Stark Individual Differences: Face Recognition Ability Influences the Relationship
Between Confidence and Accuracy in a Recognition Test of Game of Thrones Actors

Jesse H. Grabman and Chad S. Dodson

Department of Psychology, University of Virginia

Author Note:

This is the accepted version of a manuscript appearing in the Journal of Applied Research in Memory and Cognition (JARMAC). Please access the published journal article at the following DOI: <https://doi.org/10.1016/j.jarmac.2020.02.007>. Please cite this article in press as:

Grabman, J. H., & Dodson, C. S. (in press). Stark individual differences: Face recognition ability influences the relationship between confidence and accuracy in a recognition test of Game of Thrones Actors. *Journal of Applied Research in Memory and Cognition*. <https://doi.org/10.1016/j.jarmac.2020.02.007>

Correspondence should be addressed to Jesse H. Grabman, Department of Psychology, PO Box 400400, University of Virginia, Charlottesville, VA 22904-4400, email: jhg5dq@virginia.edu, Phone: 434-982-4237

Abstract

Growing evidence suggests face identifications made with high confidence are typically accurate (Wixted & Wells, 2017). However, few studies capture the complexities of real-world face recognition (e.g., non-experimental setting, varied contexts). Moreover, individual differences in face recognition ability may moderate the confidence-accuracy relationship (Grabman et al., 2019). In this study, we reanalyzed data from 32 participants who watched six seasons of the television show *Game of Thrones* for entertainment as the series aired (Devue et al., 2019). Participants provided confidence ratings on a 168-item old-new recognition test of actors and completed a standard test of face recognition ability. Highest confidence ratings were remarkably accurate -- even considering retention-intervals of >3 years and large changes in appearance. However, confidence was generally a better indicator of accuracy for stronger, as compared to weaker, face recognizers.

Keywords: face recognition; confidence; CFMT; metacognition; individual differences

General Audience Summary

How well do we know who we know? Properly assigning certainty to a face identification has real-world implications. For example, police are instructed to ask how confident an eyewitness is in a criminal identification. Research suggests that face identifications made with high confidence are generally accurate (Wixted & Wells, 2017). However, our knowledge about the correspondence between confidence and accuracy is incomplete. Few studies capture the complexities of real-world face recognition. For example, participants typically know that they are in an experiment and see faces only in limited contexts (e.g., a single video). Moreover, whereas the conclusions of most studies are based on and apply to the average individual, face recognition ability ranges considerably in the population, from people with developmental prosopagnosia (i.e., ‘face blindness’) to ‘super-recognizers’. In this study, we reanalyzed data from 32 participants who watched six seasons of the television show *Game of Thrones* for personal entertainment as the series aired (Devue et al., 2019). Participants completed a standard test of face recognition ability, and then viewed 168-pictures of either actors from the show (84 trials) or complete strangers (84 trials). On each trial, their task was to determine whether they had previously ‘seen’ the person in the show, or if the face was ‘new’. After making their decision, they provided a confidence rating using a 5-point scale (1 = *not at all confident*, 5 = *totally confident*). Mirroring previous research, highest confidence ratings were remarkably accurate, even considering responses to actors who had not appeared on the show for > 3-years and/or exhibited large changes in appearance. But, confidence was generally a better indicator of accuracy for stronger, as compared to weaker, face recognizers.

Stark Individual Differences: Face Recognition Ability Influences the Relationship Between Confidence and Accuracy in a Recognition Test of Game of Thrones Actors

Most people have experienced the embarrassment of greeting a stranger as if they were a recent acquaintance. Whether we risk this social faux pas depends on our certainty that we previously encountered this individual. In higher stakes contexts, eyewitness confidence has profound effects on the criminal justice system. Juror decisions are strongly influenced by confidence (Brewer & Burke, 2002), and judges are instructed to use certainty as an indicator of whether to admit the witness's testimony in court (*Neil vs. Biggers*, 1972). The question is how probative confidence is of face recognition accuracy.

In an influential review of the eyewitness literature, Wixted and Wells (2017) found that high confidence identifications are generally accurate. This relationship holds over changes in retention interval (i.e., the amount of time between study and test) (see Wixted, Read, et al., 2016 for a review), exposure duration (i.e., the amount of time a face is viewed at encoding) (e.g., Palmer, Brewer, Weber, & Nagesh, 2013), and a variety of other manipulations (see Wixted & Wells, 2017 for a review). However, there is a compelling need for studies of the confidence-accuracy relationship which capture the richness of the real-world face viewing experience.

The fact that the average person can recognize thousands of unique faces (Jenkins, Dowsett, & Burton, 2018) masks aspects of this task that are remarkably complex. Faces are encountered in a myriad of contexts, often with considerable changes in lighting, orientation, and other characteristics (e.g., hair, age, clothing, etc.). While the majority of people can easily recognize family members and friends in a variety of situations, this task is far more challenging for unfamiliar faces (Kramer, Young, &

Burton, 2018). As some examples of this difficulty, growing literature suggests that minimal disguises (such as sunglasses) can impair face recognition accuracy (Mansour, Beaudry, & Lindsay, 2017; Nguyen & Pezdek, 2017; Righi, Peissig, & Tarr, 2012; Terry, 1994). Moreover, studies in the face matching literature (i.e., indicating whether two simultaneously presented faces are the same person or different people), show that subtle changes in viewing conditions (e.g., photos of the same person taken with different cameras) can substantially decrease matching decision accuracy (see Young & Burton, 2017 for a review).

Given the complexity of real-world face recognition, claims about the value of high confidence are complicated by multiple factors. First, participants in past studies generally knew that they were in an experiment, which potentially alters their face encoding strategies. Second, exposure durations are shorter than those experienced in everyday life (e.g., 90-seconds), and retention-intervals rarely longer than a few weeks (though see Read, Lindsay, & Nicholls, 1998 for an exception). Third, most studies use single-trial designs, which limits conclusions to the small set of faces presented. Finally, there is typically a single context for encoding faces, whereas in practice we must learn to recognize people (often encountered more than once) in varied environments.

Additionally, a largely ignored aspect of the confidence-accuracy relationship in the eyewitness literature is heterogeneity in unfamiliar face recognition ability (Duchaine & Nakayama, 2006). Skill in this domain ranges from people with developmental prosopagnosia (i.e., face blindness), who may have difficulties recognizing even close family members (J. J. S. Barton & Corrow, 2016), to super-recognizers who are actively recruited to police departments for their face-recognition prowess (Ramon, Bobak, &

White, 2019; Russell, Duchaine, & Nakayama, 2009). These differences are highly heritable (Shakeshaft & Plomin, 2015; Wilmer et al., 2010; Zhu et al., 2010), and only weakly associated with general intelligence (Gignac, Shankaralingam, Walker, & Kilpatrick, 2016; Shakeshaft & Plomin, 2015; Wilhelm et al., 2010; Zhu et al., 2010).

Multiple studies show that higher face recognition ability predicts increased accuracy in eyewitness identification tasks (Andersen, Carlson, Carlson, & Gronlund, 2014; Bindemann, Avetisyan, & Rakow, 2012; Morgan et al., 2007). But, only our group has investigated whether this skill influences the probative value of confidence in face recognition tasks. In contrast to previous research documenting a robust confidence-accuracy relationship across a wide range of manipulations, we found that weaker face recognizers are far more likely to make high confidence errors than are stronger recognizers (Grabman, Dobolyi, Berelovich, & Dodson, 2019).

However, there are several aspects that limit the real-world applicability of Grabman et al (2019). Participants viewed static images of faces at encoding and test, which fails to capture the experience of encountering moving people in varied contexts. Moreover, the study used relatively short exposure durations (3 repetitions of 3-seconds) and retention-intervals (up to 1 day). It is possible that the impact of face recognition ability on the confidence-accuracy relationship is minimal with longer exposures or delays. Finally, the stimulus set consisted solely of young adult males, which further limits generalizability.

Given the paucity of studies of the confidence-accuracy relationship under real-world viewing conditions, there are two aims for the current study. The first aim is to determine if the results from a more naturalistic setting mirror those of carefully designed

experiments. The second aim is to assess whether differences in face recognition ability influence the confidence-accuracy relationship using a design that addresses each of the short-comings of our previous study (Grabman et al., 2019).

To accomplish these aims, we leveraged a dataset published by Devue, Wride, and Grimshaw (2019), accessed using the Open Science Framework (OSF) (<https://osf.io/wg8vx>). In this study, participants viewed the first six seasons of the popular television show *Game of Thrones* (GoT) as the series aired, then completed a recognition task of 90 pictures of actors (not in character) intermixed with 90 strangers. Importantly, participants viewed the show for personal entertainment, meaning that all faces are incidentally encoded. Moreover, as Devue et al. (2019) note, there are several additional aspects of GoT that make it an appealing way to study real-world face recognition. Characters are seen in a variety of natural viewing contexts, with often substantial changes in appearance, lighting, clothing, age, and viewpoint. Additionally, screen-time is readily accessible from internet databases, allowing for assessment of exposure duration effects. There are many character deaths throughout the series, resulting in lengthy retention intervals between encoding and test for some actors. Finally, there are over 600 actors listed in the show credits, which provides a substantial face corpus from which to prepare stimuli.

From the standpoint of the current study aims, this dataset offers some additional advantages. Each participant completed a standard test of face-recognition, the Cambridge Face Memory Test+ (CFMT+), and provided confidence ratings for each decision. While the original authors examined associations between these variables and accuracy using correlational analysis, we use calibration curves, which are superior for

assessing confidence-accuracy calibration (Wixted & Wells, 2017). And, for the first time, we analyze the conjunctive effects of confidence and face recognition ability on accuracy under real-world viewing conditions.

Additionally, whereas Grabman et al. (2019) used criminal lineups, participants in Devue et al. (2019) completed an old-new recognition task. Studies using old-new recognition paradigms are generally in alignment with those using lineups, finding a strong relationship between confidence and accuracy for identifications (i.e., ‘old’ responses), and relatively weaker correspondence for non-identifications (i.e., ‘new’ responses) (Tekin & Roediger, 2017; Weber & Brewer, 2004, 2006). The conceptually similar eyewitness ‘show-up’ procedure (where a single face is presented) exhibits a robust relationship between identifications and confidence (Key et al., 2017; Sauerland, Sagana, Sporer, & Wixted, 2018), but also non-identifications and confidence (Sauerland, Sagana, & Sporer, 2012).

Finally, the use of another group’s dataset carries the benefit of reducing ‘researcher degrees of freedom’. If stronger face recognizers continue to make fewer high confidence errors than weaker recognizers in an uncontrolled, naturalistic context then this bolsters claims that there are robust associations between face recognition ability, confidence, and accuracy.

Methods

Participants.

Characteristics of the participants are reported in Devue et al., (2019). Briefly, the results are comprised of 32 participants (20 women and 12 men), aged between 19 and 56 years ($M = 28.7$ years ± 10.5), who completed the task 3-6 months after the end of the

sixth season of GoT. All participants watched six seasons of GoT once, and in order as the show aired, with the exception of some who viewed both Seasons 1 and 2 during the same year. While the sample size is low, the large number of trials per participant ($n = 168$) fits with current recommendations for the logistic mixed effects analysis outlined in the Results section (e.g., McNeish & Stapleton, 2016).

Materials.

Cambridge Face Memory Test + (CFMT+). The CFMT+ is a frequently used test that assesses poor to superior face recognition ability (Russell et al., 2009). Participants must memorize six male faces in three separate orientations. For each trial, previously viewed faces must be selected from an array of the target face and two foils. The test phase proceeds across 102 trials in five increasingly difficult blocks. Difficulty is manipulated with the use of novel images, visual noise filters, different levels of cropping, and (eventually) the use of a profile view with extra levels of noise. Scores can range from 0 – 102 correct responses, but in practice a score of 34 represents random guessing.

Face Stimuli. Extensive details about the generation of the study materials are provided in Devue et al., (2019), with the materials themselves available on the OSF platform (<https://osf.io/wg8vx>). The researchers selected 84 actors from GoT from 15 conditions, consisting of the interaction between retention-interval since last viewing (Season 6, 5, 4, 3, 1/2) and three levels of exposure: ‘lead characters’ [20 – 90 min screen time], ‘support characters’ [9 – 19 min], and ‘bit parts’ [<9 min, but role in story for 1 – 3 episodes]. Additionally, 6 characters categorized as ‘main heroes’ [> 123 min screen time] survived to the end of the sixth season, with the actors serving as training trials for

the task. Ninety pictures of unfamiliar faces were collected to serve as foils (i.e., ‘new’ trials), and “matched the actor set in terms of head orientation, age range, facial expression, attractiveness, presence of make-up, facial hair, or glasses, hairstyle, clothing style, lighting, and picture quality” (Devue et al., 2019). While foils matched the characteristics of the sample of actors as a whole, they were not individually paired to specific actors.

In a similarity manipulation, half of the participants viewed photos of the actors which were as similar as possible to their last appearance on the show (*similar*), while the other half viewed photos that were as different as possible (*dissimilar*). Specifically, the dissimilar condition included “variations in hair length, hair color, hair style, presence of facial hair, glasses, apparent age, and differences in make-up that did not conceal internal features” (Devue et al., 2019). These similarity groups were matched on CFMT+ scores, age, and gender. Due to the scarcity of available photos for ‘bit part’ actors, all participants responded to both *similar* (17 trials) and *dissimilar* (13 trials) pictures for this exposure level, regardless of their assigned similarity condition.

Procedure.

Full details of the procedure are outlined in Devue et al., (2019), so we mention only those pertinent to the present study. Participants completed all tasks on a computer. Following the CFMT+, participants were assigned to a similarity condition, and then started the GoT face recognition task. An easy block consisting of the six ‘main heroes’ and six foils served to practice the task, and was followed by 168 test trials consisting of 84 actors intermixed with 84 foils. Each trial started with a fixation cross (500 ms), followed by a picture stimulus that remained in the center of the screen until the

participant's response or up to 3,000 ms. Participants pressed the 'K' key to indicate they had 'seen' the face before (in GoT or elsewhere), or pressed 'L' to indicate that the face was 'new'. They then provided a confidence rating for this decision using a 5-point scale (1 = *not at all confident*, 5 = *totally confident*).

Results

Data preparation.

Following the lead of the original authors, we discarded 26 trials where participants indicated they recognized an actor from outside of GoT, as well as the training trials (6 'main heroes' + 6 foils per participant). One trial was omitted due to a typo (i.e., score of '2' on accuracy, when only 0 and 1 were possible). We also removed all trials where participants responded in < 300 ms ($n = 371$; 6.9% of total trials), as this is faster than consistent findings on the time to process face identity, along with the additional time needed to perform a keystroke (e.g., Gosling & Eimer, 2011). In total, this left 4,979 responses from 32 participants. We have uploaded the data file used for the analysis to the OSF platform, along with a cleaned version of the original Devue et al. (2019) file that is more conducive toward coding environments (e.g., R, Python) (<https://osf.io/quhsg>).

Table 1 shows the breakdown of the frequency of responses into Hits ("Seen"|Actor), Misses ("New"|Actor), Correct Rejections (CR; "New"|Foil), and False Alarms (FA; "Seen"|Foil) by confidence level and a median split of CFMT+ performance, which we categorize as Weaker Face Recognizers (CFMT+ scores of 52-73) and Stronger Face Recognizers (CFMT+ scores of 74-90). Due to low frequencies of responses in confidence categories 1 and 2, we collapsed these levels to form a single

confidence level ('1-2'). Note that the same information can be visualized in a receiver operating characteristic (ROC) curve by plotting the false alarm rate ($fa/[fa + cr]$) against the hit rate ($hit / [hit + miss]$), across a range of confidence criterion (for details, see Mickes, Flowe, & Wixted, 2012). ROCs for the full data (Figure A1), manipulations (Figure A2), CFMT+ group (Figure A3), and manipulation x CFMT+ group (Figure A4) are presented in Appendix A as an Online Supplement.

Tables 2 and 3 show the frequencies of hits, misses, correct rejections, and false alarms across CFMT+ median split for the exposure duration and retention-interval manipulations, respectively. Due to the single-block design, the same foil counts (i.e., false alarms and correct rejections) are present in all levels of these within-subjects manipulations. To obtain an adequate trial count for the retention-interval contrasts (especially at the upper-end of the confidence scale), we recoded this variable into 'Long Delay' (Seasons 1-3; 34 actors), 'Medium Delay' (Seasons 4-5; 32 actors), and 'Short Delay' (Season 6; 18 actors) conditions, based on comparable discriminability within these time periods. The exposure duration contrast is composed of 'leading actors' (longest exposure; 27 actors), 'supporting actors' (medium exposure; 27 actors), and 'bit parts' (shortest exposure; 30 actors).

Finally, Table 4 shows the counts for the between-subjects similarity manipulation. We removed 'bit part' actors who did not match the condition assigned to the participant (e.g., *dissimilar* 'bit part' photos in the *similar* condition). Note that removing the 'bit part' actors causes a slight difference in the total actor counts (i.e., hits + misses) for the similarity manipulation as compared to the total count for the full sample and the other manipulations.

Is there a strong relationship between confidence and accuracy in a real-world viewing context?

Devue et al., (2019) analyzed the relationship between confidence and overall accuracy using Pearson's correlation coefficients. This analysis found minimal associations between overall accuracy (centered and scaled) and average confidence on accurate trials ($r = .125$), as well as average confidence on inaccurate trials ($r = -.096$).

One issue with defining the confidence-accuracy relationship in terms of overall accuracy is that research generally shows a stronger correspondence between confidence and accuracy for identifications (i.e., 'seen' responses) than non-identifications (i.e., 'new' responses) (e.g., Brewer & Wells, 2006). Separating these response types may reveal more robust relationships than previously reported. Additionally, correlation analysis addresses a fundamentally different question than is typically of interest to applied memory researchers (Juslin, Olsson, & Winman, 1996). Whereas correlation coefficients measure covariation, or the tendency for one variable to increase/decrease as another variable increases/decreases, applied researchers are generally more interested in the accuracy of responses made with a particular level of confidence.

As a concrete example of this difference, imagine that a participant provides the highest possible confidence rating to every trial. The correlation between confidence and accuracy is zero because, regardless of whether accuracy increases/decreases, confidence remains the same. However, despite there being zero correlation, the participant would be perfectly calibrated if they were correct on every trial. Given that the participant used the highest possible confidence rating, we observed their response to be correct 100% of the time.

An easy way to visualize the probative value of confidence is with a calibration curve (Tekin & Roediger, 2017; see also Mickes, 2015). Along the X-axis are progressively increasing confidence values. On the Y-axis is a proportion representing the number of correct items over the sum total of items at this level of confidence (i.e., $\text{correct} / (\text{correct} + \text{incorrect})$). Points are plotted representing *Y-accuracy* at *X-confidence level*. The slope of the lines connecting the points provides additional information. Upward sloping lines signal increasing accuracy with higher levels of confidence, whereas flat lines indicate little difference in predictive power between two confidence ratings.

Figure 1 shows the calibration curves for all identification ('seen') ($\text{hits} / [\text{fa} + \text{hits}]$) and non-identification ('new') ($\text{cr} / [\text{cr} + \text{misses}]$) responses in the GoT task, collapsed across participants. Replicating the eyewitness research, there is clearly a strong positive relationship between higher confidence responses and identification accuracy. The highest confidence level ('5') boasts accuracy rates of 93.5% (95% HDI¹, [89.8, 97.0]), as compared to 53.3% (95% HDI, [46.3, 61.3]) at the lowest level ('1-2'). However, as indicated by the flat line in the right panel, there is little association between confidence and accuracy for non-identifications.

¹ Highest Density Intervals (HDI) are presented for consistency with later analyses. These intervals are based on 10,000 bootstrapped resamples and reflect 95% of values where the probability density is greater than points outside these bounds.

Figure 1. Calibration curves for the full sample of responses. Notably, there is a strong relationship between confidence and accuracy for identifications (left panel), but weaker associations for non-identifications (right panel). The dashed lines at 50% reflect chance accuracy. Error bars reflect 95% HDIs.

Next, we examined the impact of exposure duration (‘leads’ vs. ‘supports’ vs. ‘bit parts’; within-subjects), retention-interval (‘long’ [S1-3] vs. ‘medium’ [S4-5] vs. ‘short’ [S6]; within-subjects), and similarity (‘*similar*’ vs. ‘*dissimilar*’; between-subjects) on the predictive value of confidence ratings. We analyzed each of these manipulations separately (i.e., main effects), as there are too few data-points per cell to assess interactions.

Because foils are not matched to specific actors in this single-block design, the same false alarms and correct rejections must be used in (non-)identification accuracy calculations for each condition. However, before computing accuracy scores, we needed to account for the unequal numbers of actor trials across conditions. Without an adjustment, the same hit/false alarm rates (at a given level of confidence) can produce different calibration curves.

To ensure comparability between conditions, we adjusted the frequency of foil trials to match the frequency of actor trials in each condition. Specifically, we multiplied the frequencies of false alarms and correct rejections in a given condition by the ratio of actor trials to the total number of foil trials ($\text{adj} = \text{actors}_{\text{condition}} / \text{foils}_{\text{total}}$). Thus, at each level of confidence, adjusted identification accuracy = $\text{hits} / [\text{hits} + \text{adj} * \text{fa}]$, and adjusted non-identification accuracy = $\text{adj} * \text{cr} / [\text{adj} * \text{cr} + \text{misses}]$.²

² Why is this adjustment necessary? Imagine that participants respond ‘seen’ to 50% of actor trials and 25% of foil trials with a given level of confidence for both short (18 actors) and medium (32 actors) retention-intervals (i.e., hit rate = 50%, false alarm rate = 25% at this level of confidence). Multiplying out, this gives 18 actors * .50 hit rate * 32 participants = 288 hits vs. 32 actors * .50 hit rate * 32 participants = 512 hits for the short and medium conditions, respectively. Naively, these trials would be compared against 84 foils * .25 false alarm rate * 32 participants = 672 false alarms for both groups. Using the formula for identification accuracy $[\text{hits} / (\text{hits} + \text{fa})]$, we would find accuracy rates of 288 hits / (288 hits + 672 fa) $\approx$ 43% and 512 hits / (512 hits + 672 fa) $\approx$ 76%, for the short and medium retention-intervals, respectively. Despite the same use of the confidence scale across conditions, a difference of $\sim 33\%$ emerges due to disparities in the number of actor trials. Moreover, both groups are far from the identification accuracy rate expected with a study design implementing equal numbers of actor to foil trials, or $.50 / (.50 + .25) \approx 67\%$.

Using the adjustment produces the correct interpretation of identification accuracy. The frequency of hits remains the same for the short (288 hits) and medium (512 hits) retention-intervals. However, the frequency of false alarms (672 fa) is adjusted in the denominator by the number of actor trials in the short (18 actors * 32 participants = 576) and medium (32 actors * 32 participants = 1024) retention-interval conditions over the total number of foil trials (84 trials * 32 participants = 2688 total foils): short ID accuracy = $288 \text{ hits} / (288 \text{ hits} + (576_{\text{short actors}} / 2688_{\text{total foils}}) * 672 \text{ fa}) \approx 67\%$ and medium ID accuracy = $512 \text{ hits} / (512 \text{ hits} + (1024_{\text{medium actors}} / 2688_{\text{total foils}}) * 672 \text{ fa}) \approx 67\%$. There is no longer a difference between the groups, and the identification accuracy rates match the nominal rate expected by equal numbers of actor and foil trials.

Due to computation of an adjusted accuracy statistic, binomial confidence intervals are inappropriate for characterizing the uncertainty around the estimated proportion. For this reason, we bootstrapped 95% highest density intervals (HDIs) around each estimate (R functions and example code are available here: <https://osf.io/quhsg>).³

Figure 2 shows calibration curves for exposure duration (panel a), retention-interval (panel b), and similarity (panel c). The left column displays curves for identifications ($\text{hits}/[\text{fa} \cdot \text{adjustment} + \text{hits}]$), whereas the right column represents non-identifications ($\text{cr} \cdot \text{adjustment}/[\text{cr} \cdot \text{adjustment} + \text{misses}]$).

³ This procedure uses a routine that simulates the experimental design by:

- a) Calculating the adjusted-accuracy at all levels of confidence for each condition.
- b) Randomly sampling participants (with replacement) from the entire dataset (for within-subject only comparisons), or in equivalent number to the size of the between-group comparison (e.g., similarity, CFMT+ score group).
- c) Re-calculating the adjusted accuracy statistic from the resampled data.
- d) Repeating Steps (b) & (c), 10,000 times.
- e) Computing the 95% HDI of the distribution of resampled adjusted-accuracy statistics, using the bayestestR package v. 0.2.5 (Makowski, Ben-Shachar, & Lüdtke, 2019). We use the HDI because it is more agnostic about the symmetry of the resampling distribution than percentile intervals. Specifically, this interval captures modal outcomes, rather than mean outcomes (which are highly influenced by skew). Note that HDI will give the same intervals as a percentile if the resampling distribution is symmetrical.

Figure 2. Calibration curves for subsets of Exposure (panel a), Retention-interval (panel b), and similarity (panel c). The dashed line at 50% represents chance accuracy. Highest confidence identifications (left column) are potentially less reliable for shorter exposure durations ('bit parts'), but impressively accurate across all similarity and retention-interval conditions. In contrast, there is little relationship between confidence and accuracy for non-identifications (right column). Comparing the 95% HDIs, for identifications, we see only a significant difference between highest confidence short vs. long exposure.

Beginning with exposure-duration, we found a strong positive relationship between confidence and accuracy for ‘lead’ (longest exposures) and ‘supporting’ (medium exposures) actor identifications. Moreover, highest confidence identifications exceed 90% accuracy for both conditions (‘lead’ = 96.9%, 95% HDI [94.8, 98.6]; ‘support’ = 92.2%, 95% HDI [87.8, 96.3]). However, there is a much weaker confidence-accuracy relationship for ‘bit part’ (shortest exposure) identifications, as demonstrated by the flat line across confidence levels. In fact, the accuracy for highest confidence responses (53.0, 95% HDI [29.5, 63.6]) is nearly identical to that of the lowest confidence responses, ‘1-2’ (53.0, 95% HDI [37.4, 58.5]). For non-identifications, there are non-overlapping HDIs at nearly all levels of confidence between ‘leads’ and ‘bit parts’, suggesting that ‘new’ responses to longer exposures are generally more accurate overall. But, replicating previous research, confidence is a weak predictor of non-identification accuracy across all exposure-durations.

Next, examining Figure 2b-c, the overlapping 95% HDIs suggest there is limited influence of retention-interval (panel b) or similarity (panel c) on accuracy. Highest confidence identifications are impressively accurate, even when considering the longest retention-intervals (91.6%, 95% HDI [87.5, 95.9]), and dissimilar trials (90.5%, 95% HDI [84.7, 97.0]). And, as with the exposure-duration analysis, there is little probative value of confidence for non-identifications. Importantly, these results do not imply that retention-interval and similarity have limited influence on accuracy at a macro-level. In fact, Devue et al. (2019) found reduced discriminability (i.e., lower d') for longer retention-intervals and dissimilar photos at test. However, it does suggest that when

participants use the same level of confidence, the impact of retention-interval and similarity are minimal.

In summary, the results of the calibration curve analyses generally support previous findings. High confidence is strongly indicative of accuracy for identifications (Wixted & Wells, 2017), though this comes with the caveat that short exposures may reduce reliability. Likewise, there are weaker relationships between confidence and accuracy for non-identifications (Brewer & Wells, 2006). Replicating these findings is encouraging, as the uncontrolled encoding context of watching GoT more closely matches our daily experiences than past experimental studies.

However, these results apply to the average individual. We now turn to the question of whether the confidence-accuracy relationship is influenced by individual-differences in face recognition ability.

Does face recognition ability influence the confidence-accuracy relationship in a real-world viewing context?

To begin, we divided participants by a median-split of CFMT+ score into weaker (scores = 52 – 73) and stronger (scores = 74 - 90) face recognition ability groups. Next, we plotted calibration curves for the totality of the data, separated by face recognition group. The positive slopes in Figure 3 (left panel) suggest that confidence is strongly related to identification accuracy for both groups. But, at the highest confidence level the stronger face recognition group (98.2%, 95% HDI [96.3, 100.0]) is about 9 percentage points more accurate than the weaker face recognition group (89.4%, 95% HDI [84.7, 94.3]). It is important to interpret this result with the caveat that the HDIs at the highest

confidence level are potentially compressed for the Stronger group due to ceiling effects. The same figure (right panel) indicates that there is little association between confidence and non-identification accuracy for both face recognition groups.

Figure 3. Calibration curves for the full dataset, comparing median splits of CFMT+ score. Stronger face recognizers (i.e., CFMT+ = 74 - 90) are more accurate at the highest-level of confidence for identifications than weaker recognizers (CFMT+ = 52-73) (left panel). There are few differences between groups for non-identifications (right panel). Error bars represent 95% HDIs.

Figure 4 shows the identification accuracy calibration curves for stronger and weaker recognizers across varying levels of exposure (panel a), retention-interval (panel b), and similarity (panel c). Starting with the effects of exposure duration, there is a noticeable gap between stronger and weaker recognizers for highest confidence ‘supporting’ actor identifications. Stronger recognizers commit about 10% fewer highest

confidence errors than do weaker recognizers in this exposure condition (97.9%, 95% HDI [95.2, 100.0] vs. 87.3%, 95% HDI [82.1, 92.7], respectively). However, the patterns in the curves are similar across face recognition groups in the other levels of exposure. Though the HDIs do not overlap, highest confidence identifications of ‘lead’ actors are extremely accurate for both stronger and weaker recognizers (99.2%, 95% HDI [98.3, 100.0] vs. 94.6%, 95% HDI [91.4, 97.5], respectively). In contrast, the flat lines for ‘bit part’ identifications signify that the probative value of confidence is limited for shorter exposure durations, regardless of face recognition ability. As reflected by the missing point on the graph, the ‘bit part’ findings should be treated with caution. Table 2 indicates that stronger recognizers never used the highest confidence rating for identifications in this exposure condition.

Next, we found that retention-interval exhibits a graded pattern of differences between stronger and weaker face recognizers. Specifically, the gap between the two groups at the highest confidence level appears to grow with longer retention intervals. Whereas stronger recognizers are about 6 percentage points more accurate than weaker recognizers for short delays (Season 6) (98.7%, 95% HDI [97.3, 100.0] vs. 92.7%, 95% HDI [89.4, 96.3], respectively), the difference increases to 11 points for long delays (Seasons 1-3) (97.7%, 95% HDI [95.4, 100.0] vs. 86.4%, 95% HDI [81.0, 92.1], respectively). As identification accuracy is essentially identical across retention-interval conditions for stronger recognizers, these differences are seemingly fueled by weaker recognizers committing increasing numbers of highest confidence errors over longer delays.

Concluding the identification accuracy results, both groups achieved high levels of accuracy for highest confidence identifications of *similar* faces, though stronger recognizers were about 5 percentage points more accurate than weak recognizers (99.5%, 95% HDI [98.6, 100.0] vs. 94.0%, 95% HDI [92.0, 96.1]). The gap between the two face recognition groups is numerically greater in the *dissimilar* condition, but the overlapping HDIs hint that these differences may not be particularly robust (96.3%, 95% HDI [92.2, 100.0] vs. 87.1%, 95% HDI [79.6, 95.7]).

Finally, Figure 5 displays the non-identification calibration curves for stronger and weaker recognizers across the conditions of exposure duration (panel a), retention-interval (panel b), and similarity (panel c). The curves are essentially flat across all conditions of the manipulations for both groups, which implies that there is limited correspondence between confidence and non-identification accuracy, regardless of face recognition ability.

Figure 4. Calibration curves for identification accuracy across manipulations, broken down by CFMT+ Median split. In almost all cases, there is a strong correspondence between confidence and accuracy. However, stronger face recognizers (CFMT+ scores = 74 – 90) commit fewer highest confidence errors than weaker recognizers (CFMT + scores = 52 - 73). The exception is for ‘bit part’ exposures (panel a, 1st column), where confidence is not predictive of accuracy for either face recognition group. Error bars represent 95% HDIs.

Figure 5. Calibration curves for non-identification accuracy, broken down by CFMT+ Median split. There is little relationship between confidence and accuracy across all manipulations. Additionally, there are few differences between the stronger (CFMT+ scores = 74 – 90) and weaker (CFMT+ scores = 52 – 73) face recognition groups. Error bars represent 95% HDIs.

Does CFMT+ predict identification accuracy above and beyond confidence?

Following Grabman et al. (2019), we fitted logistic mixed effects models to identification responses, consisting of the main effects of Confidence and CFMT+, and their interaction. This method has the advantage of analyzing the effects of CFMT+ as a continuous variable, as opposed to collapsing participants into quantiles (e.g., median split). Moreover, including random intercepts for each participant partials out systematic variance resulting from unmeasured factors that are associated with accuracy (e.g., motivation, fatigue). Combined with visual inspection of the plot of the model output, these analyses clarify whether including CFMT+ significantly improves on predictions of identification accuracy generated from confidence alone.

Due to the single block design, we could not model interactions between the within-subjects conditions (i.e., exposure duration, retention interval), confidence, and CFMT+. Instead, we separately modeled subsets of total false alarms and hits from each condition (e.g., modeling ‘lead’ actor identifications separately from ‘supporting’ actor identifications). We gauged the relative contributions of CFMT+ and Confidence to the model by computing p-values from likelihood ratio tests (LRTs), provided by the *afex* package v. 0.24-1 (Singmann, Bolker, Westfall, & Aust, 2018).

There are no consensus standards for assessing the absolute model fit for logistic mixed effects models. Thus, we employed a combination of three methods to evaluate fit. First, we used the *DHARMa* package (Hartig, 2018, version 0.2.4) to perform Kolmogorov-Smirnov goodness-of-fit tests (KS tests), comparing the observed data to a cumulative distribution of 1,000 simulations from model estimates. Second, we examined residual plots based on deviations between simulated and observed values to check for

signs of model misspecification (i.e., ensuring errors are uniformly distributed for each predicted value). And third, we calculated marginal pseudo- R^2 ($R^2_{\text{GLMM}(m)}$) for fixed-effects, using the *MuMIn* package (Barton, 2018, version 1.42.1). This statistic compares variance accounted for by fixed effects in the model to remaining error, while partialing out the variance accounted for by the random effect structure (i.e., participant intercept).

We fitted mixed logistic regression models to identifications from the full dataset, as well as each individual subset (exposure, retention-interval) using the *lme4* package v. 1.1-21 (Bates, Maechler, Bolker, & Walker, 2014). Models consisted of the main effects and interaction between Confidence (as an ordered factor) and CFMT score (centered and scaled), as well as a random intercept for each participant. In Wilkinson-Rodgers notation (1973) all models took the form: Accuracy $\sim$ Confidence + CFMT + Confidence:CFMT + (1|Participant). However, a slightly different model was needed for the analysis of the similarity manipulation because the between-subject nature of this manipulation required the inclusion of main effects and interactions of the similar vs. dissimilar contrast. None of these terms contributed significantly to the model (all p s > .144), so for brevity we discuss effects in the context of the terms common to all models.

Table 5 shows the chi-square and significance values for each of the three terms (CFMT, Confidence, CFMT:Confidence) across models. In all cases, Kolmogorov-Smirnov tests suggested that the models adequately fit the data (all p s $> .135$), with residual plots showing no signs of major misspecification. Starting with the model of the full

dataset, Figure 6 suggests there is a strong main effect of confidence, such that higher confidence values are associated with greater predicted accuracy. Critically, replicating previous research, the model predicts that weaker face recognizers will make more frequent high confidence errors (as represented by the red line) than stronger recognizers.

Figure 7 reveals that this interpretation persists across most of the manipulation analyses. Examining Table 5, CFMT+ makes a significant contribution to nearly all of the models. In general, the models predict that weaker recognizers will commit more high confidence errors than stronger recognizers. Two notable exceptions come from the analysis of exposure duration. Consistent with the calibration curves, both confidence and

Figure 6. Plot of the logistic mixed effects model for the full dataset, with predictors of CFMT+ score and confidence. Lines represent model estimates, with error shading representing the 95% confidence interval. Notably, high confidence errors (red line) are more frequent when participants are worse face recognizers.

face recognition ability are not particularly predictive of accuracy for ‘bit part’ identifications. Additionally, while confidence is strongly predictive of accuracy for ‘lead actors’, face recognition ability does not contribute significantly to this model.

In summary, the logistic mixed-effects modeling analyses largely reach the same conclusions as the calibration plots: face recognition ability influences the rate of high confidence errors across the dataset as a whole, and across the manipulations of retention-interval and similarity. Concurrently, there are minimal differences in the predictive value of confidence between stronger and weaker recognizers at the longest exposure durations (where confidence is similarly predictive across CFMT+), and the shortest exposure durations (where confidence is not predictive of accuracy across CFMT+).

Figure 7. Plots of the logistic mixed effects model Exposure (panel a), Delay (panel b), and Similarity (panel b) subsets, with predictors of CFMT+ score and confidence. Lines represent model estimates, with error shading representing the 95% confidence interval. Notably, high confidence errors (red line) are more pronounced when participants are worse face recognizers.

Discussion

The ability to evaluate whether we have encountered a person before carries real-world implications. The consequences of misplaced confidence range from the relatively mild (e.g., mistaking one actor for another on a television show) to the devastatingly severe (e.g., convicting an innocent person in a police lineup), yet knowledge about the probative value of confidence is rooted in experiments that inadequately capture the complexity of real-world face recognition. And, excepting results from one previous experiment (Grabman et al., 2019), little is known about how individual differences impact the confidence-accuracy relationship.

Participants in Devue et al (2019) grappled with many of the challenges posed by everyday face recognition. They watched the television show *Game of Thrones* for personal entertainment, incidentally encoding the faces of hundreds of actors. Retention-intervals were much longer than in previous studies (up to 6 years) and characters appeared in varied contexts, often with substantial changes in appearance.

Despite the substantial difficulties imposed by this uncontrolled viewing context, our reanalysis of the data finds that confidence ratings are generally predictive of identification accuracy (e.g., the strong positive slopes for identification in Figure 1). Highest confidence identifications were remarkably accurate for retention-intervals of >3 years (91.6%), and for photos that were as dissimilar as possible from the actors' last appearance in GoT (90.5%).

And, in agreement with most previous research, confidence ratings were less predictive of accurate responses for non-identifications, though a 'show-up' study by Sauerland et al. (2012) is a notable exception. However, Sauerland et al. (2012) differs

from the present study in some key respects. For example, Sauerland et al. (2012) argued that ‘show-ups’ improve the probative value of non-identification confidence compared to lineups, because judgments are constrained to a single face (rather than the ‘average match’ of lineup members). The large differences between Sauerland et al and the present study in number of trials (2 vs. 168) and encoded faces (e.g., 4 faces in a video vs. the full GoT cast) may make this comparison to ‘average’ considerably more difficult. Additionally, participants in Sauerland et al. (2012) could register a ‘don’t know’ response rather than a (non-)identification. Since forced reporting, such as in the present study, causes witnesses to typically reject faces rather than identify them (Weber & Perfect, 2012), it is possible that confidence is less indicative of non-identification accuracy under these conditions. Future studies should decouple these explanations to examine if one (or both) accounts for differences observed for non-identifications.

In contrast to previous research, we found weak correspondence between confidence and accuracy for ‘bit part’ identifications. Importantly, this supports the caveat that high confidence does not always imply high accuracy (Sauer, Palmer, & Brewer, 2019). We believe the most parsimonious explanation for this finding is that performance is essentially at chance for this condition (Lindsay, Read, & Sharma, 1998; Sauer, Brewer, Zweck, & Weber, 2010). As ‘bit part’ actors are nearly indistinguishable from foils (suggesting that memory strength should be comparable), the process that leads participants to assign a given level of confidence should be similar for both.

Taken together, these findings add compelling evidence to claims of a robust association between confidence and accuracy for face identifications (at least, insofar as ‘seen’ faces can be discriminated from ‘new’ faces) (Tekin & Roediger, 2017; Wixted &

Wells, 2017). However, replicating Grabman et al. (2019), high confidence identifications by stronger face recognizers (upper median of CFMT+) were more reliable (98.2% overall) than those made by weaker recognizers (89.4%). The gaps between these two groups were largest when identifying photos at delays >3 years (97.7% vs. 86.4%), for medium exposures (97.9% vs. 87.3%), and (potentially) for dissimilar appearance (96.3% vs. 87.1%). It is surprising that the effects of face recognition ability on the confidence-accuracy relationship are so clear, given that the sample did not include individuals at the extreme tails of CFMT+ performance (i.e., scores closer to chance or perfect). We expect that differences would increase when contrasting wider ranges of ability.

Importantly, face recognition ability is likely one among many individual-differences that influence the confidence-accuracy relationship for face recognition tasks. Kantner and Dobbins (2019) estimated the relative contributions of accuracy and individual-differences (broadly defined) to the confidence ratings participants assigned to recognition trials of paintings and words. The results showed both factors accounting for near-equal variance in confidence for identifications, while individual-differences explained 13-20x more variance in confidence for non-identifications. Martschuk, Sporer, & Sauerland, (2019) found that older eyewitnesses tended to make high confidence errors more frequently than younger participants (see also, Dodson, Bawa, & Krueger, 2007). Additionally, a ‘confidence trait’ may influence how people use confidence scales across many unrelated tasks (e.g., Stankov, Kleitman, & Jackson, 2015). The Wixted & Wells (2017) review convincingly argues for determining the probative value of individual levels of confidence. We believe the next step is to determine the probative value of confidence for *responses by individuals*.

Finally, it is important to acknowledge limitations of the study design. While the encoding context of watching a television show more closely matches daily life than previous studies, it is not perfect. We rarely have the same level of coherence in retrieval cues as we would for a Game of Thrones character (e.g., what is the story arc for the person we met last night?). Ideally future studies will use a combination of experimental and real-world designs, with an eye for capturing the full complexities of face recognition.

In summary, we found that confidence ratings are generally predictive of accuracy in a difficult, naturalistic test of face recognition. Importantly, we replicate that face recognition ability influences the reliability of high confidence ratings, despite using other researchers' data, a different memory paradigm (old-new vs. lineup), and a more ecologically valid context.

Author Contributions

J. H. Grabman developed the study concept and performed primary data analysis under the supervision of C. S. Dodson. J. H. Grabman drafted the manuscript, while C. S. Dodson provided critical revisions.

Open Practices Statement

The original study by Devue et al. (2019) was not formally pre-registered, but data and materials are available at: <https://osf.io/wg8vx/>. We did not formally pre-register predictions, but data and code are posted here: <https://osf.io/quhsg/>.

References

- Andersen, S. M., Carlson, C. A., Carlson, M. A., & Gronlund, S. D. (2014). Individual differences predict eyewitness identification performance. *Personality and Individual Differences, 60*, 36–40. <https://doi.org/10.1016/j.paid.2013.12.011>
- Barton, J. J. S., & Corrow, S. L. (2016). The problem of being bad at faces. *Neuropsychologia, 89*, 119–124. <https://doi.org/10.1016/j.neuropsychologia.2016.06.008>
- Barton, K. (2018). *MuMIn: Multi-model inference*.
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). *lme4: Linear mixed-effects models using Eigen and S4*.
- Bindemann, M., Avetisyan, M., & Rakow, T. (2012). Who can recognize unfamiliar faces? Individual differences and observer consistency in person identification. *Journal of Experimental Psychology: Applied, 18*(3), 277–291. <https://doi.org/10.1037/a0029635>
- Brewer, N., & Burke, A. (2002). Effects of Testimonial Inconsistencies and Eyewitness Confidence on Mock-Juror Judgments. *Law and Human Behavior, 26*(3), 353–364. <https://doi.org/10.1023/A:1015380522722>
- Brewer, N., & Wells, G. L. (2006). The confidence-accuracy relationship in eyewitness identification: Effects of lineup instructions, foil similarity, and target-absent base rates. *Journal of Experimental Psychology: Applied, 12*(1), 11–30. <https://doi.org/10.1037/1076-898X.12.1.11>

- Devue, C., Wride, A., & Grimshaw, G. M. (2019). New insights on real-world human face recognition. *Journal of Experimental Psychology: General*, Vol. 148, pp. 994–1007. <https://doi.org/10.1037/xge0000493>
- Dodson, C. S., Bawa, S., & Krueger, L. E. (2007). Aging, metamemory, and high-confidence errors: A misrecollection account. *Psychology and Aging*, Vol. 22, pp. 122–133. <https://doi.org/10.1037/0882-7974.22.1.122>
- Duchaine, B. C., & Nakayama, K. (2006). The Cambridge Face Memory Test: Results for neurologically intact individuals and an investigation of its validity using inverted face stimuli and prosopagnosic participants. *Neuropsychologia*, 44(4), 576–585. <https://doi.org/10.1016/j.neuropsychologia.2005.07.001>
- Gignac, G. E., Shankaralingam, M., Walker, K., & Kilpatrick, P. (2016). Short-term memory for faces relates to general intelligence moderately. *Intelligence*, 57, 96–104. <https://doi.org/10.1016/j.intell.2016.05.001>
- Gosling, A., & Eimer, M. (2011). An event-related brain potential study of explicit face recognition. *Neuropsychologia*, 49(9), 2736–2745. <https://doi.org/10.1016/j.neuropsychologia.2011.05.025>
- Grabman, J. H., Dobolyi, D. G., Berelovich, N. L. N. L., & Dodson, C. S. (2019). Predicting High Confidence Errors in Eyewitness Memory: The Role of Face Recognition Ability, Decision-Time, and Justifications. *Journal of Applied Research in Memory and Cognition*, 8(2), 233–243. <https://doi.org/10.1016/j.jarmac.2019.02.002>
- Hartig, F. (2018). *DHARMA: Residual diagnostics for hierarchical (multi-level/mixed) regression models*.

Jenkins, R., Dowsett, A. J., & Burton, A. M. (2018). How many faces do people know?

Proceedings of the Royal Society B: Biological Sciences.

<https://doi.org/10.1098/rspb.2018.1319>

Juslin, P., Olsson, N., & Winman, A. (1996). Calibration and diagnosticity of confidence

in eyewitness identification: Comments on what can be inferred from the low

confidence–accuracy correlation. *Journal of Experimental Psychology: Learning,*

Memory, and Cognition, Vol. 22, pp. 1304–1316. [https://doi.org/10.1037/0278-](https://doi.org/10.1037/0278-7393.22.5.1304)

[7393.22.5.1304](https://doi.org/10.1037/0278-7393.22.5.1304)

Kantner, J., & Dobbins, I. G. (2019). Partitioning the sources of recognition confidence:

The role of individual differences. *Psychonomic Bulletin and Review.*

<https://doi.org/10.3758/s13423-019-01586-w>

Key, K. N., Wetmore, S. A., Neuschatz, J. S., Gronlund, S. D., Cash, D. K., & Lane, S.

(2017). Line-up Fairness Affects Postdictor Validity and ‘Don’t Know’ Responses.

Applied Cognitive Psychology, 31(1), 59–68. <https://doi.org/10.1002/acp.3302>

Kramer, R. S. S., Young, A. W., & Burton, A. M. (2018). Understanding face familiarity.

Cognition. <https://doi.org/10.1016/j.cognition.2017.12.005>

Lindsay, D. S., Read, J. D., & Sharma, K. (1998). Accuracy and Confidence in Person

Identification: The Relationship Is Strong When Witnessing Conditions Vary

Widely. *Psychological Science*, 9(3), 215–218. [https://doi.org/10.1111/1467-](https://doi.org/10.1111/1467-9280.00041)

[9280.00041](https://doi.org/10.1111/1467-9280.00041)

- Mansour, J. K., Beaudry, J. L., & Lindsay, R. C. L. (2017). Are multiple-trial experiments appropriate for eyewitness identification studies? Accuracy, choosing, and confidence across trials. *Behavior Research Methods*, 49(6), 2235–2254. <https://doi.org/10.3758/s13428-017-0855-0>
- Martschuk, N., Sporer, S. L., & Sauerland, M. (2019). Confidence of Older Eyewitnesses: Is It Diagnostic of Identification Accuracy? *Open Psychology*, 1(1), 132–151. <https://doi.org/10.1515/psych-2018-0010>
- McNeish, D. M., & Stapleton, L. M. (2016). The Effect of Small Sample Size on Two-Level Model Estimates: A Review and Illustration. *Educational Psychology Review*, 28(2), 295–314. <https://doi.org/10.1007/s10648-014-9287-x>
- Mickes, L. (2015). Receiver operating characteristic analysis and confidence-accuracy characteristic analysis in investigations of system variables and estimator variables that affect eyewitness memory. *Journal of Applied Research in Memory and Cognition*, 4(2), 93–102. <https://doi.org/10.1016/j.jarmac.2015.01.003>
- Mickes, L., Flowe, H. D., & Wixted, J. T. (2012). Receiver operating characteristic analysis of eyewitness memory: Comparing the diagnostic accuracy of simultaneous versus sequential lineups. *Journal of Experimental Psychology: Applied*, Vol. 18, pp. 361–376. <https://doi.org/10.1037/a0030609>
- Morgan, C. A., Hazlett, G., Baranoski, M., Doran, A., Southwick, S., & Loftus, E. (2007). Accuracy of Eyewitness Identification is significantly associated with performance on a standardized test of face recognition. *International Journal of Law and Psychiatry*, 30(3), 213–223. <https://doi.org/https://doi.org/10.1016/j.ijlp.2007.03.005>

- Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R² from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, 4(2), 133–142. <https://doi.org/10.1111/j.2041-210x.2012.00261.x>
- Nguyen, T. B., & Pezdek, K. (2017). Memory for disguised same- and cross-race faces: The eyes have it. *Visual Cognition*, 25(7–8), 762–769. <https://doi.org/10.1080/13506285.2017.1329762>
- Palmer, M. A., Brewer, N., Weber, N., & Nagesh, A. (2013). The confidence-accuracy relationship for eyewitness identification decisions: Effects of exposure duration, retention interval, and divided attention. *Journal of Experimental Psychology: Applied*, 19(1), 55–71. <https://doi.org/10.1037/a0031602>
- Ramon, M., Bobak, A. K., & White, D. (2019). Super-recognizers: From the lab to the world and back again. *British Journal of Psychology*, 0(0). <https://doi.org/10.1111/bjop.12368>
- Read, J. D., Lindsay, D. S., & Nicholls, T. (1998). The relation between confidence and accuracy in eyewitness identification studies: Is the conclusion changing? *Eyewitness Memory: Theoretical and Applied Perspectives.*, pp. 107–130. Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers.
- Righi, G., Peissig, J. J., & Tarr, M. J. (2012). Recognizing disguised faces. *Visual Cognition*, 20(2), 143–169. <https://doi.org/10.1080/13506285.2012.654624>
- Russell, R., Duchaine, B. C., & Nakayama, K. (2009). Super-recognizers: People with extraordinary face recognition ability. *Psychonomic Bulletin and Review*, 16(2), 252–257. <https://doi.org/10.3758/PBR.16.2.252>

- Sauer, J. D., Brewer, N., Zweck, T., & Weber, N. (2010). The effect of retention interval on the confidence-accuracy relationship for eyewitness identification. *Law and Human Behavior*, 34(4), 337–347. <https://doi.org/10.1007/s10979-009-9192-x>
- Sauer, J. D., Palmer, M. A., & Brewer, N. (2019). Pitfalls in using eyewitness confidence to diagnose the accuracy of an individual identification decision. *Psychology, Public Policy, and Law*, Vol. 25, pp. 147–165. <https://doi.org/10.1037/law0000203>
- Sauerland, M., Sagana, A., & Sporer, S. L. (2012). Assessing nonchoosers' eyewitness identification accuracy from photographic showups by using confidence and response times. *Law and Human Behavior*, Vol. 36, pp. 394–403. <https://doi.org/10.1037/h0093926>
- Sauerland, M., Sagana, A., Sporer, S. L., & Wixted, J. T. (2018). Decision time and confidence predict choosers' identification performance in photographic showups. *PLOS ONE*, 13(1), e0190416. <https://doi.org/10.1371/journal.pone.0190416>
- Shakeshaft, N. G., & Plomin, R. (2015). Genetic specificity of face recognition. *Proceedings of the National Academy of Sciences*, 112(41), 12887–12892. <https://doi.org/10.1073/pnas.1421881112>
- Singmann, H., Bolker, B., Westfall, J., & Aust, F. (2018). *afex: Analysis of factorial experiments*.
- Stankov, L., Kleitman, S., & Jackson, S. A. (2015). Measures of the Trait of Confidence. In *Measures of Personality and Social Psychological Constructs*. <https://doi.org/10.1016/B978-0-12-386915-9.00007-3>

- Tekin, E., & Roediger, H. L. (2017). The range of confidence scales does not affect the relationship between confidence and accuracy in recognition memory. *Cognitive Research: Principles and Implications*, 2(1), 1–13. <https://doi.org/10.1186/s41235-017-0086-z>
- Terry, R. L. (1994). Effects of facial transformations on accuracy of recognition. *The Journal of Social Psychology*, 134(4), 483–492.
- Weber, N., & Brewer, N. (2004). Confidence-accuracy calibration in absolute and relative face recognition judgments. *Journal of Experimental Psychology: Applied*, 10(3), 156–172. <https://doi.org/10.1037/1076-898X.10.3.156>
- Weber, N., & Brewer, N. (2006). Positive versus negative face recognition decisions: confidence, accuracy, and response latency. *Applied Cognitive Psychology*, 20(1), 17–31. <https://doi.org/10.1002/acp.1166>
- Weber, N., & Perfect, T. J. (2012). Improving eyewitness identification accuracy by screening out those who say they don't know. *Law and Human Behavior*, Vol. 36, pp. 28–36. <https://doi.org/10.1037/h0093976>
- Wilhelm, O., Herzmann, G., Kunina, O., Danthiir, V., Schacht, A., & Sommer, W. (2010). Individual differences in perceiving and recognizing faces-One element of social cognition. *Journal of Personality and Social Psychology*, 99(3), 530–548. <https://doi.org/10.1037/a0019972>
- Wilmer, J. B., Germine, L. T., Chabris, C. F., Chatterjee, G., Williams, M., Loken, E., ... Duchaine, B. C. (2010). Human face recognition ability is specific and highly heritable. *Proceedings of the National Academy of Sciences*, 107(11), 5238–5241. <https://doi.org/10.1073/pnas.0913053107>

- Wixted, J. T., Read, J. D., & Lindsay, D. S. (2016). The Effect of Retention Interval on the Eyewitness Identification Confidence–Accuracy Relationship. *Journal of Applied Research in Memory and Cognition*, 5(2), 192–203.
<https://doi.org/https://doi.org/10.1016/j.jarmac.2016.04.006>
- Wixted, J. T., & Wells, G. L. (2017). The Relationship Between Eyewitness Confidence and Identification Accuracy: A New Synthesis. *Psychological Science in the Public Interest*, 18(1), 10–65. <https://doi.org/10.1177/1529100616686966>
- Young, A. W., & Burton, A. M. (2017). Recognizing Faces. *Current Directions in Psychological Science*, 26(3), 212–217. <https://doi.org/10.1177/0963721416688114>
- Zhu, Q., Song, Y., Hu, S., Li, X., Tian, M., Zhen, Z., ... Liu, J. (2010). Heritability of the Specific Cognitive Ability of Face Perception. *Current Biology*, 20(2), 137–142.
<https://doi.org/https://doi.org/10.1016/j.cub.2009.11.067>

Table 1

Frequency of Hits, Misses, Correct Rejections (CR), and False Alarms (FA) Categorized by Confidence Level and CFMT+ Median Split

CFMT+	Confidence	hit	miss	fa	cr
Weaker Face Recognizers [52,73]	1-2	77	142	81	193
	3	196	257	149	348
	4	174	212	75	384
	5	236	117	28	141
Stronger Face Recognizers [74,90]	1-2	44	96	25	112
	3	104	189	52	290
	4	103	183	28	349
	5	222	131	4	213

Table 2

Frequency of Hits, Misses, Correct Rejections (CR), and False Alarms (FA), Categorized by Level of Exposure, as well as CFMT+ Median Split.

CFMT+	Confidence	Exposure	hit	miss	fa	cr
Weaker Face Recognizers [52,73]	1-2	'Bit Parts'	28	61	81	193
		'Supports'	25	51		
		'Leads'	24	30		
	3	'Bit Parts'	73	138	149	348
		'Supports'	63	76		
		'Leads'	60	43		
	4	'Bit Parts'	26	115	75	384
		'Supports'	75	60		
		'Leads'	73	37		
	5	'Bit Parts'	13	53	28	141
		'Supports'	62	37		
		'Leads'	161	27		
Stronger Face Recognizers [74,90]	1-2	'Bit Parts'	15	41	25	112
		'Supports'	21	31		
		'Leads'	8	24		
	3	'Bit Parts'	34	97	52	290
		'Supports'	38	62		
		'Leads'	32	30		
	4	'Bit Parts'	19	104	28	349
		'Supports'	43	49		
		'Leads'	41	30		
	5	'Bit Parts'	0	73	4	213
		'Supports'	58	33		
		'Leads'	164	25		

Table 3

Frequency of Hits, Misses, Correct Rejections (CR), and False Alarms (FA) Categorized by Retention-interval, as well as CFMT+ Median Split

CFMT+	Confidence	Delay	hit	miss	fa	cr
Weaker Face Recognizers [52,73]	1-2	Long	33	71	81	193
		Medium	29	44		
		Short	15	27		
	3	Long	77	122	149	348
		Medium	89	91		
		Short	30	44		
	4	Long	55	98	75	384
		Medium	74	70		
		Short	45	44		
	5	Long	72	37	28	141
		Medium	86	56		
		Short	78	24		
Stronger Face Recognizers [74,90]	1-2	Long	18	47	25	112
		Medium	19	32		
		Short	7	17		
	3	Long	43	77	52	290
		Medium	45	79		
		Short	16	33		
	4	Long	36	74	28	349
		Medium	32	78		
		Short	35	31		
	5	Long	68	59	4	213
		Medium	85	44		
		Short	69	28		

Table 4

Frequency of Hits, Misses, Correct Rejections (CR), and False Alarms (FA) Categorized by Actor Similarity, as well as CFMT+ Median Split.

Similarity	CFMT+	Confidence	hit	miss	fa	cr
Similar	Weaker Face Recognizers [52,73]	1-2	28	62	27	92
		3	54	102	58	175
		4	57	87	36	181
		5	96	73	18	122
	Stronger Face Recognizers [74,90]	1-2	23	39	16	54
		3	41	82	21	144
		4	24	85	5	162
		5	62	64	3	122
Dissimilar	Weaker Face Recognizers [52,73]	1-2	38	48	54	101
		3	105	89	91	173
		4	106	61	39	203
		5	136	12	10	19
	Stronger Face Recognizers [74,90]	1-2	16	32	9	58
		3	51	62	31	146
		4	72	45	23	187
		5	160	30	1	91

Note. Trial counts do not match Table 1 because of the removal of ‘bit part’ actors who did not match the condition assigned to the participant (e.g., *dissimilar* ‘bit part’ photos in the *similar* condition).

Table 5

Output for Logistic Mixed Effects Models Examining the Main Effects of CFMT+, Confidence, and their Interaction

		CFMT		Confidence		CFMT:Confidence		R ² _{GLMM(m)}
<i>Model</i>		<i>X²(1)</i>	<i>p</i>	<i>X²(2)</i>	<i>p</i>	<i>X²(3)</i>	<i>p</i>	
All Data		6.91	.009	197.64	< .001	6.51	.089	.221
Delay	<i>Long</i>	6.43	.011	96.80	< .001	6.26	.100	.184
	<i>Medium</i>	4.66	.031	133.48	< .001	5.67	.128	.230
	<i>Short</i>	7.33	.007	195.55	< .001	6.12	.106	.323
Exposure	<i>Bit Parts</i>	0.24	.627	2.86	.414	3.27	.351	.013
	<i>Supports</i>	7.06	.008	99.85	< .001	5.50	.139	.189
	<i>Leads</i>	2.99	.084	326.96	< .001	7.61	.055	.408
Similarity		7.68	.006	64.26	< .001	4.46	.216	.279

Note. Models are fitted to the full dataset, as well as subsetting conditions of exposure, retention-interval, and similarity. Significance values are computed by Likelihood ratio tests. CFMT+ contributes to all, except in the cases of the shortest ('bit parts') and longest ('leads') exposures.