The language of accurate and inaccurate eyewitnesses

Travis M. Seale-Carlisle^{1,2}, Jesse H. Grabman³, & Chad S. Dodson³

¹University of Aberdeen, School of Psychology

²Duke University, Wilson Center for Science and Justice, School of Law

³University of Virginia, Department of Psychology

Correspondence: Travis M. Seale-Carlisle, University of Aberdeen, School of Psychology,
tcarlisl1688@gmail.com

Acknowledgements: We thank David Dobolyi, Gurpreet Reen, and Melissa Colloff for their thoughts on an earlier version of this manuscript. We also thank Ian Dobbins and Justin Kantner for their helpful scripts. Lastly, we thank Caroline Levenson, Kelis Johnson, Emma Balsam, and Annie Han for their valuable assistance.

The authors have no conflict of interest.

The data for this study can be found here: <https://osf.io/s64dj/>

Abstract

This paper uses machine-learning techniques to examine people's use of verbal expressions of confidence. Across the field of academic psychology, it is often assumed that such statements reflect the same underlying information as numeric confidence ratings. We show that verbal confidence is not redundant with numeric confidence, but instead contributes unique diagnostic value in predicting the accuracy of a response. We use eyewitness confidence in a lineup identification as our model paradigm. There is potentially great applied value in developing a machine-learning algorithm that can predict eyewitness identification accuracy, such as by reducing false convictions. To this end, we applied a machine-learning methodology to investigate the natural language of accurate and inaccurate eyewitnesses. This method revealed that verbal confidence statements provide rich diagnostic information about the likely accuracy of eyewitness identifications. Moreover, verbal confidence statements provide unique diagnostic information that traditional indicators of identification accuracy such as numeric confidence ratings and response times do not provide. However, the diagnostic value of an eyewitness confidence statement depends in part on the face recognition ability of the eyewitness: the natural language of strong face recognizers is more diagnostic than the natural language of weak face recognizers. These results are theoretically interesting, but from an applied perspective, this machine-learning methodology may prove useful to those in the criminal justice system that must evaluate eyewitnesses' verbal confidence statements.

Keywords: eyewitness identifications, language and memory, machine learning, discriminability, confidence-accuracy relationship, positive predictive value, confidence

The language of accurate and inaccurate eyewitnesses

A considerable body of empirical work suggests that eyewitnesses have insight into the likely accuracy of their identifications. For example, decades of laboratory research (Wixted & Wells, 2017) and more recent field research (Wixted, Mickes, Dunn, Clark, & Wells, 2016) show that there can be a strong relationship between eyewitness confidence and identification accuracy. This relationship between confidence and accuracy can be especially strong when an eyewitness identifies the suspect from an initial, uncontaminated, and properly administered lineup (Wixted, Mickes, Clark, Gronlund, & Roediger, 2015; Wixted & Wells, 2017). If these conditions are met, a suspect identified with high confidence can have a high likelihood of being the person who actually committed the crime (i.e., a high likelihood of being guilty). That likelihood can be substantially reduced if a suspect is identified with low confidence instead.

The laboratory experiments that have led to our understanding of the generally strong relationship between confidence and accuracy have required participants to express their confidence in their identification decisions by selecting a number on a rating scale, such as 80% from a 0% to 100% scale (e.g., Dodson & Dobolyi, 2016; Tekin, Lin, & Roediger, 2018; Tekin & Roediger, 2017; Wixted & Wells, 2017). Outside of the laboratory and in the real world, eyewitnesses in police investigations typically express their confidence in their own words. For example, in the United States, federal law enforcement agencies are required to obtain an initial statement of confidence in the eyewitnesses' own words (Yates, 2017). Likewise, in the United Kingdom, eyewitnesses are asked for a statement of confidence during the administration of a showup procedure (but, curiously, not during the administration of a lineup procedure). The guidelines require that, "... they shall be asked how sure they are that the

person is the one they saw on the earlier occasion” (PACE Code Annex D, 1984; 2017, p. 47). In Sweden, however, eyewitnesses are first asked for a statement of confidence in their own words and then are asked to rate their confidence on a scale of 0-100% (Fitzgerald, Rubinova, & Juncu, 2021).

While there is much information about the value that confidence holds when a numeric rating scale is used, less is known about the diagnostic value of verbal confidence statements (Behrman & Richards, 2005; Cash & Lane, 2017; 2021; Dobolyi & Dodson, 2018; Dodson & Dobolyi, 2015; Grabman, Dobolyi, Berelovich, & Dodson, 2019; Klobuchar, Steblay, & Caliguiri, 2006; Mansour, 2020). Wixted and Wells (2017) highlighted this as a priority for future research, about which they asked:

Should a confidence statement be taken in the witness’s own words (as in Klobuchar et al., 2006), or should confidence be recorded using an explicit 3-point rating scale (as in Wixted et al., 2016) – or should a 100-point scale be used? Given the clear information value of initial confidence, this issue seems important to pursue. (p. 54)

We agree that this is an important issue and we aim to answer this question in this paper. We briefly review the relevant literature below.

Behrman and Richards (2005) in Study 1 analyzed verbal confidence statements from eyewitnesses to crimes who made an identification from a lineup as part of an ongoing criminal investigation. Their coding scheme put the statements into low, medium, or high confidence bins (see their Table 1). Although the suspects’ guilt or innocence was unknown in these cases, every filler placed in the lineup was innocent and every filler identification was therefore a

known error. Behrman and Richards found that eyewitnesses rarely identified fillers with high confidence; most eyewitnesses identified fillers with medium or low confidence instead (see their Table 3). Results like these are consistent with other police department field studies (e.g., Wixted et al., 2016) and suggest that a verbal confidence statement, much like a confidence rating from a rating scale, is indicative of identification accuracy.

Mansour (2020) investigated whether verbal confidence statements and numeric confidence ratings yield a similarly strong confidence-accuracy relationship. In Experiment 1, participants viewed a perpetrator (i.e., the target) commit a mock crime and later a 6-person, target-present or target-absent simultaneous photo lineup. Participants either identified a lineup member or selected the “not there” option. After participants made an identification decision, they expressed their confidence in one of four ways. In the “scale only” condition, participants provided a confidence rating using a scale ranging from 0% – not confident at all – to 100% – completely confident. In the “own words” condition, participants expressed their confidence in their own words. In the “why” condition, participants explained the reasons why they made their identification decision. Lastly, in the “selection” condition, participants selected a pre-written confidence statement from a collection of confidence statements. Mansour binned confidence ratings less than 50%, 50-79%, and 80-100% into low, medium, and high confidence bins, respectively. Verbal confidence statements collected from the “own words”, the “why”, and the “selection” conditions were placed into low, medium, or high confidence bins based on the Behrman and Richards (2005) coding scheme. Mansour then conducted confidence-accuracy characteristic (CAC) analysis (Mickes, 2015), which simply plots suspect identification accuracy for each level of confidence, and found that the confidence-

accuracy relationship was similarly strong across all four conditions (though the relationship was strongest in the “own words” condition). Participants in Experiment 2 saw a sequential lineup instead of a simultaneous lineup and were randomly assigned to either the “scale only” condition or the “own words” condition. CAC analysis revealed a nearly identical confidence-accuracy relationship regardless of whether participants expressed their confidence in their own words or selected their confidence from the rating scale.

Dobolyi and Dodson (2018) investigated whether the diagnostic value of numeric confidence ratings depends on how participants justify their identification decisions. In Experiment 1, participants studied a list of faces and were later presented a series of 6-person, simultaneous lineups half of which contained a target from the study list. Participants either identified a lineup member or rejected the lineup by clicking the “not present” option. After making an identification decision, participants were asked “in your own words, please explain how certain you are in your response” and then were asked to “provide specific details about why you made your written response above” (p. 546). Participants then translated their verbal confidence statements onto a six-point percentage scale ranging from 0% – not certain at all – to 100% – completely certain. A majority of participants (80%) expressed their confidence entirely in their own words without using any numbers. Dobolyi and Dodson placed each of these participants’ verbal confidence statements into one of six bins: familiarity (e.g., “He just seems familiar”), recognition (e.g., “I am positive that is him”), one observable feature (e.g., “I remember his eyes”), many observable features (e.g., “I remember his eyes and hair”), one unobservable feature (e.g., “He looks like my neighbor”), and mixed (which contained statements that fell into multiple categories).

To analyze the relationship between the type of confidence statement and the accuracy of an identification, Dobolyi and Dodson (2018) fit a generalized linear mixed effects model to the “chooser” data, which included correct identifications from target-present lineups as well as filler identifications from target-present and target-absent lineups. The model predicted a strong confidence-accuracy relationship when participants justified their identification by referring to an unobservable feature (e.g., “I am positive that is him”). The model predicted a comparably strong confidence-accuracy relationship when participants justified their identification by referring to one or many observable features (e.g., “I recognize his eyebrows”). However, when participants justified their identification based on familiarity (e.g., “He just seems familiar”), the model predicted a much weaker confidence-accuracy relationship (though there were only a few of these types of identifications made with high confidence). Dobolyi and Dodson also conducted CAC analysis (Mickes, 2015), which does not require any model fitting and focuses exclusively on the accuracy of an identified suspect (i.e., filler identifications from target-present lineups are excluded from this analysis), and found nearly identical results.

The findings from these studies suggest that eyewitness confidence statements collected from an initial and properly administered lineup are a strong indicator of identification accuracy. However, it is unclear whether eyewitness confidence statements are more or less indicative of identification accuracy compared to numeric confidence ratings or other indicators of identification accuracy such as response times. It is generally assumed that eyewitness confidence statements capture the same underlying information as numeric confidence ratings, but eyewitness confidence statements may provide unique diagnostic

information that numeric confidence ratings and other indicators of identification accuracy do not provide. The current study addresses these gaps in the literature.

The Current Study

We use a fine-grained analysis to assess the diagnostic value of eyewitness confidence statements. We follow the footsteps of Dobbins and Kantner (2019) and use their machine-learning methodology, which they applied to list-learning recognition data, to translate the language of an eyewitness into a classifier. Using this classifier, we capture the diagnostic value contained within eyewitness confidence statements in richer detail than the broad categorizations used in past research (e.g., Behrman & Richards, 2005; Dobolyi & Dodson, 2018; Mansour, 2020).

To illustrate how the classifier works, we apply it to laboratory data from an experiment investigating face recognition ability and the strength of the confidence-accuracy relationship (Grabman et al., 2019). We chose these data because a large number of eyewitness confidence statements were collected (in order to train the classifier) and because we wanted to explore the relationship between face recognition ability and the ability to provide a diagnostic confidence statement.

To preface our results, we find that eyewitness confidence statements contain rich diagnostic information, replicating Dobolyi and Dodson (2018) and others (e.g., Behrman & Richards, 2005; Mansour, 2020). Even more noteworthy, we find that eyewitness confidence statements contain unique diagnostic information that numeric confidence ratings and other indicators of identification accuracy, such as response times, do not provide. We also find that

the diagnostic value of an eyewitness confidence statement depends in part on the face recognition ability of the eyewitness: the natural language of strong face recognizers is more diagnostic than the natural language of weak face recognizer. This finding is consistent with recent work in this area of research (e.g., Gettleman, Grabman, Dobolyi, & Dodson, 2020; Grabman & Dodson, 2020; Grabman et al., 2019) and suggests that strong face recognizers are better able to select the words that best describe the strengths of their memories.

These results have theoretical implications, but there are two potential applied implications as well. First, these results show that collecting an eyewitness confidence statement and a numeric confidence rating provides more information about the likely accuracy of an identification than either provide on their own. Both types of confidence expressions can therefore be useful to the criminal justice system. Second, machine-learning technology may help people accurately interpret eyewitness confidence statements. Decades of laboratory research demonstrate numerous problems in understanding someone's verbal expression of confidence (e.g., Budescu, Por, & Brumell, 2014). We know of several cognitive biases that cause people to misinterpret eyewitness confidence statements specifically (e.g., Cash & Lane, 2017, 2020; Dodson & Dobolyi, 2015, 2017; Grabman, Cash, Slane, & Dodson, *in press*; Grabman & Dodson, 2019; Slane & Dodson, 2019). If actors within the criminal justice system misinterpret an eyewitness confidence statement, they can jeopardize innocent suspects and safeguard guilty suspects. Machine-learning technology may minimize this problem if it can help interpret eyewitness confidence statements as intended by eyewitnesses.

Method

Very few studies have documented the natural language that accompanies eyewitness identifications. Grabman et al. (2019) recently conducted one of the few studies that has. In this study, participants ($N = 569$; all of which self-reported as being White/Caucasian and 68.5% identified as female) were recruited online via Amazon Mechanical Turk. Participants ranged in age from 18 to 50 years old ($M = 31.66$, $SD = 6.08$) and were compensated for completing the study.

Briefly, participants studied a list of 12 target photos. Each photo depicted a person wearing street clothes with a casual facial expression in a “head and shoulders” shot. Half of the targets were Black and half were White. Each photo was shown for three seconds and, after either a 5-minute distractor task or a 24-hour delay, participants were shown twelve 6-person, simultaneous lineups. Half of the lineups consisted of Black faces and the other half consisted of White faces. Half of the lineups were target-present and half were target-absent. Target-present lineups consisted of five filler images and a new image of a previously seen target. Target-absent lineups consisted of six filler images. Participants had the option of identifying a lineup member or selecting the “not present” option if they did not recognize anyone. The materials used in this study were the same materials used by Dobolyi and Dodson (2013; 2018) and were thoroughly evaluated for lineup fairness.

After each lineup decision, participants were instructed to, “[i]n your own words, please explain how certain you are in your response” and, following this, were instructed to provide

“specific details about why you made your written response above” (p. 546).¹ Participants then translated their written expressions of certainty onto a 6-point rating scale ranging from 0% (not certain at all) to 100% (completely certain) in increments of 20%.

After participants saw the series of lineups, they completed the Cambridge Face Memory Test (Duchaine & Nakayama, 2006) to assess their face recognition ability. The ability to recognize a previously seen but otherwise unfamiliar face varies considerably between individuals (for a review, see Wilmer, 2017). Those with exceptionally strong face recognition abilities, often termed “super recognizers”, routinely achieve stellar performance in perceptual discrimination and face recognition tasks (Russell, Duchaine, & Nakayama, 2009). In contrast, those with exceptionally poor face recognition abilities, often termed “developmental prosopagnosia”, are effectively face-blind (e.g., Barton & Corrow, 2016). On study trials of the CFMT, participants attempt to memorize six faces in three separate orientations (e.g., left profile view, front view, right-profile view). On test trials, participants see one target face mixed amongst two fillers. Participants must indicate which one of the three faces is the target. There are 72 test trials in total. Thus, participants can receive a score ranging from 0, if every decision was incorrect, to 72, if every decision was correct. The median CFMT score in Grabman et al. (2019) was 61 with a median absolute deviation of 8.9.

Grabman et al. (2019) were primarily interested in investigating factors that systematically cause eyewitnesses to misidentify someone with a high rating of confidence.

¹ The confidence statements we use in this study combines the text from these two responses. Grabman et al. (2019) found that 78% of participants provided confidence statements entirely in their own words without using any numbers.

While it is true that a high rating of confidence is generally a good indicator that the identification is correct (Wixted & Wells, 2017), it is possible that certain conditions increase the likelihood of a high confidence misidentification. Indeed, they found, for example, that incorrect identifications made with a high rating of confidence were more likely to occur if the eyewitness was a poor face recognizer, if the decision to identify someone took a long time, or if the eyewitness justified their identification decision by describing a feeling of familiarity.

Our focus will be on the words participants used to explain their certainty in their identification decisions. Does this verbal expression of certainty contain predictive value about the accuracy of the identification? If so, is that predictive value partially independent of the predictive value of numeric confidence or other indicators of identification accuracy? The following sections describe a fine-grained analysis that answers these questions.

Building the Classifier

There are many ways to build a language classifier. We followed the “bag-of-words” approach used by Dobbins and Kantner (2019). This approach treated each individual word as a potentially predictive feature. For the first step, we created a document-feature-matrix (DFM) in which each row is a confidence statement and each column is a word present in the entire collection of statements (i.e., the corpus). To increase classifier generalizability, we removed words from the DFM that were present in less than .5% of statements. Although it is common to also remove “stop words” in the DFM (e.g., “and”, “but”, “very”, etc.), we decided to keep them because some of these words may be psychologically relevant. For example, participants may use the stop word “very” as an amplifier in a verbal expression of confidence (e.g., “very

confident”), and the stop word “but” may indicate a failure to retrieve (e.g., Dobbins & Kantner, 2019). We also did not reduce words down to their stem (e.g., reducing “remembered” down to “remember”) because the slight differences in these words may be psychologically relevant, especially for participants discussing their memories of the past.

For the second step, we counted the number of times each word in the corpus appeared in a statement. We did this for each confidence statement (i.e., for every row in the DFM) and for each word in the corpus (i.e., for every column in the DFM) until every cell in the DFM contained a number. We then trained a logistic regression model to use the resulting word predictors to distinguish between correct and incorrect identification decisions. We fit the logistic regression model to the data under the LASSO (Least Absolute Shrinkage and Selection Operator) constraint (Tibshirani, 1996). Given the large number of predictors available (in this case, each word in the corpus), we sought to avoid the model overfitting the data. At the same time, we wished to identify a small subset of words that carried the most predictive value. The LASSO constraint (λ) was able to achieve both goals by constraining the sum of the absolute values of the estimated regression coefficients. This process helped to avoid overfitting and had the added benefit of excluding a large number of non-predictive words entirely from the model.

We considered models under a range of λ values. For each λ considered, we trained the model via 10-fold cross-validation, which randomly divided the data into 10 separate folds.² During each fold, the model trained on 9/10's of the available data. The model then used these training data to predict the remaining 1/10th of data. A goodness-of-fit was assessed for each

² We chose to divide the folds randomly during cross-validation to be consistent with Dobbins and Kantner (2019).

fold by comparing the difference in fit between the tested model and a completely saturated model that perfectly recovered the observed data. The goodness-of-fit for each fold was averaged together to establish the model's mean cross-validation error; the lower the mean cross-validation error, the better the model fit the data. For the last step of this process, we used the λ that achieved the minimum mean cross-validation error to fit a final model to all 10 folds in the data set.

Using these methods, we built two models: one for lineup identifications and another for lineup rejections. Lineup identifications can result in one of three decision outcomes: a correct identification (ID), a filler ID from a target-present lineup, and a filler ID from a target-absent lineup. A correct ID occurs when participants identify the target and a filler ID occurs when participants identify a filler (either from a target-present or target-absent lineup). Lineup rejections can result in one of two outcomes: a correct rejection or a miss. A correct rejection occurs when participants reject the lineup and the target is absent from the lineup. A miss occurs when participants reject the lineup but the target is present in the lineup. The first model, referred to as the Hits-vs-False-Alarms (HFA) model, aimed to distinguish between correct IDs and filler IDs regardless of whether the fillers were in target-present or target-absent lineups. We combined filler IDs in this manner to use as much data as possible. The second model, referred to as the Correct-Rejections-vs-Misses (CRM) model, aimed to distinguish between correct rejections and misses.

We analyzed 5,256 confidence statements from the 569 participants that took part in the Grabman et al. (2019) study. Participants provided 2,731 confidence statements after they identified a lineup member. There were 1,185 correct IDs and 1,546 filler IDs. We fit the HFA

model to those data. We fit the CRM model to the remaining 2,525 confidence statements that participants provided after they had rejected the lineup. There were 1,634 correct rejections and 891 misses. Figure 1 displays the mean cross-validation errors across a range of λ values for the HFA and CRM classifiers. The size of the points in the plot reflect the number of predictors (i.e., words) that remained in the classifier at that particular λ . This figure clearly shows that the mean cross-validation error decreases as λ increases. It is also clear that a stricter λ excludes more predictors from the model. Ideally, a minimum is reached such that further increases in λ come at the cost of increases in mean cross-validation error. The HFA model reached a minimum while retaining 104 words (the square border indicates this classifier). This is the final HFA model and we use this model in the remainder of this study. The CRM model reached a minimum mean cross-validation error only after having excluded all words from the model. This result indicates that the confidence statements participants made after they rejected the lineup did not contain enough useful predictive information as to whether that decision was a correct rejection or a miss. We do not consider the CRM model further for this reason. Interestingly, this pattern of results is precisely what Dobbins and Kantner (2019) found in their list-learning recognition study and further highlights the limited diagnostic information available for lineup rejections. Now that we have built a final HFA classifier, we move onto the results.

Figure 1. Mean cross-validation error during training for the CRM (yellow) and HFA (blue) classifiers across a range of λ values.

Note. The size of the points indicate the number of words retained in a model. The shading reflects bootstrapped 95% confidence intervals.

Results

The data for this experiment including the results from the language classifier are located at this repository: <https://osf.io/s64dj/files/>. We start this section by inspecting the words retained in the language classifier.

Words Diagnostic of Identification Accuracy

The final Hits-vs-False-Alarms (HFA) classifier retained 104 words from the corpus in order to classify correct IDs and filler IDs. The left and right panels of Figure 2 show the words – along with the value of their regression coefficients – that are predictive of a correct ID and a filler ID, respectively. Many of the words predictive of correct IDs come as no surprise. Words such as “completely”, “positive”, and “certain” convey a strong match-to-memory and the HFA classifier properly reflects the diagnostic value that these words contain. Similarly, the words that convey a weak match-to-memory such as “uncertain”, “somewhat”, “slightly”, and “maybe” are properly considered to be error-prone and predictive of filler IDs.

Figure 2. Regression coefficients for words retained in the final HFA Classifier.

Note. Left panel illustrates the words predictive of a correct ID. Right panel illustrates words predictive of a filler ID. In both panels, the words most predictive are displayed at the top and the words least predictive are displayed at the bottom.

It is surprising however, that the classifier did not retain words referencing internal facial features such as “eyes”, “nose”, and “mouth” even though these features are widely considered diagnostic in the face recognition literature (e.g., Davies, Ellis, & Shepherd, 1978). One possible reason why these features were not retained is because they were hardly mentioned: the word “eyes” was mentioned in 274 statements, the word “eye” was mentioned in 69 statements, the word “nose” was mentioned in 88 statements, and the word “mouth” was mentioned in 35 statements (out of a total of 2,731 statements). These results are consistent with a large-scale archival analysis of perpetrator descriptions provided by eyewitnesses to real crimes, which found that eyewitnesses mentioned perpetrators’ internal facial features less than 5% of the time (Van Koppen & Lochun, 1997). Similarly, Lindsay, Martin, and Webber (1994) found that the “eyes” were mentioned only 3% of the time and the “nose” and “mouth” were never mentioned in a sample of 105 eyewitness descriptions from real crimes. The only facial features the classifier retained were external facial features such as “chin”, “jaw”, “ears”, and “hair”. The classifier considered “hair” and “haircut” diagnostic of a correct ID and “chin”, “jaw”, and “ears” diagnostic of a filler ID.

Words in Context

Next, we inspect these words in context – within the raw confidence statements. The HFA classifier provides log odds for each confidence statement it evaluates. This value reflects the odds of a correct ID or a filler ID. High log odds mean the classifier is highly confident of a correct ID. Log odds equal to or near zero indicate the classifier is uncertain of a filler ID or a correct ID. Log odds far below zero mean the classifier is highly confident of a filler ID. We show the three confidence statements that received the highest log odds below:

1. “I am again **very certain because I remember his face and hair**. I **remember** thinking that **this** guy's **hair** looked fairly retro **and it is very** recognizable.” (3.46)
2. “I'm **very certain** that I saw **this** man **in** the previous lineup. I **remember his haircut and facial hair**.” (2.88)
3. “**Very certain** look **in his** eyes, way **he** holds **his** mouth **and haircut**.” (2.85)

The words the classifier used to evaluate the statements are in bold and are underlined. The bold words are predictive of correct IDs and the underlined words are predictive of filler IDs. The log odds are in parenthesis. These statements rightfully convey a strong recognition experience and the machine correctly classified each statement (i.e., a correct ID occurred). Next, we show the confidence statements that received the three lowest log odds below:

4. “Little certainty on **this** one, **but** more than not at all. **This face looks to be the** most familiar to me out **of all of** them based on facial structure and head shape.” (-5.74)
5. Not at all. These pictures all kind of look alike. (-3.97)
6. “About 40% certainty, **so not certain** at all, **but this is the** only one that seems familiar. It's all in the eyes. Otherwise, not much is familiar **but I know I saw** those eyes before.” (-3.95)

These statements rightfully convey a weak recognition experience and, again, the machine correctly classified each statement (i.e., a filler ID occurred).

An important point to make is that the classifier correctly categorized each statement even though they often mentioned an observable feature about the identified person. Others often perceive eyewitnesses as being less confident and less accurate in their identification from a fair lineup if they mentioned an observable feature (e.g., “I remember his eyes”) compared to when they mentioned an unobservable feature (e.g., “He looks like someone I

know”) or no feature whatsoever. This cognitive bias is known as the featural justification effect (Dodson & Dobolyi, 2015). Here, the classifier regards some observable features as diagnostic of correct IDs (e.g., “hair”) and others as diagnostic of filler IDs (e.g., “head” and “shape”). This suggests to us that the classifier may be less susceptible to the featural justification effect, though we need further research to investigate that possibility.

We also suspect that a person reading these statements might find it difficult to interpret them accurately. An investigator for example might misinterpret confidence statement #6 due to the phrase “but I know I saw those eyes before.” Focusing exclusively on that phrase might lead one to believe that the eyewitness has correctly identified the perpetrator, when, in fact, they have not. We do not know if the classifier is more capable of distinguishing correct from incorrect identifications than human evaluators. Based on these results, we suspect that could be the case. We are currently investigating that possibility. Together, these few examples highlight some of the classifier’s capabilities. In the next section, we investigate the classifier’s performance.

Classifier Performance

We assessed the performance of the retained HFA classifier in several ways and we discuss each analysis in turn.

Discrimination Accuracy

Mickes, Flowe, and Wixted (2012) introduced receiver operating characteristic (ROC) analysis to the field of eyewitness identification to assess how well eyewitnesses discriminate guilty from innocent suspects. Since then, researchers have used ROC analysis to assess

eyewitness discriminability in lineups (Dobolyi & Dodson, 2013; Mickes et al., 2012; Seale-Carlisle, Wetmore, Flowe, & Mickes, 2019) and in showups (Wilson, Seale-Carlisle, & Mickes, 2018). Here, we conduct ROC analysis, using the pROC package (Robin et al., 2011) in R v. 3.6.3 (R Core Team, 2020), to assess the classifier's ability to discriminate between correct IDs and filler IDs.

As mentioned previously, the HFA classifier provides a log odds score for each verbal confidence statement it evaluates. The left panel of Figure 3 shows the distribution of classifier log odds for all statements associated with correct and filler IDs, whereas the right panel of Figure 3 shows that distribution broken down by correct and filler IDs. As illustrated in the right panel, the log odds for correct IDs are, on average, greater ($M = 0.18$, $SD = 0.89$) than the log odds for filler IDs ($M = -0.70$, $SD = 0.85$). This result indicates the classifier has some discriminative capability.

Figure 3. Classifier log odds distribution collapsed across ID outcome (left panel) and broken down by correct and filler IDs (right panel)

The ROC demonstrates the classifier's discriminative capability across all response thresholds. Consider a classifier log odds equal to one. At this conservative threshold, 24 filler IDs and 196 correct IDs achieved this score or higher. In total, there were 1,185 correct IDs and 1,546 filler IDs, which means the hit rate (or sensitivity) is 16% ($196/1,185$) and the false alarm

rate (or $1 - \text{specificity}$) is 1.6% (24/1,546) at this threshold. Additional hit and false alarm rates are calculated for every possible threshold, thereby producing an ROC curve. The area underneath the curve (AUC) provides a nonparametric measure of the classifier's discrimination ability (Macmillan & Creelman, 2005). Figure 4A shows the ROC for the HFA classifier trained on the Grabman et al. (2019) data along with 95% confidence intervals calculated across all thresholds. The HFA classifier is well above the chance line, achieving an AUC of 77%, which is comparable to the AUC of 69% reported by Dobbins and Kantner (2019).

Recall that we divided the folds randomly during the 10-fold cross-validation to be consistent with Dobbins and Kantner (2019). This means that data from individual participants were likely split across each of the 10 folds. This cross-validation procedure may have therefore introduced a type of data leakage, where participant-related information from testing folds leaked into training folds. Data leakage is problematic because it may influence the generalizability of the classifier. For these reasons, we tested the HFA classifier on an entirely new sample of eyewitness confidence statements collected by Dobolyi and Dodson (2018) who used the same eyewitness paradigm as Grabman et al (2019). To generate predictions, words in the Dobolyi and Dodson confidence statements were frequency coded in the same manner as the words in the Grabman et al. (2019) confidence statements. However, the words used for predictions and their respective weights were limited to the words and regression coefficients shown in Figure 2. Figure 4B shows the ROC for the HFA classifier when tested on the Dobolyi and Dodson data along with 95% confidence intervals calculated at all thresholds. The HFA classifier achieved an AUC of 75%, which is slightly, but not significantly ($D = 1.13$, $p = 0.26$), less than the AUC it achieved on the Grabman et al. (2019) data. Together, these ROC results clearly

demonstrate the discriminative capabilities of the classifier and show that there is minimal (if any) data leakage.

Figure 4. ROC curves for the retained HFA classifier for the Grabman et al. (training) data (A) and for the Dobolyi and Dodson (out-of-sample) data (B).

Positive Predictive Value of Verbal Confidence

Knowing the classifier achieved an AUC of 75%-77% is important to establish as it demonstrates the classifier correctly discriminated between correct and filler IDs 75%-77% of the time. However, discriminability is only one type of “accuracy” that is important to the criminal justice system. The second type of accuracy that is important to the criminal justice system is positive predictive value. Positive predictive value is the probability that a positive identification is accurate. This measure can answer the following question: how accurate are eyewitnesses when the classifier is very certain of a correct identification? Over the years, researchers have introduced measures of positive predictive value for numeric confidence and response times (e.g., Juslin, Olsson, & Winman, 1996; Mickes, 2015; Seale-Carlisle, Colloff,

Flowe, Wells, Wixted, & Mickes, 2019). For example, Mickes (2015) introduced CAC analysis, which is concerned with the accuracy of suspect identifications made with a certain level of confidence (e.g., 100% confidence). Similarly, Seale-Carlisle et al. (2019) introduced response-time accuracy characteristic (RAC) analysis, which is concerned with the accuracy of suspect identifications made within a certain amount of time (e.g., within 5 seconds). Here, we introduce a similar analysis to CAC and RAC analyses by calculating the relationship between the classifier log odds and overall identification accuracy.

For this analysis, we binned the classifier log odds into one of several bins and then calculated identification accuracy for each bin. We calculated identification accuracy in each bin by counting the number of correct IDs and dividing that number by the sum of correct IDs, target-present filler IDs, and target-absent filler IDs (i.e., $\text{accuracy} = \# \text{ correct IDs} / [\# \text{ target-present filler IDs} + \# \text{ target-absent filler IDs} + \# \text{ correct IDs}]$). Here, the first bin consisted of identifications made with low log odds (i.e., equal to or less than -1). The second and third bin consisted of identifications made with slightly higher log odds (i.e., ranging from -0.99 to -0.1 for the second bin and 0 to +0.99 for the third bin). Lastly, the fourth bin consisted of identifications made with high log odds (i.e., equal to or greater than 1).³ Figure 5 shows the log odds-accuracy relationship for the Grabman et al. (2019) data (Figure 5A) and the Dobolyi and Dodson (2018) data (Figure 5B). In both cases, identifications with high log odds (i.e., equal to or greater than 1) achieved high identification accuracy: 89% for the Grabman et al. data and

³ We chose these cutoffs for several reasons. First, these cutoffs approximate classifier certainty of <.25, .25 - .50, .50-.75, > .75, while still being nice round numbers for the log odds values (technically the probabilities are, < .27, .27 - .50, .50 - .73, >.73). Moreover, they contained a reasonable number of identifications in each bin to calculate positive predictive value statistics. Lastly, they kept the distinction between positive and negative log odds intact.

83% for the Dobolyi and Dodson data. Identification accuracy decreased substantially as the log odds decreased. These results indicate a strong log-odds-accuracy relationship.

Figure 5. ID accuracy plotted for each log-odds bin. The size of the points reflect relative frequencies.

Suspect Identification Accuracy

The previous analysis included target-present and target-absent filler IDs in the calculation of identification accuracy, but Mickes (2015) made the point that suspect identification accuracy is most pertinent to the criminal justice system given that the suspect is the person under investigation. Neither Grabman et al. (2019) nor Dobolyi and Dodson (2018) included a designated innocent suspect in their target-absent lineups, meaning that we are unable to compute the number of innocent suspect identifications directly. However, given that their lineups were fair, it is possible to estimate the number of innocent suspects that theoretically would have been identified by dividing the number of target-absent filler IDs by

the size of the lineup (6). This is legitimate because, from the eyewitness's point of view, the innocent suspect is just another filler (the eyewitness has not seen either).

We therefore calculated suspect identification accuracy by counting the correct IDs and dividing that number by the sum of the correct IDs and the number of target-absent filler IDs divided by 6 (i.e., suspect ID accuracy = # correct IDs / [(# target-absent filler IDs / 6) + # correct IDs]). Figure 6 shows the relationship between the classifier log odds and suspect identification accuracy for the Grabman et al. (2019) data (Figure 6A) and the Dobolyi and Dodson (2018) data (Figure 6B).⁴ In both cases, confidence statements that received high log odds (i.e., equal to or greater than 1) achieved very high suspect identification accuracy: 99% for the Grabman et al. data and 98% for the Dobolyi and Dodson data. Suspect identification accuracy decreased substantially as the log odds decreased. These results indicate a strong relationship between the classifier log odds and the accuracy of suspect identifications. Given that the suspect is the person under investigation, the ability to predict the accuracy of a suspect identification adds tremendous value to the applied utility of the classifier.

Figure 6. Suspect ID accuracy plotted for each log odds bin. The size of the points reflect relative frequencies, which we calculated using the formula in Seale-Carlisle et al. (2019).

⁴ Because of the suspect-ID adjustment, it is not appropriate to use binomial confidence intervals for characterizing uncertainty. Instead, we bootstrapped 10,000 resampled datasets, and calculated the Highest Density Interval (HDI) around each proportion. The HDI is less susceptible to skew than percentile intervals and represents 95% of values where the probability density is greater than points outside these bounds.

Unique Positive Predictive Value of Verbal Confidence

Until this point, we have considered the performance of the HFA classifier in isolation. One major question is whether the classifier log odds contain unique positive predictive value after we take other indicators of identification accuracy, such as numeric confidence ratings, into account. We investigated this question using a logistic mixed effects model, but before describing this inferential analysis we explore the data descriptively first.

Descriptive Analysis

Descriptively, the information provided by the classifier log odds seem to differ somewhat from numeric confidence ratings. The left column of Figure 7 shows the distribution of classifier log odds for each numeric confidence rating. As shown in this figure, there exists a wide distribution of classifier log odds within each rating. For example, when participants identify someone with a rating of 100% confidence, the log odds of their verbal confidence statement can vary widely from approximately -2 to 3. Similarly, when participants identify

someone with a lower rating of 60% confidence, the log odds for their verbal confidence statement can range between -3 and 2.

Figure 7. Classifier log odds distributions collapsed across ID outcome for each numeric confidence rating (left column) and broken down by correct and filler IDs (right column).

The right column of Figure 7 shows these distributions broken down by correct and filler IDs. These distributions seem to interact with numeric confidence in an important way: as numeric confidence increases, the correct ID distribution shifts to the right and gains separation from the filler ID distribution. For example, when participants identify someone with 20% confidence, the correct and filler ID distributions overlap extensively. However, when participants identify someone with 100% confidence, the correct ID distribution shifts further to the right of the filler ID distribution. This suggests that verbal confidence adds positive predictive value to numeric confidence especially as numeric confidence increases. Ignoring verbal confidence and only measuring numeric confidence therefore seems to lose diagnostic information.

Inferential Analysis

To test whether verbal confidence adds significant positive predictive value to numeric confidence ratings and other indicators of identifications accuracy, we adopted a multi-model comparison approach (Burnham & Anderson, 2002). We constructed a series of logistic mixed effects models that predict the accuracy of an identification from the chooser data collected by Grabman et al. (2019), which included correct IDs from target-present lineups and filler IDs from target-present and target-absent lineups. Identification accuracy is therefore modeled as $\text{accuracy} = \# \text{ correct IDs} / [\# \text{ target-present filler IDs} + \# \text{ target-absent filler IDs} + \# \text{ correct IDs}]$. The predictors in the models included the retention interval condition (i.e., delay of either 5 minutes or 1 day) and the lineup race condition (i.e., same race or cross race) as well as participants' numeric confidence ratings, response times, and CFMT scores. Importantly, the classifier log odds were also included as a predictor. Because each participant provided multiple

responses, we added the participant ID as a random intercept. All continuous predictors (i.e., numeric confidence ratings, CFMT scores, and response times) were centered and scaled across all participants prior to model fitting.

To start, we fit separate starting models consisting of all 4-way, 3-way, 2-way, and main effects using the *lme4* package (Bates, Machler, Bolker, & Walker, 2015) in *R* v. 3.6.3 (R Core Team, 2020). In order to obtain the most parsimonious model fit, we performed a backward stepwise procedure based on Akaike's Information Criterion (AIC) that selected the best fitting model from each start-point. This method removes model terms that demonstrate any improvement in AIC, while maintaining the principles of marginality (e.g., not removing a three-way term if it is nested in a higher four-way term). This process yielded four reduced models (one for each start point) that we compared based on AIC (Table 1).

Table 1. Fit statistics of models generated by the backward stepwise procedure.

Model	Terms	df	AIC	BIC	-2LL
Main Effects	Classifier+Confidence+ResponseTime+CFMT+Delay+LineupRace+ (1 Participant)	8	2814.3	2861.6	2798.3
Two-way	Classifier + Confidence + ResponseTime + CFMT + Delay + LineupRace + Classifier:Confidence + Classifier:ResponseTime + Confidence:ResponseTime + Confidence:CFMT + Confidence:Delay + Confidence:LineupRace + (1 Participant)	14	2774.6	2857.4	2746.6
Three-way	Classifier + Confidence + ResponseTime + CFMT + Delay + LineupRace + Classifier:Confidence + Classifier:ResponseTime + Classifier:CFMT + Classifier:Delay + Confidence:ResponseTime + Confidence:CFMT + Confidence:Delay + Confidence:LineupRace + ResponseTime:Delay + CFMT:Delay + Classifier:Confidence:CFMT + Classifier:CFMT:Delay + Confidence:ResponseTime:Delay + (1 Participant)	22	2773.7	2903.8	2729.7
Four-way	Classifier + Confidence + ResponseTime + CFMT + Delay + LineupRace + Classifier:Confidence + Classifier:ResponseTime + Classifier:CFMT + Classifier:Delay + Classifier:LineupRace + Confidence:ResponseTime + Confidence:CFMT + Confidence:Delay + Confidence:LineupRace + ResponseTime:CFMT + ResponseTime:Delay + CFMT:Delay + CFMT:LineupRace + Delay:LineupRace + Classifier:Confidence:ResponseTime + Classifier:Confidence:CFMT + Classifier:ResponseTime:CFMT + Classifier:ResponseTime:Delay + Classifier:CFMT:Delay + Classifier:CFMT:LineupRace + Classifier:Delay:LineupRace + Confidence:ResponseTime:CFMT + Confidence:CFMT:Delay + Confidence:CFMT:LineupRace + Confidence:Delay:LineupRace + CFMT:Delay:LineupRace + Classifier:Confidence:ResponseTime:CFMT + Classifier:CFMT:Delay:LineupRace + Confidence:CFMT:Delay:LineupRace + (1 Participant)	37	2785.5	3004.3	2711.5

A model that incorporated three-way terms provided the best fit to the data (AIC = 2773.7). However, given that a less complex two-way model contained far fewer parameters – with only a negligible loss in fit (AIC = 2774.6) – we decided to adopt it as the final model. In Wilkinson-Rogers (1973) notation, this final model consisted of several main effects and two-way interactions: $\text{Accuracy} \sim \text{Classifier} + \text{Confidence} + \text{ResponseTime} + \text{CFMT} + \text{Delay} + \text{LineupRace} + \text{Classifier:Confidence} + \text{Classifier:ResponseTime} + \text{Confidence:ResponseTime} + \text{Confidence:CFMT} + \text{Confidence:Delay} + \text{Confidence:LineupRace} + (1 | \text{Participant})$.

To our knowledge, there are no consensus standards about evaluating the absolute fit of generalized mixed effects models. For this reason, we examined the fit of the final model in two ways. First, we used the *MuMIn* package to calculate both the marginal ($R^2\text{GLMM}(m)$) and conditional ($R^2\text{GLMM}(c)$) pseudo- R^2 . Whereas the marginal statistic includes only the variance accounted for by the fixed effects, the conditional statistic incorporates variance from the random effects. Both statistics suggested a good fit to the data ($R^2\text{GLMM}(m) = .386$, $R^2\text{GLMM}(c) = .493$). Additionally, we used the *DHARMa* package (Hartig, 2020, version 0.2.7) to perform a Kolmogorov-Smirnov test for uniformity, comparing the observed data to a cumulative distribution of 1,000 simulated datasets from the model estimates. This test also suggested an adequate model fit (KS $D = .023$, $p = .114$).

Satisfied with the fit of the model, we used the *afex* package (Singmann, Bolker, Westfall, & Aust, 2016) to perform likelihood ratio tests (LRTs) to compute the significance of model terms. This yielded multiple significant main effects and interactions (Table 2). Many of these effects are similar to those reported in Grabman et al. (2019). For example, we and Grabman et al. both observe interactions between numeric confidence ratings and response

times. We do not expand more on those effects here and instead focus on the novel results, which are those that involve the HFA classifier.

Table 2. Likelihood ratio tests (LRTs) on the final logistic mixed-effects model.

Effect	Chi df	χ^2	<i>p</i> -value
Classifier	1	125.57	<.001
Confidence	1	91.75	<.001
ResponseTime	1	21.50	<.001
CFMT	1	40.64	<.001
Delay	1	10.95	.001
LineupRace	1	6.72	.010
Classifier:Confidence	1	8.09	.004
Classifier:ResponseTime	1	2.97	.085
Confidence:ResponseTime	1	6.75	.009
Confidence:CFMT	1	3.15	.076
Confidence:Delay	1	2.95	.086
Confidence:LineupRace	1	5.25	.022

Consistent with the descriptive analysis discussed in the previous section, the model revealed a significant interaction between the classifier log odds and numeric confidence ratings, $\chi^2(1) = 8.09$, $p = .004$. Figure 8 shows this interaction. We used the *effects* package to calculate model estimates (lines) and 95% confidence intervals (error shading). There are several interesting aspects of this figure. First, the model predicts a strong relationship between the classifier log odds and identification accuracy. This result replicates the results shown in Figure 5, which did not rely on any model fitting and simply calculated identification accuracy across several log odds bins. Second, identification accuracy increases as numeric confidence increases, which is in line with decades of empirical research (Wixted & Wells, 2017). However, the interaction suggests that the diagnostic value of numeric confidence depends on the verbal classifier log odds (and vice versa). When the classifier is highly confident of a correct ID (e.g., a

verbal log odds score of 2), the model predicts that greater numeric confidence (e.g., the 100% 'red' line) is associated with much higher identification accuracy, as compared to lower levels of numeric confidence (e.g., the 0% 'purple' line). Yet, when the classifier is highly confident about a filler ID (e.g., a verbal log odds score of -2), the model predicts that identification accuracy is nearly indistinguishable between the numeric confidence ratings. Together, these results suggest that both numeric and verbal confidence provide unique positive predictive value.

Figure 8. Interaction between the Classifier log odds values and Confidence ratings when predicting identification accuracy.

Note. Lines indicate model estimates, whereas the error shading represents 95% confidence intervals. The graph clearly shows that the probative value of a confidence rating differs based on the words a witness uses to justify their identification from a lineup.

Next, we investigate whether verbal confidence statements add positive predictive value to the other predictors in the model. Figure 9 shows the classifier log odds (x-axis) plotted against the significant main effects of (A) delay, (B) lineup race, (C) face recognition ability (CFMT), and (D) response time. Figure 9A shows the model predicted a similar log odds-accuracy relationship across the delay manipulation. According to the model, the classifier log odds are similarly predictive of identification accuracy regardless of whether participants saw the target 5 minutes or 1 day ago. Figure 9B shows the model predicted a similar log odds-accuracy relationship for same race and cross race identifications. Figure 9C shows the model predicted a stronger log odds-accuracy relationship for strong face recognizers (e.g., CFMT score of 70) than for weak face recognizers (e.g., CFMT score of 30). Lastly, Figure 9D shows the

model predicted a weaker log odds-accuracy relationship for slow identifications (e.g., after 30 seconds) compared to fast identifications (e.g., within 5 seconds).

Figure 9. Identification accuracy as a function of Verbal classifier log odds and Delay (A), Lineup Race (B), Face recognition ability (i.e., CFMT score; C) and Response Time (D).

Note. Classifier log odds plotted against the main effects of (a) Delay, (b) Lineup Race, (c) Face recognition ability (CFMT), and (d) Response Time. Lines indicate model estimates, whereas the error shading represents 95% confidence intervals. Delay and Lineup race show limited impact on the utility of the classifier. However, misidentifications become more frequent when participants are poor recognizers or respond slowly.

As a whole, these results highlight the unique positive predictive value contained within verbal confidence statements. Knowing a participant's numeric confidence rating, response time, or CFMT score can help shed light on whether their identification is accurate or inaccurate, but collecting their verbal confidence statement and evaluating that statement using the HFA classifier adds additional predictive information.

Relationship between face recognition ability and language

In the previous sections that focused on positive predictive value, we showed that the classifier log odds are informative of identification accuracy: high log odds can indicate high accuracy and low log odds can indicate low accuracy (see Figure 5 and 6). However, the strength of that relationship depends on certain factors. For example, a high numeric confidence rating strengthens the log odds-accuracy relationship, whereas a low numeric confidence rating weakens it (see Figure 8). Here, we turn our focus from the log-odds accuracy relationship to address *which eyewitnesses* are more or less likely to provide diagnostic confidence statements.

When eyewitnesses correctly identify the target from a lineup, they ideally use language that is diagnostic of a correct ID. Conversely, when eyewitnesses incorrectly identify a filler, they ideally use language that is diagnostic of a filler ID. Stronger face recognizers achieved a stronger log odds-accuracy relationship than weaker face recognizers (see Figure 9C), but do stronger face recognizers provide more diagnostic confidence statements, on average, compared to weaker face recognizers? We investigate this question in a similar manner to the previous model fitting exercise, but instead of predicting the accuracy of an identification, we predict the classifier log odds.

We built a series of linear mixed effects models to predict the classifier log odds based on how well participants score on the CFMT. The predictors in the models included participants' CFMT scores, trial outcome (i.e., correct ID or filler ID), and participants' overall discrimination accuracy, which we measured by calculating d' (i.e., d prime) for every participant (Macmillan & Creelman, 2005; Mickes, Moreland, Clark, & Wixted, 2014). It is crucial to account for

participants' discrimination accuracy in these models. Failing to do so may test what we already know: stronger face recognizers experience stronger memory-match signals, on average, than weaker face recognizers. Including participants' d' as a predictor therefore allows us to test for differences in language across face recognition ability while holding discrimination accuracy constant.⁵ Lastly, we added the participant ID as a random effect because each participant made multiple responses. We centered and scaled the CFMT and d' scores across all participants prior to model fitting.

As in the previous model fitting exercise, we started by fitting models consisting of all 3-way, 2-way, and main effects using the *lme4* package (Bates et al., 2015) in R v. 3.6.3 (R Core Team, 2020). The most parsimonious fit, as determined by AIC, was obtained by performing a backward stepwise procedure from each of these starting points. This process resulted in three reduced models. The two models that initially incorporated all 3-way and 2-way interactions reduced to the same, best-fitting model (AIC = 6812.3). This process removed the 3-way interaction (i.e., CFMT:TrialOutcome:dprime) as well as two 2-way interactions (i.e., CFMT:dprime and TrialOutcome:dprime). The final model in Wilkinson-Rogers (1973) notation is: Classifier ~ TrialOutcome + dprime + CFMT + CFMT:TrialOutcome + (1 | Participant). The marginal and conditional pseudo- R^2 statistics both suggested an adequate fit to the data ($R^2_{GLMM(m)} = .215$, $R^2_{GLMM(c)} = .337$). Likelihood ratio tests revealed that all terms in the final model were significant (Table 4).

Table 3. Fit statistics of models generated by the backward stepwise procedure.

⁵ We corrected for correct and false ID rates equal to 0 or 1 by following the recommendation discussed by Macmillan and Kaplan (1985). That is, we replaced 0 by $1 / (2 * N)$ and 1 by $1 - (1 / (2 * N))$, where N equals the number of trials.

Model	Terms	df	AIC	BIC	-2LL
Main Effects	TrialOutcome + dprime + (1 Participant)	5	6825.3	6854.9	3407.7
Two-way	TrialOutcome + dprime + CFMT + CFMT:TrialOutcome + (1 Participant)	7	6812.3	6853.6	3399.1
Three-way	TrialOutcome + dprime + CFMT + CFMT:TrialOutcome + (1 Participant)	7	6812.3	6853.6	3399.1

Table 4. Likelihood ratio tests (LRTs) on the final model.

Effect	Chi df	χ^2	<i>p-value</i>
CFMT	1	4.35	.037
Trial Outcome	1	417.52	< .001
d prime	1	6.88	.009
CFMT : Trial Outcome	1	15.28	< .001

The question of interest is whether face recognition ability (i.e., CFMT score) predicts the diagnostic value of participants' verbal confidence statements. The significant main effect of CFMT suggests that it does: stronger face recognizers generally provide verbal confidence statements that yield higher classifier log odds, $\chi^2(1) = 4.35$, $p = .037$. Crucially, Figure 10 shows that the effect is qualified by a significant interaction with Trial Outcome (color lines), $\chi^2(1) = 15.28$, $p < .001$. Again, we used the *effects* package to calculate model estimates (lines) and 95% confidence intervals (error shading). This figure shows that as face recognition ability (i.e., CFMT score) increases so too does the classifier log odds when the identification is accurate (i.e., for correct IDs). When the identification is inaccurate (i.e., for filler IDs), the model predicts a slight decrease in classifier log odds with increases in face recognition ability (i.e., CFMT score). This results in a much larger difference in the verbal classifier log odds between correct

IDs and filler IDs for stronger, as compared to weaker, face recognizers. Importantly, this effect of face recognition ability on classifier log odds for correct and filler IDs persists across different levels of discrimination accuracy (as measured by d'). The panels in Figure 10 (from left to right) show the same pattern of results as d' varies from poor ($d' = .5$), moderate ($d' = 1$), to good ($d' = 1.5$). These results suggest that stronger face recognizers compared to weaker face recognizers have superior metacognitive awareness, which they use to select the words that best describe the strength of their memories.⁶ That is, stronger face recognizers express higher confidence when making a correct ID but express lower confidence when making a filler ID, which is a key sign that they are more cognizant of the strength of their memory.

Figure 10. Increasing face recognition ability (i.e., CFMT score) is associated with increasing diagnostic value of language (i.e., verbal classifier log odds) for correct IDs but not for filler IDs.

⁶ Note that metacognitive ability depends in part on discrimination accuracy. When discrimination accuracy is at chance levels (i.e., $d' = 0$), poor metacognitive performers and high metacognitive performers are indistinguishable. The key takeaway from this analysis is that face recognition ability appears to play a much stronger role than discriminability.

Note. Lines indicate model estimates, whereas the error shading represents 95% confidence intervals. This interaction effect remains even when we hold discrimination accuracy (d') constant across face recognition ability.

Discussion

What is the relative value of verbal and numeric expressions of confidence when predicting eyewitness identification accuracy? To answer this question, we adapted the machine-learning methodology used by Dobbins and Kantner (2019) to investigate the natural language of accurate and inaccurate eyewitnesses. This method followed a “bag of words” approach, which treated each individual word in an eyewitness confidence statement as a potentially predictive feature of identification accuracy. In contrast to previous investigations (e.g., Behrman & Richards, 2005), this method provided a finer-grained analysis of eyewitness verbal confidence statements. We then took a series of steps to evaluate the overall performance of the final Hits-vs-False-Alarms (HFA) classifier.

In the first step, we inspected the words retained in the HFA classifier. Reassuringly, the words that convey high confidence (e.g., “completely”) were determined to be strong indicators of correct identifications, whereas the words that convey low confidence (e.g., “semi”) were determined to be strong indicators of filler identifications. We then inspected these words in context (i.e., within the raw confidence statements). The classifier provided log odds for each statement it evaluated. That log odds represents the classifier’s certainty as to whether a decision was a correct identification or a filler identification. Qualitatively, the classifier log odds seemed to hold diagnostic value based on the statements we inspected.

These results suggest the classifier was able to distinguish between strong and weak recognition experiences.

The second step consisted of several analyses that evaluated the performance of the classifier. Receiver operating characteristic (ROC) analysis determined that the classifier correctly distinguished between correct identifications and filler identifications 77% of the time, when trained and tested on the Grabman et al. (2019) data (Figure 4A). The classifier achieved a similarly impressive score of 75% when tested on an out-of-sample data set collected by Dobolyi and Dodson (2018; see Figure 4B). The ROC results are important for showing that the classifier is far above chance discrimination. We also found that the classifier had a tremendous capability to predict the accuracy of a suspect identification. When the classifier was highly certain (i.e., when the log-odds was equal to or greater than 1), there was a 99% and a 98% likelihood that the suspect identification was accurate for the Grabman et al (2019) data and the Dobolyi and Dodson (2018) data, respectively (see Figure 6). As the classifier became less certain of a correct identification, suspect identification accuracy dramatically decreased.

Once these two steps were completed, it became clear that the classifier was capturing the rich, diagnostic information contained within the eyewitness verbal confidence statements. The third step we took tested whether that diagnostic information was unique to the verbal confidence statements or if that information was shared amongst other indicators of identification accuracy, such as numeric confidence ratings or response times. For this step, we used a logistic mixed effects model to predict identification accuracy based on various factors including the classifier log odds, numeric confidence ratings, response times, CFMT score, retention interval, and lineup race. The best-fitting model found that the main effect of the

classifier log odds was qualified by a significant interaction with numeric confidence ratings (see Figure 7). This figure shows that eyewitness verbal confidence statements contain diagnostic information that numeric confidence ratings do not capture. Thus, knowing participants' numeric confidence rating can help shed light on whether their identification is accurate or inaccurate, but we can gather additional diagnostic information by asking participants to express and justify their level of confidence in their own words and then having the HFA classifier evaluate their verbal confidence statement.

The fourth and final step we took investigated the relationship between face recognition ability and the ability to provide a diagnostic verbal confidence statement. We found that strong face recognizers were better at justifying their identification than weak face recognizers (see Figure 10). That is, strong face recognizers as compared to weak face recognizers expressed higher confidence in their correct identifications and expressed slightly lower confidence in their filler identifications. Crucially, this pattern held even when strong and weak face recognizers had the same overall discrimination accuracy. Ultimately, this result suggests that stronger face recognizers have superior meta-cognitive awareness, which they use to better describe the strength of their memories -- in line with recent work examining usage of the numeric confidence scale across face recognition ability (Gettleman et al., 2020).

Theoretical Implications

Theoretically, what explains the interaction between verbal confidence statements and numeric confidence ratings observed in this study? One possible explanation for this result is that an eyewitness's verbal confidence statement contains finer-grain information about the accuracy of an identification than a numeric confidence rating from a 6-point rating scale. To

illustrate this point, we make use of signal-detection theory, which has long guided thinking about recognition memory ever since Egan (1958) first introduced it several decades ago.

Consider the signal-detection model for a fair, 6-person, simultaneous lineup shown in Figure 11A. This figure illustrates two Gaussian distributions of memory signals resting along a memory strength axis. The distribution on the left represents the memory signals generated by the fillers and the distribution on the right represents the memory signals generated by the targets. A 6-person target-present lineup consists of five random draws from the filler distribution and one random draw from the target distribution, whereas a 6-person target-absent lineup consists of six random draws from the filler distribution. Because the two distributions overlap, eyewitnesses must draw a decision criterion at some point along the memory strength axis in order to make an identification decision. In this illustration, we placed the decision criterion at the midpoint between the two distributions, but it is possible for eyewitnesses to place this criterion at more conservative locations (right of the midpoint) or liberal locations (left of the midpoint) along the memory strength axis. For simplicity, this model assumes the independent-observations decision rule (Macmillan & Creelman, 2005). That is, the eyewitness identifies the lineup member that generates the strongest memory signal only if that signal surpasses their decision criterion. Otherwise, the eyewitness identifies no one and selects the “not present” option instead.

A straightforward extension of this basic signal-detection model holds that confidence judgments are reached in much the same way as the decision to “identify” or “not identify” is reached. That is, a different decision criterion for each level of confidence is placed somewhere along the memory strength axis. Participants in the Grabman et al. (2019) study used a 6-point

rating scale. Figure 11B illustrates the possible placements of these 6 decision criteria, which, from left to right, reflect confidence ratings of 0%, 20%, 40%, 60%, 80%, and 100% confidence, respectively. According to this model, eyewitnesses make an identification with 100% confidence when the strongest memory signal generated from the lineup surpasses the 100% criterion (located on the far right). If none of the lineup members generate a strong enough signal to surpass the 0% criterion (i.e. the left-most criterion), the eyewitness identifies no one and chooses the “not present” option. The key insight here is that memory signals that differ in strength can receive the same confidence rating. For example, a memory signal that barely surpasses the 80% criterion would be given the same confidence rating as a memory signal that far exceeds that criterion but fails to surpass the 100% criterion. Although these memory signals differ in strength, both identifications are made with 80% confidence.

Figure 11. Signal-detection models illustrating one (A), six (B), and twenty (C) response criteria.

One way to combat this issue is to increase the grain-size of the rating scale. For example, consider a 20-point rating scale ranging from 0% – not certain at all – to 100% – completely certain with increments of 5%. Figure 11C illustrates the possible placements of these 20 decision criteria along the memory strength axis. With this finer-grain scale, a memory signal strong enough to surpass the 80% criterion, but not strong enough to surpass the 100% criterion, can now be assigned a confidence rating of 80%, 85%, 90%, or 95% (depending on where that memory signal lands). In this sense, the 20-point rating scale provides finer-grain information about the accuracy of an identification than the 6-point rating scale. A similar argument can be made in regards to eyewitness verbal confidence statements. Allowing the eyewitness to express their confidence in their own words and then evaluating that statement using the HFA classifier may help pinpoint the exact strength of that memory signal. In this example, the verbal confidence statement may help determine whether that memory signal landed near the 80% criterion or closer to the 100% criterion.

If we take this argument to the extreme, then why not use a very fine-grain rating scale? For example, a 100-point rating scale might provide even more fine-grained responses than the previous options discussed. Mickes, Wixted, and Wais (2007) conducted a standard recognition memory test during which targets and lures were presented one at a time, and participants were asked to rate each item using a 99-point rating scale ranging from 1 – the item was definitely *not* on the list – to 99 – the item was definitely *on* the list (yielding 100 confidence bins). They found that participants often provided ratings on the fives (i.e., 5, 10, 15, 20, etc.), which essentially converted the 99-point scale back into a 20-point scale. In a similar experiment, Mickes, Hwe, Wais, and Wixted (2011) conducted a recognition memory test, but

allowed participants to choose their own rating scale. Although participants were free to use whatever numbers they wished to scale the strength of their memories, most participants spontaneously chose a 20-point scale and, on average, only 17 separate ratings were used. These results suggest that ratings scales with more than 20 points yield no apparent gain than a 20-point scale. If the rating scale is more fine-grained than that, participants might find it difficult to differentiate the strengths of their memories (Mickes et al., 2011). Such a scale may also be cognitively burdensome for participants to use because they would need to maintain a large number of decision criteria along the memory strength axis (Benjamin, Tullis, & Lee, 2013). Thus, increasing the grain-size of the rating scale to an extreme degree may not achieve the same level of fine-grained responses that language may naturally provide.

Regardless of which rating scale is used, it is possible that numeric confidence ratings are simply poorer proxies of memory strength compared to verbal confidence statements. There are two reasons why this might be the case.⁷ First, participants must learn to translate their recognition experience to a number on a scale, which likely depends on how numerate they are (Winman, Juslin, Lindskog, Nilsson & Kerimi, 2014). However, language is a highly practiced skill, and it is presumably easier for participants to express certainty in their decisions using their own words. After all, humans have used language in this way for thousands of years, but have only started using numbers and probabilities to express confidence within the last few hundred years (Zimmer, 1983). In line with this view, surveys have found that people prefer to express their certainty using language because they can communicate nuance and imprecision

⁷ We are grateful to an anonymous reviewer for making these points.

better with words than with numbers (Brun & Teigen, 1988; Wallsten, Budescu, Zwick, & Kemp, 1993). Second, recent studies suggest that people show 'trait-like' use of the numeric confidence scale across a variety of tasks (Kantner & Dobbins, 2019; Kantner & Lindsay, 2014; Stankov, Kleitman, & Jackson, 2015). That is, participants' numeric confidence ratings can carryover from task to task even when the tasks are quite different. For example, Kantner, Solinger, Grybinas, and Dobbins (2019) found that participants' numeric confidence ratings for their perceptual discrimination judgments were indicative of their numeric confidence ratings for their recognition memory judgments. A 'confidence trait' like this could mean that numeric confidence ratings are poorer proxies of memory strength compared to verbal confidence statements.

One can also make the argument that verbal confidence statements are poorer proxies of memory strength compared to numeric confidence ratings. There are several reasons why this might be the case. First, there are features of faces that are not easy to describe (e.g., the distance between features), which suggests that some of the reasons why an eyewitness made an identification may not be fully captured in words (e.g., Schooler & Engstler-Schooler, 1990). Second, individuals likely differ in the vocabulary they use to justify their identifications and the classifier may have eliminated some of the words that are idiosyncratic, yet diagnostic. Lastly, in some cases, people may not know the exact reasons why they made an identification, and so they undergo a process of inference to come up with a sufficient explanation (e.g., Nisbett & Wilson, 1977). In these scenarios, a numeric confidence rating could succeed in capturing diagnostic information where a verbal confidence statement fails.

Distinguishing Between Recollection and Familiarity

Central to all dual-process theories of recognition memory is the claim that recognition decisions are based on two qualitatively distinct retrieval processes: recollection and familiarity (Yonelinas, 1994; Wixted, 2007). The process of recollection consists of retrieving contextual details associated with a studied item, whereas the process of familiarity allows one to appreciate the fact that an item has been previously encountered even though no contextual detail can be retrieved (Mandler, 1980). Critically, researchers often rely on participants to report when they are experiencing recollection and when they are experiencing familiarity (e.g. Gardiner, Ramponi, & Richardson-Klavehn, 1998; Mickes, Seale-Carlisle, & Wixted, 2013; Tulving, 1985). Rather than rely on participants' self-reports, Dobbins and Kantner (2019) tested whether their language classifier could parse recollection-based and familiarity-based memories based on participants' verbal expressions of confidence. According to their language classifier, the words that referenced context, such as the time and place when participants originally encountered the items (e.g., "when", "me", and "remember"), were strong indicators of hits, whereas the words that referenced a failure to retrieve context (e.g., "not", "but", and "think") were strong indicators of false alarms. Based in part on these results, Dobbins and Kantner concluded, "the HFA classifier appears to recover a distinction between a recollective versus non-recollective basis (R/NR) for the two outcomes" (p. 4).

The HFA classifier applied to the Grabman et al. (2019) data appears to distinguish between recollection-based and familiarity-based memories as well. Recollection-based words were usually diagnostic of correct IDs, whereas familiarity-based words were usually diagnostic of filler IDs. For example, words that reference when the target was encountered (e.g.,

“remember”, “yesterday”, “earlier”, and “previously”) were diagnostic of correct IDs. However, words that reference a failure to retrieve context (e.g., “not”, “maybe”, and “don’t”) or an outright feeling of familiarity (e.g., “feel” and “familiar”) were diagnostic of filler IDs.

Counterintuitively, a number of words that reference specific details about the identified person (e.g., “ears”, “chin”, “jaw”, “person”, “facial”, “head”, “features”, and “shape”) were diagnostic of filler IDs. This is a counterintuitive result considering these participants may have been experiencing recollection when they made their identification. Many researchers emphasize that recollection is a continuous memory process (e.g., Johnson, Hashtroudi, & Lindsay, 1993; Slotnick & Dodson, 2005; Wixted, 2007). That is, people can have strong, vivid recollections as well as weak, vague recollections. Proponents of this view often conceptualize the well-known tip-of-the-tongue state as reflecting trace amounts of recollection, for example. Perhaps when participants justified their identification by referring to these particular features, they were experiencing only trace amounts of recollection and the classifier therefore determined these words to be error prone.

Some may reject this dual-process account in favor of either a single-process account (e.g. Donaldson, 1996; Dunn, 2004) or an account that can aggregate the contributions of recollection and familiarity onto a single, continuous dimension (Wixted, 2007; Wixted & Mickes, 2010). From this perspective, the HFA classifier might be distinguishing between strong and weak memories rather than two underlying memorial processes such as recollection and familiarity. One challenge to this view is that the CRM classifier failed to distinguish between correct rejections and misses. This should not be the case if participants base their decision to “identify” or “not identify” on the strength of one underlying, continuous memory signal.

However, Yilmaz and Wixted (under review) find evidence that participants use a combination of memory signals generated from multiple faces in the lineup to reject the lineup. If that is true, then participants' confidence in lineup rejections reflects a much noisier signal than the signal participants use as a basis for their confidence in lineup identifications. That possible difference between lineup rejections and identifications may partly explain why the CRM classifier failed and why the HFA classifier succeeded. Future work that incorporates verbal confidence measures is needed to address these issues.

Applied Implications

The results reported in this study are theoretically interesting, but from an applied perspective, this machine-learning methodology may prove useful to those who work in the criminal justice system and must evaluate eyewitnesses' verbal confidence statements. Although we and others (e.g., Behrman & Richards, 2005; Mansour, 2020) have found that verbal confidence statements are informative of identification accuracy, people often have trouble accurately interpreting them (e.g., Dodson & Dobolyi, 2015; Dobolyi & Dodson, 2018; Grabman & Dodson, 2019; Grabman, Cash, Slane, & Dodson, *in press*). This machine-learning methodology may minimize this problem by providing people a more accurate interpretation of eyewitness confidence statements. More broadly, machine-learning technology may help people interpret verbal confidence statements across a wide range of decisions outside the field of eyewitness identifications (e.g., climate science; Budescu et al., 2014).

From a policy standpoint, these results suggest that law enforcement should collect both a numeric confidence rating as well as a verbal confidence statement, which includes the reason why eyewitnesses made their identification. Although we caution against using the

particular machine learning classifier generated by this study for real-world cases, our results suggest that law enforcement can gather more diagnostic information about the likely accuracy of eyewitness identifications by collecting both types of confidence expressions than by collecting one and not the other. As far as we know, the only country with this guideline is Sweden (Fitzgerald et al., 2021).

Other Machine-learning Classifiers

The HFA classifier made predictions based on the individual words contained within participants' verbal confidence statements (i.e., unigrams). How does this unigram classifier compare to other classifiers that consider features that are more complicated? Following Dobbins and Kantner (2019), we built two additional HFA classifiers: one that considered *only* bigrams, which are features comprised of two words, and another classifier that considered *only* trigrams, which are features comprised of three words. We built the bigram and trigram classifiers using the same method we used to build the unigram classifier. In Appendix A and B, we show the features retained in the best-fitting bigram and trigram classifiers, respectively. The bigrams most diagnostic of correct and filler IDs were "his hair" and "at all", respectively. Whereas, the trigrams most diagnostic of correct and filler IDs were "because he looks" and "certain at all", respectively.

One benefit of bigrams and trigrams compared to unigrams is that they are relatively easy to interpret because they do a better job of capturing context (see Appendix A and B). However, the cost of using these features is that they can be increasingly specific to the language within the training data set. They can therefore generalize poorly to out-of-sample data. To test the performance of these classifiers, we conducted ROC analysis. We show these

ROCs in Figure 12. The top row contains the ROCs applied to the Grabman et al. data, whereas the bottom row contains the ROCs applied to the Dobolyi and Dodson data, which are the out-of-sample data. Each classifier discriminated correct IDs from filler IDs far above chance performance, but the unigram classifier, which is the focus of this paper, yielded the highest out-of-sample AUC (Figure 12D), whereas the trigram classifier yielded the lowest AUC (Figure 12F). This pattern of results replicates Dobbins and Kantner (2019) who found that the AUCs decreased as the feature lengths increased.

Figure 12. ROCs for the unigram, bigram, and trigram classifiers for the Grabman et al. (2019) training data (top row) and the Dobolyi and Dodson (2018) out-of-sample data (bottom row).

An alternative way of exploring the added predictive value of bigrams and trigrams is to build a bigram classifier that incorporates unigrams *and* bigrams, or a trigram classifier that

incorporates unigrams, bigrams, *and* trigrams (e.g., Tan, Wang, & Lee, 2002). Using this alternative approach, ROC analysis revealed that both the uni/bi-gram and uni/bi/tri-gram classifiers achieved an AUC of .76 (when tested on the Grabman et al., and Dobolyi and Dodson data). This AUC score is slightly less than the AUC for the unigram classifier when tested on the Grabman et al. data (.77), but is slightly higher than the AUC for the unigram model when tested on the Dobolyi and Dodson data (.75). Ultimately, we found that adding in these more complicated features did not yield a worthwhile improvement in discrimination accuracy, nor did they change the foregoing conclusions drawn from unigram model. While this is true for these data sets, we encourage researchers interested in using machine-learning approaches to test alternative classifiers.

Limitations

A few limitations to our study are worth emphasizing. First, future studies should collect participants' verbal confidence statements using a broader range of stimuli to minimize item-based effects. For example, the HFA classifier retained words such as "hair", "haircut", and "afro" which means these features were diagnostic within the stimulus set used by Grabman et al. (2019) and Dobolyi and Dodson (2018). While these features were diagnostic of correct IDs within these stimuli, we suspect these words would be less predictive (or not predictive at all) if hairstyles were changed or if the lineup members were photographed days, months, or years after an initial witnessing event.

Second, generalizability of these results is potentially limited by the relatively homogenous demographic characteristics of the participants (e.g., all self-reported their race as White). There is the possibility that language usage differs across (and within) demographic

groups (e.g., Ahmad et al., 2020), with consequences for the predictive properties of certain words. Incorporating these individual differences into future work examining verbal confidence classifiers is a crucial next step for research in this domain.

Conclusion

This study used machine-learning techniques to examine people's understanding and use of verbal expressions of confidence. Although such statements are often assumed to reflect the same underlying information as numeric confidence ratings, we showed that this is not the case. We find that verbal confidence is not redundant with numeric confidence, but can contribute unique diagnostic value in predicting the accuracy of an eyewitness identification. While we have applied this machine-learning methodology to eyewitness memory, we anticipate that these findings would extend to all domains in which people express confidence verbally or numerically. These techniques (with development of requisite theory) could eventually be applied to other areas, such as evaluating claims of intelligence analysts, climate forecasters, medical professionals, or other professions where verbal confidence expressions are ubiquitous.

References

- Ahmad, F., Abbasi, A., Li, J., Dobolyi, D. G., Netemeyer, R. G., Clifford, G. D., & Chen, H. (2020). A deep learning architecture for psychometric natural language processing. *ACM Transactions on Information Systems*, 38(1), 1–29. <https://doi.org/10.1145/3365211>
- Bates D, Mächler M, Bolker B, Walker S (2015). “Fitting Linear Mixed-Effects Models Using lme4.” *Journal of Statistical Software*, 67(1), 1–48. doi: 10.18637/jss.v067.i01.
- Barton, J. J. S., & Corrow, S. L. (2016). Recognizing and identifying people: A neuropsychological review. *Cortex*, 75, 132–150. <https://doi.org/10.1016/j.cortex.2015.11.023>
- Behrman, B. W., & Richards, R. E. (2005). Suspect/foil identification in actual crimes and in the laboratory: A reality monitoring analysis. *Law and Human Behavior*, 29(3), 279-301.
- Benjamin, A. S., Tullis, J. G., & Lee, J. H. (2013). Criterion noise in ratings-based recognition: Evidence from the effects of response scale length on recognition accuracy. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(5), 1601-1608.
- Brun, W., & Teigen, K. H. (1988). Verbal Probabilities : Ambiguous , or Both ? *Organizational Behavior and Human Decision Processes*, 41, 390–404.
- Budescu, D. V., Por, H. H., Broomell, S., & Smithson, M. (2014). The interpretation of IPCC probabilistic statements around the world. *Nature Climate Change*, 4, 508–512. <http://dx.doi.org/10.1038/nclimate2194>

- Burnham, K. P., & Anderson, D. R. (2002). Model selection and multimodel inference: A practical information-theoretic approach (2nd ed.). New York, NY: Springer-Verlag.
- Cash, D. K., & Lane, S. M. (2017). Context influences interpretation of eyewitness confidence statements. *Law and Human Behavior*, 41(2), 180–190.
<https://doi.org/10.1037/lhb0000216>
- Cash, D. K., & Lane, S. M. (2021). Judging memory: strong verbal confidence and lineup context influence inferences about eyewitnesses. *Psychology, Crime and Law*, 27(1), 40–60.
<https://doi.org/10.1080/1068316X.2020.1774588>
- Davies, G., Ellis, H. D., & Shepherd, J. (1978). Face recognition accuracy as a function of mode of representation. *Journal of Applied Psychology*, 63(2), 180–187.
<https://doi.org/10.1037/0021-9010.63.2.180>
- Dobbins, I. G., & Kantner, J. (2019). The language of accurate recognition memory. *Cognition*, 192, 103988. <https://doi.org/10.1016/j.cognition.2019.05.025>
- Dobolyi, D. G., & Dodson, C. S. (2013). Eyewitness confidence in simultaneous and sequential lineups: A criterion shift account for sequential mistaken identification overconfidence. *Journal of Experimental Psychology: Applied*, 19(4), 345–357.
<https://doi.org/10.1037/a0034596>
- Dobolyi, D. G., & Dodson, C. S. (2018). Actual vs. perceived eyewitness accuracy and confidence and the featural justification effect. *Journal of Experimental Psychology: Applied*, 24(4), 543–563. <https://doi.org/10.1037/xap0000182>

Dodson, C. S., & Dobolyi, D. G. (2015). Misinterpreting eyewitness expressions of confidence: The featural justification effect. *Law and Human Behavior*, 39(3), 266–280.

<https://doi.org/10.1037/lhb0000120>

Dodson, C. S., & Dobolyi, D. G. (2016). Confidence and eyewitness identifications: The cross-race effect, decision time and accuracy. *Applied Cognitive Psychology*, 30, 113–125.

doi:10.1002/acp.3178

Dodson, C. S., & Dobolyi, D. G. (2017). Judging guilt and accuracy: highly confident eyewitnesses are discounted when they provide featural justifications. *Psychology, Crime & Law*,

23(5), 487–508. <https://doi.org/10.1080/1068316x.2017.1284220>

Donaldson, W. (1996). The role of decision processes in remembering and knowing. *Memory & cognition*, 24(4), 523–533.

Duchaine, B., & 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>

Dunn, J. C. (2004). Remember-know: a matter of confidence. *Psychological review*, 111(2), 524–542.

Egan, J. P. (1958). Recognition memory and the operating characteristic (Technical Note AFCRC-TN-58-51). Bloomington: Indiana University, Hearing and Communication Laboratory.

- Fitzgerald, R. J., Rubinova, R., & Juncu, S. (2021). Eyewitness identification around the world. In A. M. Smith, M. P. Toglia, & J. M. Lampinen (Eds.), *Methods, measures, and theories in eyewitness identification tasks*. Taylor & Francis.
<https://doi.org/10.17605/OSF.IO/KN6R5>
- Gardiner, J. M., Ramponi, C., & Richardson-Klavehn, A. (1998). Experiences of remembering, knowing, and guessing. *Consciousness and cognition*, 7(1), 1-26.
- Gettleman, J. N., Grabman, J. H., Dobolyi, D. G., & Dodson, C. S. (2020). A decision processes account of the differences in the eyewitness confidence-accuracy relationship between strong and weak face recognizers under suboptimal exposure and delay conditions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 47(3), 402-421. <https://doi.org/10.1037/xlm0000922>
- Grabman, J. H., Cash, D. K., Slane, C. R., & Dodson, C. S. (*in press*). Improving the Interpretation of Verbal Eyewitness Confidence Statements by Distinguishing Perceptions of Certainty from Those of Accuracy. *Journal of Experimental Psychology: Applied*.
- Grabman, J. H., Dobolyi, D. G., Berelovich, 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>
- Grabman, J. H., & Dodson, C. S. (2019). Prior knowledge influences interpretations of eyewitness confidence statements: 'The witness picked the suspect, they must be 100%

sure'. *Psychology, Crime and Law*, 25(1), 50–68.

<https://doi.org/10.1080/1068316X.2018.1497167>

Grabman, J. H., & Dodson, C. S. (2020). 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*.

Hartig, F. (2020). DHARMA: Residual Diagnostics for Hierarchical (Multi-Level / Mixed) Regression Models. R package version 0.3.3.0. <http://florianhartig.github.io/DHARMA/>

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, 114(1), 3-28. <https://psycnet.apa.org/doi/10.1037/0033-2909.114.1.3>

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*, 22, 1304–1316. doi:10.1037/0278-7393.22.5.1304

Kantner, J., Solinger, L. A., Grybinas, D., & Dobbins, I. G. (2019). Confidence carryover during interleaved memory and perception judgments. *Memory & cognition*, 47(2), 195-211.

Klobuchar, A., Steblay, N. K., & Caligiuri, H. L. (2006). Improving eyewitness identification: Hennepin County's blind sequential lineup pilot project. *Cardozo Public Law, Policy, and Ethics Journal*, 2, 381–414.

- Lindsay, R. C., Martin, R., & Webber, L. (1994). Default values in eyewitness descriptions. *Law and Human Behavior, 18*(5), 527-541.
- Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). Lawrence Erlbaum Associates Publishers.
- Macmillan, N. A., & Kaplan, H. L. (1985). Detection theory analysis of group data: estimating sensitivity from average hit and false-alarm rates. *Psychological bulletin, 98*(1), 185.
- Mandler, G. (1980). Recognizing: The judgment of previous occurrence. *Psychological review, 87*(3), 252-271.
- Mansour, J. K. (2020). The Confidence-Accuracy Relationship Using Scale Versus Other Methods of Assessing Confidence. *Journal of Applied Research in Memory and Cognition, 9*(2), 215-231.
- 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.
- 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, 18*(4), 361–376.
<https://doi.org/10.1037/a0030609>

Mickes, L., Hwe, V., Wais, P. E., & Wixted, J. T. (2011). Strong memories are hard to scale.

Journal of Experimental Psychology: General, 140(2), 239–257.

<https://doi.org/10.1037/a0023007>

Mickes, L., Moreland, M. B., Clark, S. E., & Wixted, J. T. (2014). Missing the information needed

to perform ROC analysis? Then compute d' , not the diagnosticity ratio. *Journal of*

Applied Research in Memory and Cognition, 3(2), 58–62.

<https://doi.org/10.1016/j.jarmac.2014.04.007>

Mickes, L., Seale-Carlisle, T. M., & Wixted, J. T. (2013). Rethinking familiarity: Remember/Know

judgments in free recall. *Journal of memory and language*, 68(4), 333-349.

Mickes, L., Wixted, J. T., & Wais, P. E. (2007). A direct test of the unequal-variance signal

detection model of recognition memory. *Psychonomic Bulletin & Review*, 14(5), 858-

865.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: verbal reports on mental

processes. *Psychological review*, 84(3), 231-259.

Police and Criminal Evidence Act. (1984a). Codes of Practice, Code D. 2011. Retrieved from

<https://www.gov.uk/government/publications/pace-code-d-201>

R Core Team (2020). R: A language and environment for statistical computing. R Foundation for

Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.

Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J. C., & Müller, M. (2011).

pROC: An open-source package for R and Sto analyze and compare ROC curves. *BMC Bioinformatics*, 12, 77. [http:// dx.doi.org/10.1186/1471-2105-12-77](http://dx.doi.org/10.1186/1471-2105-12-77)

Russell, R., Duchaine, B., & Nakayama, K. (2009). Super-recognizers: People with extraordinary face recognition ability. *Psychonomic Bulletin & Review*, 16(2), 252–257.

<https://doi.org/10.3758/pbr.16.2.252>

Schooler, J. W., & Engstler-Schooler, T. Y. (1990). Verbal overshadowing of visual memories: Some things are better left unsaid. *Cognitive psychology*, 22(1), 36–71.

Seale-Carlisle, T. M. (2021, August 24). Language of accurate and inaccurate eyewitnesses. Retrieved from osf.io/s64dj

Seale-Carlisle, T. M., Colloff, M. F., Flowe, H. D., Wells, W., Wixted, J. T., & Mickes, L. (2019). Confidence and Response Time as Indicators of Eyewitness Identification Accuracy in the Lab and in the Real World. *Journal of Applied Research in Memory and Cognition*, 8(4), 420–428. <https://doi.org/10.1016/j.jarmac.2019.09.003>

Seale-Carlisle, T. M., Wetmore, S. A., Flowe, H. D., & Mickes, L. (2019). Designing police lineups to maximize memory performance. *Journal of Experimental Psychology: Applied*, 25(3), 410–430. <https://doi.org/10.1037/xap0000222>

Singmann, H. Bolker, B., Westfall, J., & Aust, F. (2016). afex: Analysis of Factorial Experiments. R package version 0.16-1. <https://CRAN.R-project.org/package=afex>

- Slane, C. R., & Dodson, C. S. (2019). Two eyewitnesses are more persuasive than one except when they remember a suspect's feature. *Journal of Applied Research in Memory and Cognition*, 8(1), 60–67. <https://doi.org/10.1016/j.jarmac.2018.12.002>
- Slotnick, S. D., & Dodson, C. S. (2005). Support for a continuous (single-process) model of recognition memory and source memory. *Memory & Cognition*, 33(1), 151–170. <https://doi.org/10.3758/BF03195305>
- Stankov, L., Kleitman, S., & Jackson, S. A. (2015). Measures of the Trait of Confidence. In *Measures of Personality and Social Psychological Constructs* (Issue 2014). <https://doi.org/10.1016/B978-0-12-386915-9.00007-3>
- Tan, C.-M., Wang, Y.-F., & Lee, C.-D. (2002). The use of bigrams to enhance text categorization. *Information Processing & Management*, 38(4), 529–546. [https://doi.org/https://doi.org/10.1016/S0306-4573\(01\)00045-0](https://doi.org/https://doi.org/10.1016/S0306-4573(01)00045-0)
- Tekin, E., Lin, W., & Roediger, H. L. (2018). The relationship between confidence and accuracy with verbal and verbal + numeric confidence scales. *Cognitive Research: Principles and Implications*, 3(1), 41. <https://doi.org/10.1186/s41235-018-0134-3>
- 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), 49. <https://doi.org/10.1186/s41235-017-0086-z>

- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288.
<https://doi.org/10.1111/j.2517-6161.1996.tb02080.x>
- Tulving, E. (1985). How many memory systems are there? *American psychologist*, 40(4), 385.
- Van Koppen, P. J., & Lochun, S. K. (1997). Portraying perpetrators; the validity of offender descriptions by witnesses. *Law and Human Behavior*, 21(6), 661–685.
- Wallsten, T. S., Budescu, D. V., Zwick, R., & Kemp, S. M. (1993). Preferences and reasons for communicating probabilistic information in verbal or numerical terms. In *Bulletin of the Psychonomic Society* 31(2), 135–138. <https://doi.org/10.3758/BF03334162>
- Wilmer, J. B. (2017). Individual Differences in Face Recognition: A Decade of Discovery. *Current Directions in Psychological Science*, 26(3), 225–230.
<https://doi.org/10.1177/0963721417710693>
- Wilson, B. M., Seale-Carlisle, T. M., & Mickes, L. (2018). The effects of verbal descriptions on performance in lineups and showups. *Journal of Experimental Psychology: General*, 147(1), 1–12.
- Wixted, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*, 114(1), 152–176. <https://doi.org/10.1037/0033-295x.114.1.152>

- Wixted, J. T., & Mickes, L. (2010). A Continuous Dual-Process Model of Remember/Know Judgments. *Psychological Review*, 117(4), 1025–1054.
<https://doi.org/10.1037/a0020874>
- Wixted, J. T., Mickes, L., Clark, S. E., Gronlund, S. D., & Roediger III, H. L. (2015). Initial eyewitness confidence reliably predicts eyewitness identification accuracy. *American Psychologist*, 70(6), 515-526.
- Wixted, J. T., Mickes, L., Dunn, J. C., Clark, S. E., & Wells, W. (2016). Estimating the reliability of eyewitness identifications from police lineups. *Proceedings of the National Academy of Sciences*, 113(2), 304-309.
- 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>
- Yates, S. Q. (2017). Eyewitness identification: Procedures for conducting photo arrays. *US Department of Justice*.
- Yilmaz, A. & Wixted, J. T. (under review). The Reveal Procedure: A Way to Enhance Evidence of Innocence from Police Lineups.
- Yonelinas, A. P. (1994). Receiver-operating characteristics in recognition memory: evidence for a dual-process model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(6), 1341-1354.

Zimmer, A. C. (1983). Verbal vs. numerical processing of subjective probabilities. In R. W. B. T.-
A. in P. Scholz (Ed.), *Decision Making Under Uncertainty*, 16, 159 = 182. North-Holland.
[https://doi.org/https://doi.org/10.1016/S0166-4115\(08\)62198-6](https://doi.org/https://doi.org/10.1016/S0166-4115(08)62198-6)

Accepted Version

Appendix A

Figure A1. Regression coefficients for words retained in the Bigram HFA Classifier.

Note. Left panel illustrates the bigrams predictive of a correct ID. Right panel illustrates bigrams predictive of a filler ID. In both panels, the bigrams most predictive are displayed at the top and the bigrams least predictive are displayed at the bottom.

Appendix B

Figure B1. Regression coefficients for words retained in the Trigram HFA Classifier.

Note. Left panel illustrates the trigrams predictive of a correct ID. Right panel illustrates trigrams predictive of a filler ID. In both panels, the trigrams most predictive are displayed at the top and the trigrams least predictive are displayed at the bottom.