for research purposes and to the data being uploaded and stored in an online public repository (e.g., Open Science Framework), available to other researchers.

3. I consent to any data gathered from this participation to be presented to non-academic bodies (e.g., fingerprint examiners) if the research is deemed be useful to their discipline.

4. I acknowledge that:
   (a) the project is for the purpose of research;
   (b) I have been informed that my involvement is voluntary and that I am free to withdraw from the project at any time without explanation or prejudice and to withdraw any unprocessed data previously supplied;
   (c) the possible effects of the tasks have been explained to me to my satisfaction; and
   (d) I have been informed that the confidentiality of the information I provide will be protected subject to any legal requirements.

5. I understand that:
   (a) my real name or any other identifiable data will not be used in any publications arising from the research without my consent; and
   (b) my participation in the research will have no effect on my academic grades, enrolment or future employment.

Name of participant:
____

Signature: Date:
____

Rachel Searston, PhD
School of Psychology
The University of Adelaide
rachel.searston@adelaide.edu.au

[redacted]

Appendix C

Similarity-matched identity pairs, matched by computing the mean, min, and max properties of each identity


Identity
Mean
Maximum
Minimum
Example
Mean
Maximum
Minimum
Example
13
14
15
16
17
18
19
20
21
22
23
24

Appendix D

Histograms of participants' proportion correct scores for Experiment 1

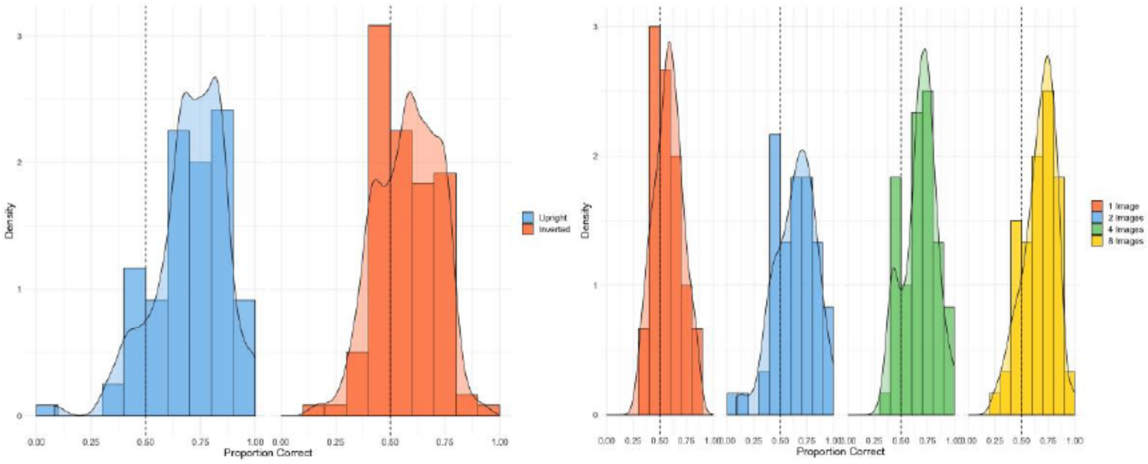

Boxplot of participants' proportion correct scores for Experiment 1

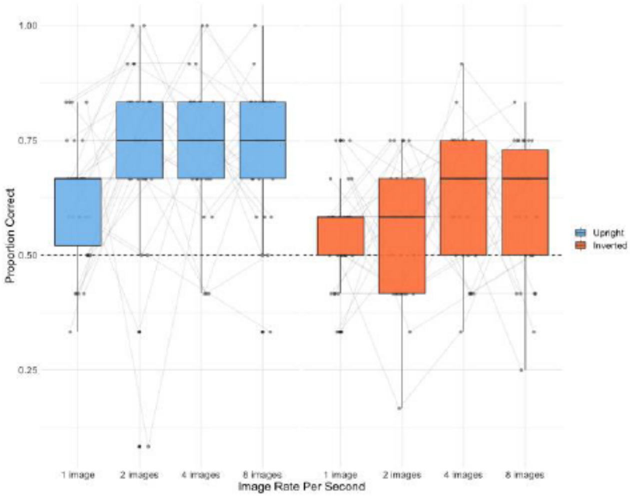

R code chunk and output for participants' proportion correct comparison to chance analysis

```
t.test(
  sum_data_faces$mean_PC[sum_data_faces$Mode == "sim"],
  sum_data_faces$mean_PC[sum_data_faces$Mode == "participant"],
  paired = TRUE,
  alternative = "two.sided"
)

## Cohen's d
sampleMean1<-mean(sum_data_faces$mean_PC[sum_data_faces$Mode == "sim"], na.rm = TRUE)
sampleMean2<-mean(sum_data_faces$mean_PC[sum_data_faces$Mode == "participant"], na.rm = TRUE)
sampleSD1<-sd(sum_data_faces$mean_PC[sum_data_faces$Mode == "sim"], na.rm = TRUE)
sampleSD2<-sd(sum_data_faces$mean_PC[sum_data_faces$Mode == "participant"], na.rm = TRUE)
sampleSDpooled<-sqrt((sampleSD1*sampleSD1+sampleSD2*sampleSD2)/2)
d<-(sampleMean1-sampleMean2)/sampleSDpooled
d
```

```
        Paired t-test

data:  sum_data_faces$mean_PC[sum_data_faces$Mode == "sim"] and
sum_data_faces$mean_PC[sum_data_faces$Mode == "participant"]
t = -10.812, df = 239, p-value < 2.2e-16
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
 -0.1810237 -0.1252263
sample estimates:
mean of the differences
             -0.153125

[1] -1.015922
```

Appendix E

Descriptive statistics for participants' proportion correct scores in Experiment 1

Orientation <fctr>	Image_Rate <ord>	mean <dbl>	variance <dbl>	SD <dbl>
upright	1 image	0.6194444	0.01900543	0.1378602
upright	2 images	0.7333333	0.03611111	0.1900292
upright	4 images	0.7333333	0.02270115	0.1506690
upright	8 images	0.7333333	0.01934866	0.1390995
inverted	1 image	0.5416667	0.01376916	0.1173421
inverted	2 images	0.5472222	0.02092114	0.1446414
inverted	4 images	0.6250000	0.02047414	0.1430879
inverted	8 images	0.6027778	0.02092114	0.1446414

Descriptive statistics for "sim" participants' proportion correct scores in Experiment 1

Orientation <fctr>	Image_Rate <ord>	mean <dbl>	variance <dbl>	SD <dbl>
upright	1 image	0.4722222	0.01979566	0.1406970
upright	2 images	0.4944444	0.01577267	0.1255893
upright	4 images	0.4944444	0.01864623	0.1365512
upright	8 images	0.4694444	0.01268359	0.1126214
inverted	1 image	0.5222222	0.02056194	0.1433944
inverted	2 images	0.4638889	0.01756865	0.1325468
inverted	4 images	0.5083333	0.02507184	0.1583409
inverted	8 images	0.4861111	0.02015485	0.1419678

Repeated measures ANOVA code chunk and output for participants' proportion correct scores in Experiment 1

```
options(contrasts=c("contr.sum","contr.poly"))
ezANOVA(data=sum_data_faces_p, dv=AUC, wid= Participant, within=
.(Orientation, Image_Rate))
ezANOVA(data=sum_data_faces_p, dv=mean_PC, wid= Participant, within=
.(Orientation, Image_Rate))
ezANOVA(data=sum_data_faces_p, dv=mean_conf, wid= Participant,
within= .(Orientation, Image_Rate))
```

	Effect <chr>	DFn <dbl>	DFd <dbl>	F <dbl>	p <dbl>	p<.05 <chr>	ges <dbl>
2	Orientation	1	29	55.465526	3.305118e-08	*	0.15872501
3	Image_Rate	3	87	5.506375	1.645666e-03	*	0.06518006
4	Orientation:Image_Rate	3	87	1.570819	2.022256e-01		0.01837855

	Effect <chr>	W <dbl>	p <dbl>	p<.05 <chr>
3	Image_Rate	0.8374307	0.4262010	
4	Orientation:Image_Rate	0.9716609	0.9772413	

Trend analysis code chunk and output for participants' proportion correct scores

```
sum_data_faces_p$Image_Rate = factor(sum_data_faces_p$Image_Rate,levels = c("1 image","2 images","4 images","8
images"), ordered = TRUE)
contrasts<-contr.poly(4,c(1,2,4,8))
summary(lm(AUC ~ Image_Rate, data=sum_data_faces_p, contrasts = list(f = contrasts)))
summary(lm(mean_PC ~ Image_Rate, data=sum_data_faces_p, contrasts = list(f = contrasts)))
```

```
Call:
lm(formula = mean_PC ~ Image_Rate, data = sum_data_faces_p, contrasts = list(f = contrasts))

Residuals:
     Min       1Q   Median       3Q      Max
-0.55694 -0.08472  0.00278  0.09201  0.35972

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)   0.642014   0.010349  62.036   <2e-16 ***
Image_Rate.L  0.067393   0.020698   3.256   0.0013 **
Image_Rate.Q -0.035417   0.020698  -1.711   0.0884 .
Image_Rate.C -0.006522   0.020698  -0.315   0.7530
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 0.1603 on 236 degrees of freedom
Multiple R-squared:  0.0546,	Adjusted R-squared:  0.04258
F-statistic: 4.543 on 3 and 236 DF,  p-value: 0.004073
```

Treatment-control contrast code chunk and output for participants' proportion correct scores

```
sum_data_faces_p$Image_Rate = factor(sum_data_faces_p$Image_Rate,levels = c("1 image","2 images","4 images","8
images"), ordered = FALSE)
options(contrasts=c("contr.treatment","contr.poly"))
contrasts_treatment1_faces<-contr.treatment(4,c(1,2,4,8))
summary(lm(AUC ~ Image_Rate, data=sum_data_faces_p, contrasts = list(f = contrasts_treatment1_faces)))
summary(lm(mean_PC ~ Image_Rate, data=sum_data_faces_p, contrasts = list(f = contrasts_treatment1_faces)))
```

```
Call:
lm(formula = mean_PC ~ Image_Rate, data = sum_data_faces_p, contrasts = list(f = contrasts_treatment1_faces))

Residuals:
     Min       1Q   Median       3Q      Max
-0.55694 -0.08472  0.00278  0.09201  0.35972

Coefficients:
                 Estimate Std. Error t value Pr(>|t|)
(Intercept)       0.58056    0.02070  28.049  < 2e-16 ***
Image_Rate2 images  0.05972    0.02927   2.040 0.042436 *
Image_Rate4 images  0.09861    0.02927   3.369 0.000882 ***
Image_Rate8 images  0.08750    0.02927   2.989 0.003092 **
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 0.1603 on 236 degrees of freedom
Multiple R-squared:  0.0546,	Adjusted R-squared:  0.04258
F-statistic: 4.543 on 3 and 236 DF,  p-value: 0.004073
```

Appendix F

Histograms of participants' discriminability scores for Experiment 1

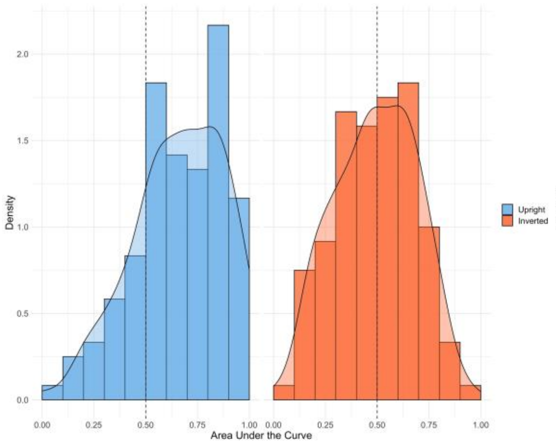

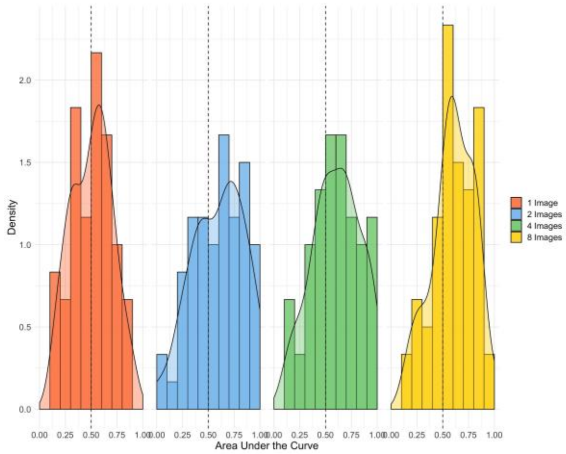

Histograms of "sim" participants' discriminability scores for Experiment 1

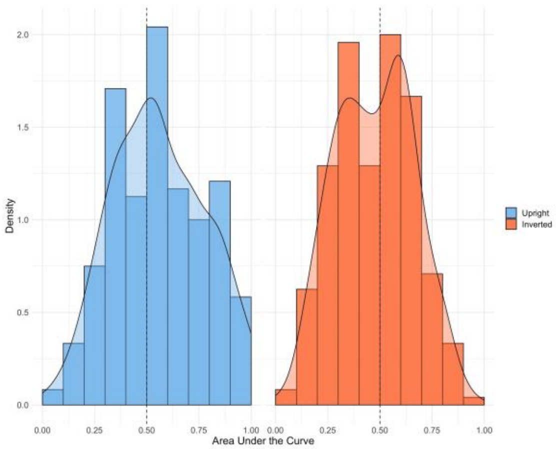

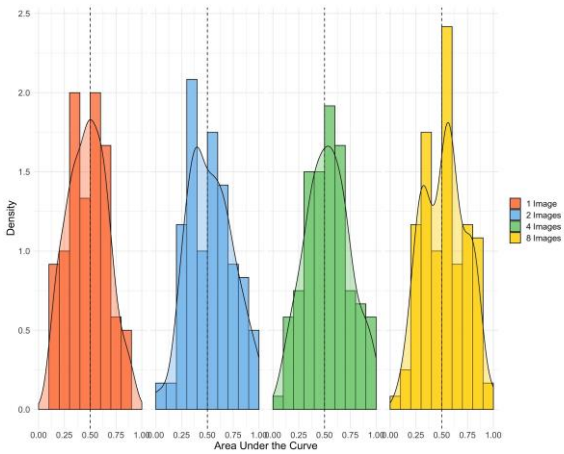

Histograms of participants' confidence scores for Experiment 1

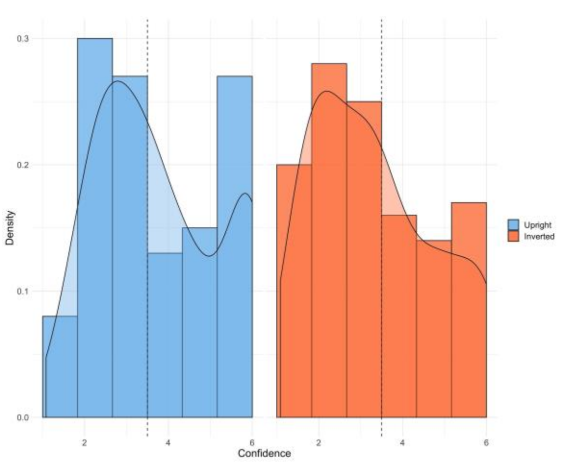

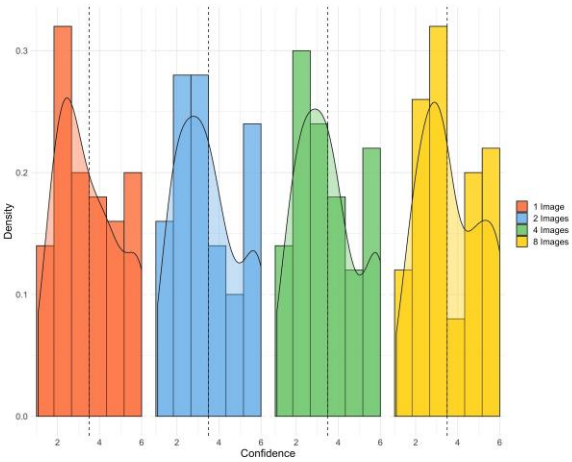

Appendix G

Histograms of participants' proportion correct scores for Experiment 2

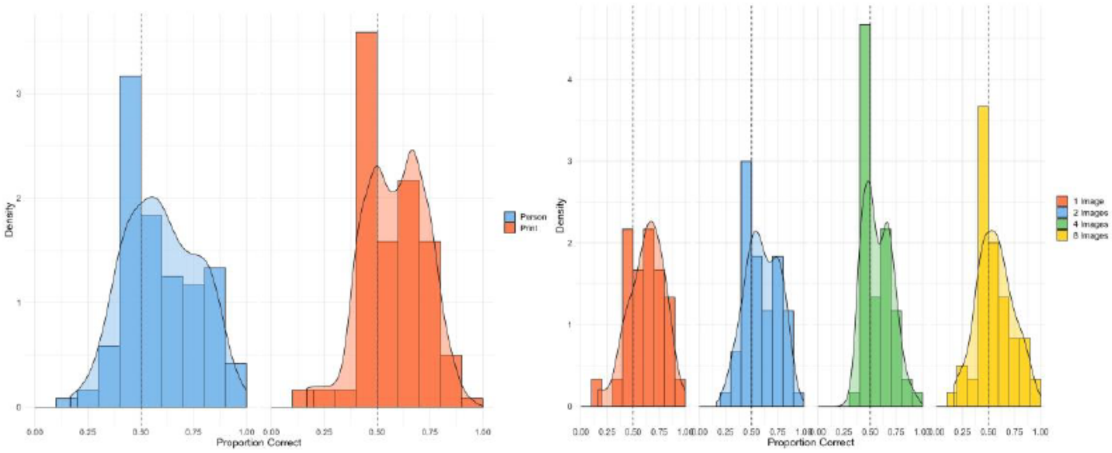

Boxplot of participants' proportion correct scores for Experiment 2

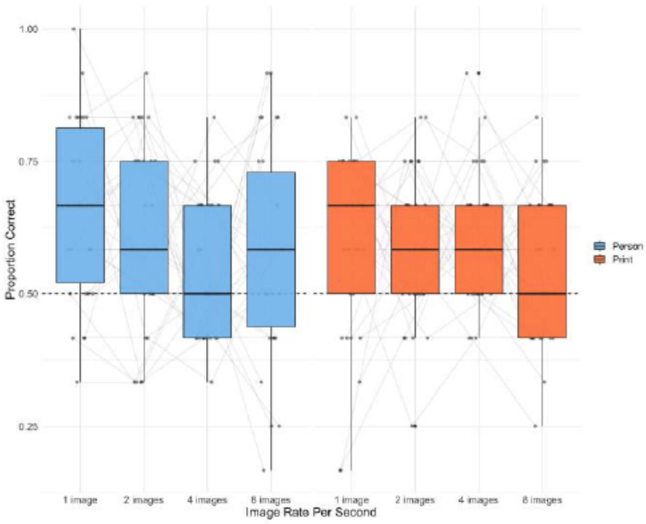

R code chunk and output for participants' proportion correct comparison to chance analysis

```r
t.test(
  sum_data_prints$mean_PC[sum_data_prints$Mode == "sim"],
  sum_data_prints$mean_PC[sum_data_prints$Mode == "participant"],
  paired = TRUE,
  alternative = "two.sided"
)

## Cohen's d
sampleMean1<-mean(sum_data_prints$mean_PC[sum_data_prints$Mode == "sim"], na.rm = TRUE)
sampleMean2<-mean(sum_data_prints$mean_PC[sum_data_prints$Mode == "participant"], na.rm = TRUE)
sampleSD1<-sd(sum_data_prints$mean_PC[sum_data_prints$Mode == "sim"], na.rm = TRUE)
sampleSD2<-sd(sum_data_prints$mean_PC[sum_data_prints$Mode == "participant"], na.rm = TRUE)
sampleSDpooled<-sqrt((sampleSD1*sampleSD1+sampleSD2*sampleSD2)/2)
d<-(sampleMean1-sampleMean2)/sampleSDpooled
d
```

```
	Paired t-test

data:  sum_data_prints$mean_PC[sum_data_prints$Mode == "sim"] and
sum_data_prints$mean_PC[sum_data_prints$Mode == "participant"]
t = -6.9752, df = 239, p-value = 2.975e-11
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
 -0.12245344 -0.06851878
sample estimates:
mean of the differences
            -0.09548611

[1] -0.6325786
```

Appendix H

Descriptive statistics for participants' proportion correct scores in Experiment 2

Specificity <fctr>	Image_Rate <ord>	mean <dbl>	variance <dbl>	SD <dbl>
person	1 image	0.6555556	0.02765006	0.1662831
person	2 images	0.6194444	0.03097861	0.1760074
person	4 images	0.5416667	0.01807950	0.1344600
person	8 images	0.5750000	0.03896073	0.1973847
print	1 image	0.5944444	0.03100255	0.1760754
print	2 images	0.5861111	0.01986750	0.1409521
print	4 images	0.6055556	0.01577267	0.1255893
print	8 images	0.5472222	0.02044221	0.1429763

Descriptive statistics for sim participants' proportion correct scores in Experiment 2

Specificity <fctr>	Image_Rate <ord>	mean <dbl>	variance <dbl>	SD <dbl>
person	1 image	0.4888889	0.02286079	0.1511979
person	2 images	0.5027778	0.01843072	0.1357598
person	4 images	0.4916667	0.02028257	0.1424169
person	8 images	0.5194444	0.01373723	0.1172059
print	1 image	0.4638889	0.02427363	0.1558000
print	2 images	0.5111111	0.01424010	0.1193319
print	4 images	0.4750000	0.01970785	0.1403847
print	8 images	0.5083333	0.02650862	0.1628147

Repeated measures ANOVA code chunk and output for participants' proportion correct scores

```
options(contrasts=c("contr.sum","contr.poly"))
ezANOVA(data=sum_data_prints_p, dv=AUC, wid= Participant, within= .(Specificity, Image_Rate))
ezANOVA(data=sum_data_prints_p, dv=mean_PC, wid= Participant, within= .(Specificity, Image_Rate))
ezANOVA(data=sum_data_prints_p, dv=mean_conf, wid= Participant, within= .(Specificity,
Image_Rate))
```

	Effect <chr>	DFn <dbl>	DFd <dbl>	F <dbl>	p <dbl>	p<.05 <chr>	ges <dbl>
2	Specificity	1	29	0.6322112	0.43300519		0.002165491
3	Image_Rate	3	87	2.1485643	0.09990996		0.024777920
4	Specificity:Image_Rate	3	87	2.4281803	0.07077994		0.022074641

	Effect <chr>	W <dbl>	p <dbl>	p<.05 <chr>
3	Image_Rate	0.8841831	0.6368961	
4	Specificity:Image_Rate	0.9016814	0.7203279	

Appendix I

Histograms of participants' discriminability scores for Experiment 2

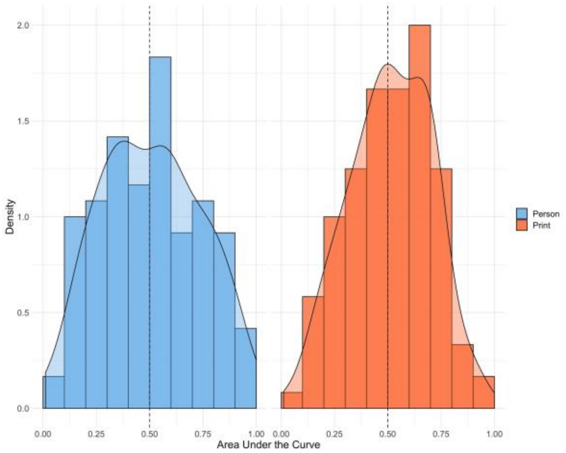

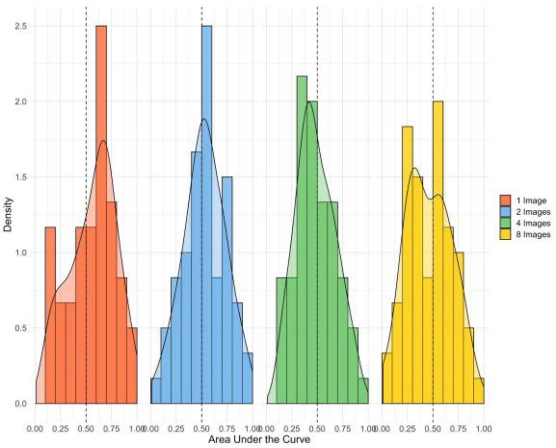

Histograms of "sim" participants' discriminability scores for Experiment 2

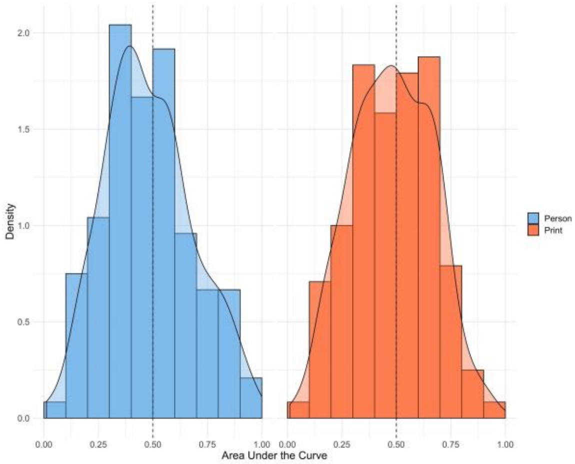

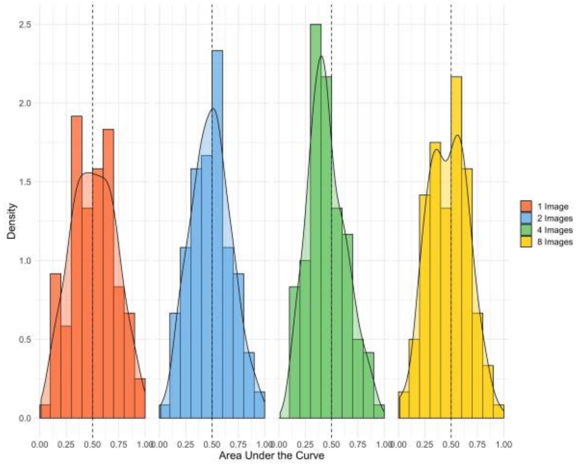

Histograms of participants' confidence scores for Experiment 2