A model of position effects in the sequential lineup

by

John C. Dunn^{1,2}, Matthew Kaesler³, Carolyn Semmler³

¹ University of Western Australia

² Edith Cowan University

³ University of Adelaide

Running head: Target position model

Address for correspondence:

Prof John C. Dunn
University of Western Australia
35 Stirling Hwy
Crawley, WA, 6009
Email: john.dunn@uwa.edu.au

Word count: 10,305

Author Note:

This work was supported by the Australian Research Council grant DP160101048. Data from the experiment may be downloaded from <https://osf.io/ujfxd/>.

Keywords:

Eyewitness memory, sequential lineup, signal detection theory, position effects, discriminability.

Abstract

What is the effect of placing the suspect in different positions in a sequential lineup?

To explore this question, we developed and applied a model called the Independent Sequential Lineup model which analyzes a sequential lineup in terms of both identification position, the position at which the witness identifies a lineup item as the target, and target position, the position at which the target or suspect appears. We conducted a large-scale online eyewitness memory experiment with 7,204 participants each of whom was tested on a 6-item sequential lineup with an explicit stopping rule. The model fit these data well and revealed systematic effects of lineup position on underlying discriminability and response criteria. We also fit the model to data from a similar pair of experiments conducted recently by Wilson, Donnelly, Christenfeld and Wixted (2019; *Journal of Memory and Language*, 104, 108-125) both with and without application of a stopping rule. In all data sets, if a stopping rule is applied, underlying discriminability was found to be constant, or to increase slightly, across target position. In the absence of a stopping rule, discriminability was found to decrease substantially. We also observed a substantial increase in response criteria following presentation of the target. We discuss the implications of these findings for current theories of recognition memory and current applications of the sequential lineup in different jurisdictions.

A police lineup may be presented simultaneously or sequentially. In a sequential lineup, items are presented one at a time. On each presentation, the witness is asked if the item is the target. If a stopping rule is used, as soon as an item is identified, whether correctly or incorrectly, the lineup is terminated. Otherwise, the next item is presented and the procedure repeated until either an identification is made or all the items have been presented. The sequential lineup was first proposed to reduce the risk of mistaken identification (Lindsay & Wells, 1985) and is currently used by police in several different jurisdictions, albeit implemented with considerable variation.

An important and as yet unresolved question concerns the effect of the position of the target¹ on identification performance in a sequential lineup with a stopping rule. Do some positions lead to better performance in terms of higher hit rates and lower false alarm rates than others? An additional question is whether false alarm rates change systematically, either increasing or decreasing, over the course of the lineup. Research addressing these questions has yielded mixed results. Some early studies found that choosing rates (the proportion of times witnesses attempted an identification) did not differ across target position if the size of the lineup was concealed from the witness (Lindsay & Wells, 1985; Sporer, 1993). However, more recent studies using a range of different methodologies have found that target position may affect overall choosing rates, as well as independently affecting rates of target and non-target identification (e.g., Gronlund, Carlson, Dailey & Goodsell, 2009; Meisters, Diedenhofen & Musch, 2018; Wilson, Donnelly, Christenfeld & Wixted, 2019). Aside from practical concerns, these results raise important theoretical questions about how such effects might occur. In terms of signal detection theory, the effect of target position on performance

¹ In our usage, the target refers to the suspect in a target-present lineup. Non-targets or foils refer to all other lineup members.

may be due to changes in either response bias or underlying discriminability, or both. Given the variety of analysis techniques used in existing research and the mixed pattern of results, the extent to which response bias and underlying discriminability change over the course of the sequential lineup remains unclear. Our aim in this paper is to develop a model of the sequential lineup that can be used to estimate discriminability and response bias at each target position and to examine target position effects in new experimental data and in data from Wilson et al. (2019)². We hope to clarify the effects of target position in the sequential lineup and to shed light on the mechanisms by which they occur.

Prior to the introduction of Receiver Operating Characteristic (ROC) analysis to lineup research (Mickes, Flowe, & Wixted, 2012), position effects were investigated using several different approaches. Some studies compared the rates of different decisions at each position. Using this approach, Clark and Davey (2005, Experiment 2) found an increase in overall choosing rates across lineup position. Similarly, Carlson et al. (2008, Experiment 2) found that target identification rates were greater at later compared to earlier positions. They did not find changes in overall choosing rates in target absent lineups although innocent suspect identifications tended to be lower at later positions. Other studies evaluated the *diagnosticity ratio* – the ratio of target to innocent suspect identification rates – at each position, observing that this tended to increase with target position (e.g., Carlson, Gronlund, & Clark, 2008; Gronlund et al., 2009).

The effect of target position on choosing rates and the diagnosticity ratio is difficult to interpret because of the impossibility of attributing it to a change in memory strength or to a

² We are grateful to Brent Wilson and colleagues for making the data from their two experiments available on the OSF at <https://osf.io/bq5y7/>.

change in response bias or both (Mickes & Wixted, 2012). To distinguish between these different possibilities, researchers have more recently turned to the analysis of ROC curves.

ROC Analysis

An important feature of ROC analysis of eyewitness memory is the ROC curve. This is a plot of the cumulative proportion of correct *target identifications* (TID) from a target present lineup against the cumulative proportion of incorrect innocent *suspect identifications* (SID) from a corresponding target absent lineup. A measure of performance, *empirical discriminability*, can also be obtained by calculating the area under the ROC curve (AUC). Several studies have investigated target position effects by constructing an ROC curve separately for each position. While some have reported greater empirical discriminability for later target positions compared to earlier positions (Gronlund, Carlson, Neuschatz, Goodsell, Wetmore, Wooten & Graham, 2012; Meisters et al., 2018; Neuschatz, Wetmore, Key, Cash, Gronlund & Goodsell, 2016), others have reported no difference (Carlson & Carlson, 2014; Dobolyi & Dodson, 2013).

In the analysis of identification performance, it is important to distinguish between empirical discriminability and *underlying discriminability* or memory strength (Wixted & Mickes, 2018). This is because the shape of an ROC curve, and hence empirical discriminability as measured by AUC, may be affected by characteristics of the identification task that are unrelated to memory strength (Rotello & Chen, 2016; Kaesler, Dunn, Ransom & Semmler, 2020). This effect was shown by Wilson et al. (2019) who simulated ROC curves for targets and suspects appearing at different lineup positions based on the first-above criterion sequential lineup model described in detail in Kaesler et al. (2020). Although underlying discriminability was held constant across all target positions, the resulting ROC

curves varied substantially and empirical discriminability as measured by AUC was found to decline substantially across target position (see Figure 6A of Wilson et al., 2019).

Model-Based Studies

To measure underlying discriminability as well as response bias, it is necessary to use a theoretical model. Within the framework of signal detection theory, underlying discriminability is identified with d' , the mean difference in familiarity between targets and non-targets, and response bias is identified with c , the decision criterion. Different models within this framework map these parameters onto identification rates in different ways. Two such models that have been used by researchers are the SDT-MAX and SDT-INT models³ described in detail by Kaesler et al (2020). These were initially proposed as models of the simultaneous lineup and implement different rules on which witnesses may base their decision to identify an item. In both models, if the witness decides to identify an item, they select the lineup item with the greatest familiarity – the one that most closely matches their memory of the perpetrator or target. In the SDT-MAX model, this decision is triggered if the item with highest familiarity exceeds a specified decision criterion. In the SDT-INT model, the decision is triggered if the total familiarity of all the items exceeds the decision criterion⁴.

Carlson, Carlson, Weatherford, Tucker, and Bednarz (2016) used the SDT-MAX model (which they called the SDT-CD model) to estimate underlying discriminability which they found to be greater when the target was in position 2 in a six-member sequential lineup

³ SDT-MAX is also called the independent observations model, and SDT-INT is also called the integration model.

⁴ Other decision rules are possible. For example, Wixted, Vul, Mickes and Wilson (2018) have proposed an *ensemble rule* in which a decision is triggered if the item with highest familiarity exceeds the mean familiarity of the remaining items by a specified amount.

than when it was in position 5, as long as witnesses were not told how many lineup members to expect. Horry, Palmer, and Brewer (2012) using the SDT-INT model also found that underlying discriminability was greater for targets in position 2 than in position 6 although, in this case, participants knew how many lineup items to expect (either 6, 12 or 30). In addition, both studies reported a decrease in the decision criterion over the course of the lineup based on choosing rates, although this was not reflected in the estimated model parameters.

The results reported by Carlson et al. and Horry et al. are, however, difficult to interpret. This is because the models they used are not easily applied to a sequential lineup with a stopping rule as they incorporate decision rules that depend on access to and comparison of the familiarity or memory strength of all lineup items, either to calculate the maximum strength (for SDT-MAX) or the total strength (for SDT-INT). This logic does not apply to a sequential lineup in which witnesses make an identification decision after each item is presented because access to items not yet presented is impossible (Kaesler et al., 2020).

SDT-SEQ Model

Recently, Kaesler et al. (2020) described a signal detection model that can be applied to a sequential lineup with a stopping rule. This model, which they called the SDT-SEQ model, proposes that witnesses identify the *first* item in the lineup whose familiarity exceeds a decision criterion and was used by Wilson et al. (2019) to estimate underlying discriminability and response bias at each position of a six-member lineup. Wilson et al. also used ROC analysis to calculate empirical discriminability for each target position and, as expected, they found a dissociation between this and underlying discriminability. In their experiment 1, they found that empirical discriminability decreased across target position

while underlying discriminability *increased* from position 1 to position 2 after which it remained constant. In contrast, in their experiment 2, they did not find any change in underlying discriminability across target position. They also found that while response bias did vary across target position, there was no obvious pattern of either increase or decrease.

The change in underlying discriminability found by Wilson et al. (2019, experiment 1) is of theoretical interest because it is consistent with a key prediction of Diagnostic Feature Detection Theory (DFDT) proposed by Wixted and Mickes (2014). This proposes that underlying discriminability is a function of the extent to which a witness can isolate unique, diagnostic features shared only by their memory of the perpetrator and the lineup target. Non-diagnostic features are those shared by the target and all other lineup items and so cannot be used to identify the target. Which features are diagnostic and which non-diagnostic depends on determining what the lineup items have in common. For example, if the target has dark hair but so have all the lineup items then this feature is non-diagnostic. The features shared by all lineup items in a simultaneous lineup is immediately apparent but must be built up item-by-item when the lineup is presented sequentially. The result found by Wilson et al. that there is no increase in underlying discriminability after position 2 suggests that comparison of the first two items may be sufficient to identify and to eliminate from consideration non-diagnostic features.

Although the methodology used by Wilson et al. (2019) represents a major advance, limitations of the SDT-SEQ model meant that they were unable to address some of their key questions. The SDT-SEQ model was developed by Kaesler et al. (2020) to account for sequential lineup identifications that have been aggregated over both identification position (where in the lineup an identification is made) and target position (where in the lineup the target appears). For the model to be tractable, it assumes that the probability of a non-target

identification is the same at all identification positions. We will see later that there is strong evidence that this assumption is false and may account for the poor model fits found by Wilson et al. (2019).

If the SDT-SEQ model is applied separately to identification rates at each target position, two further complications result. First, because the model requires foil identifications in a target-present (TP) lineup to be aggregated over position, these necessarily correspond to different sets of positions for each target location. For example, in a six-item lineup, if the target is in position 1 then foil identifications are drawn from positions 2 to 6. But if the target is in position 2 then foil identifications are drawn from position 1 and from positions 3 to 6. This means that if the model is applied to data from each target position to see if these differ in underlying discriminability, interpretation is confounded by possible changes in response criteria across lineup position. Second, Wilson et al. compared TP identifications for each target position with TA identifications at the corresponding identification position. This leads to the problem that the corresponding count of correct rejections of TA lineups cannot be uniquely allocated to any target position. If the decision is made to assign it to each position then the observations are no longer independent, compromising the resulting statistical analysis.

The Independent Sequential Lineup (ISL) Model

To examine the effect of target position on sequential lineup identifications, we propose a new probabilistic model, which we call the *Independent Sequential Lineup (ISL) model*. The unique contribution of this model is that it reflects both the position of the target *and* the position at which an identification is made. As we explain below, it can be viewed as a generalization of the SDT-SEQ model that overcomes the problems identified above. Table

1 shows the structure of this model for a 6-item lineup. Each column corresponds to a position j of the target in a TP lineup, with the final column, labelled “none”, corresponding to a TA lineup. Notionally, the TA lineup can be thought of as the special case of a target presented at position $n + 1$, where n is the lineup size. Each row of the table corresponds to the position i at which an identification is made, with the last row, also labelled “none”, corresponding to no identification or rejection of the lineup. We presume that if an identification is made then it can be associated with some level of confidence k . The model therefore has three kinds of parameter; $(1 - G_{ik})$ is the probability of identifying a target at position $i = j$ with confidence level k ; $(1 - F_{ik})$ is the probability of identifying a foil at position $i < j$ with confidence level k ; and $(1 - F_{ik}^*)$ is the probability of identifying a foil at position $i > j$ with confidence level k . We write $(1 - G_i)$, $(1 - F_i)$ and $(1 - F_i^*)$ for the corresponding probabilities of identifying targets and foils at the lowest confidence level – the minimal level at which an item is identified. We distinguish between foil identifications made before and after the presentation of the target for the simple reason that the model does not fit the data otherwise. Later, we investigate possible reasons for this. Lastly, in order for the model to be identifiable, we also assume that both the pre-target false alarm rate, F_{ik} , and the post-target false alarm rate, F_{ik}^* , while allowed to vary across identification position, do not change with target position. In other words, the probability of identifying a foil in position 1 does not depend upon whether the target will appear in position 2 or any later

position. Similarly, the probability of identifying a foil in position 6 does not depend upon whether the target was presented in position 5 or any earlier position⁵.

Because of the stopping rule, the probability of identifying an item at identification position i is equal to the joint probability of identifying this item and having not identified any previous item. We assume that these probabilities are independent and is why we call the model the *independent sequential lineup* model. It functions as a measurement model that enables the underlying hit and false alarm rates to be extracted from the observed identification rates. These, in turn, can be interpreted using detection theory which we describe later.

Formally, the ISL model is described as follows. Let $P_{TD}(k)_j$ be the conditional probability of correctly identifying the target presented at position j with confidence level k . Suppose $j = 2$. Then $P_{TD}(k)_2 = F_1(1 - G_{2k})$. That is, it is equal to the product of F_1 the probability of not identifying the foil presented in position 1 and $1 - G_{2k}$ the probability of identifying the target in position 2 at confidence level k . More generally, let $P_{ID}(k)_{ij}$ be the conditional probability of identifying item i at confidence level k given that item j is the target. Then

$$P_{ID}(k)_{ij} = (1 - Q_{ijk}) \prod_{m=1}^{i-1} Q_{mj}, \quad (1)$$

where,

⁵ Note that this is necessarily true for pre-target false alarm rates because participants cannot predict which item is the target. In contrast, it is possible that post-target false alarm rates may vary systematically.

$$Q_{ijk} = \begin{cases} F_{ik}, & i < j \\ G_{ik}, & i = j \\ F_{ik}^*, & i > j \end{cases} \quad \text{and} \quad Q_{ij} = \begin{cases} F_i, & i < j \\ G_i, & i = j \\ F_i^*, & i > j. \end{cases}$$

It follows that

$$P_{TID}(k)_j = (1 - G_{jk}) \prod_{m=1}^{j-1} F_m. \quad (2)$$

Let $P_{FID}(k)_{ij}$ be the probability of identifying a foil in position i at confidence level k in a target present lineup given that the target is in position j . Then, from Equation (1),

$$P_{FID}(k)_{ij} = \begin{cases} (1 - F_{ik}) \prod_{m=1}^{i-1} F_m, & i < j \\ (1 - F_{ik}^*) G_j \prod_{m=1}^{j-1} F_m \prod_{m=j+1}^{i-1} F_m^*, & i > j. \end{cases} \quad (3)$$

Let P_j^{MISS} be the probability of failing to identify any item in an n -item lineup given that item j is the target. Then

$$P_j^{MISS} = G_j \prod_{m=1}^{j-1} F_m \prod_{m=j+1}^n F_m^*. \quad (4)$$

Let $P_{FA}(k)_i$ be the probability of identifying item i at confidence level k given that no item is the target. Then

$$P_{FA}(k)_i = (1 - F_{ik}) \prod_{m=1}^{i-1} F_m. \quad (5)$$

Finally, let P^{CR} be the conditional probability of failing to identify any item in an n -item lineup given that no item is the target. Then

$$P^{CR} = \prod_{m=1}^n F_m . \quad (6)$$

Each column in Table 1 contains the set of conditional probabilities for a given target position, i.e., each column in Table 1 sums to one. To convert these to absolute probabilities, each column entry should be multiplied by p_j , the probability that the target appears in position j in a TP lineup, or the probability of a TA lineup.

The ISL model reduces to the SDT-SEQ model if the false alarm rates are assumed to be constant across lineup position. That is, $F_{ik} = F_{ik}^* = F_k$, and $F_i = F_i^* = F$, for all i and k . Then by substituting into the relevant equations, we can derive the equations for the SDT-SEQ model. From Equation (2),

$$P_{TD}(k)_j = (1 - G_{jk}) F^{j-1},$$

and from Equation (3),

$$P_{FID}(k)_{ij} = \begin{cases} (1 - F_k) F^{i-1}, & i < j \\ (1 - F_k) G_j F^{i-2}, & i > j \end{cases}.$$

As mentioned previously, foil identification rates in the SDT-SEQ model are aggregated over position. Thus, the quantity of interest for fitting the model given that the target is in position j is,

$$P_{FID}(k)_j = (1 - F_k) \left[\sum_{i=1}^{j-1} F^{i-1} + \sum_{i=j+1}^n G_j F^{i-2} \right].$$

Similarly, false alarm rates are also aggregated over position. Thus, from Equation (5)

$$P_{FA}(k) = (1 - F_k) \sum_{i=1}^n F^{i-1}.$$

And finally, from Equation (6),

$$P^{CR} = F^n.$$

In its present form, the ISL model is not a signal detection model. However, it can be recast in these terms by treating $1 - G_{ik}$ as a hit rate and $1 - F_{ik}$ and $1 - F_{ik}^*$ as pre-target and post-target false alarm rates, respectively, for each identification position i and confidence level k . If the corresponding distributions are assumed to be normal then we can parameterize the model as follows,

$$\begin{aligned} F_{ik} &= \Phi(c_{ik}) \\ F_{ik}^* &= \Phi(c_{ik}^*) \\ G_{ik} &= \Phi\left(\frac{c_{ik} - d_i}{s_i}\right), \end{aligned} \tag{7}$$

where c_{ik} and c_{ik}^* are decision criteria corresponding to confidence level k at identification position i , d_i is the difference between the means of the target and foil familiarity distributions at target (equivalently, identification) position i , and s_i is the standard deviation of the target distribution (relative to the foil distribution) at target (equivalently, identification) position i . We call the model described by Equation (7), the SDT-ISL model. Although it is obviously related to the ISL model, the two models are not identical as they make different assumptions and have different numbers of free parameters. If there are n lineup items and $K + 1$ confidence categories then the ISL model has $(3n - 1)K$ parameters

while the SDT-ISL model has $(2n-1)K + 2n$ parameters. For $K \leq 2$, the two models are identical. Furthermore, the SDT-ISL model is identifiable only for $K > 1$.

Aims and Hypotheses

The aim of this paper is to measure the effects of target and identification position in a sequential lineup on the probabilities of target and foil identification rates using the ISL model, and on underlying discriminability (d) and response bias (c and c^*) using the SDT-ISL model. This will enable us to test the prediction of the diagnostic feature detection hypothesis that underlying discriminability should increase over the course of the sequential lineup as well as the hypothesis, based on earlier research, that response bias may decrease across lineup position (Carlson et al., 2016; Carlson et al., 2008; Clark & Davey, 2005; Horry et al., 2012). In the experiment described below, participants were given a sequential lineup task with an explicit stopping rule that terminated the lineup as soon an identification was made. In contrast, Wilson et al. (2019) did not employ an explicit stopping rule and participants responded to each lineup item by rating their confidence that it is the target. A stopping rule was applied retrospectively, counting as an identification the first confidence rating that exceeded a specified value. This difference in procedure raises the question whether the results found by Wilson et al. would be the same if participants had been instructed that the lineup would terminate once an identification was made. It is possible that with an explicit stopping rule, participants may respond more conservatively. Alternatively, previous research has shown that when witnesses are permitted to make multiple identifications from a lineup, overall accuracy is reduced (Horry, Brewer, Weber, & Palmer, 2015; Steblay, Dietrich, Ryan, Raczynski, & James, 2011). This suggests that underlying discriminability may decline across the lineup. To address these questions, we fit the ISL and SDT-ISL models to the data from our experiment as well as to the data from Wilson et al. (2019).

Experiment

In this experiment, participants studied a short video involving the target individual. Their memory for the target was subsequently tested by a six-item sequential lineup with an explicit stopping rule. In a TP lineup, the target could be presented in any position from 1 to 6.

Participants

Participants were $N = 7,525$ Amazon Mechanical Turk workers who were compensated \$1 USD for the five- to ten-minute experiment. We excluded 321 participants for failing the attention check questions relating to the content of the stimulus video, leaving $N = 7,204$ ($N_{TP} = 5,784$, $N_{TA} = 1,420$) participants whose responses were available for analysis. Two cohorts of participants were tested over two separate periods. Data from the first cohort ($N = 1,846$; $N_{TP} = 952$, $N_{TA} = 894$) were collected in May 2019. At this point, it became clear to us that because the TP lineup required separate analysis of each target position, it would be necessary to present more TP lineups than TA lineups to participants. For this reason, we recruited a second cohort of participants ($N = 5,358$; $N_{TP} = 4,832$, $N_{TA} = 526$) in January 2020. The data from both cohorts were pooled for analysis.

Materials

This study used the stimulus set described by Kaesler et al. (2020). This consisted of a set of headshots and videos of sixteen females selected from a larger corpus of ninety candidates based on their overall similarity to each other. For example, all were young and dark-haired (for further details, see Kaesler et al. 2020). All sixteen items acted as both targets and foils across different six-item lineups. This approach was designed to mitigate the

possibility of spurious stimulus effects that may result from using a single target and single set of accompanying foils.

Procedure

The procedure was largely the same as described by Kaesler et al. (2020) for their sequential lineup condition. The entire procedure took place within Amazon Mechanical Turk (AMT) using the participants' web browsers. Participants were allocated to either a target present or target absent lineup on a round robin basis. After viewing instructions and providing their consent to participate, participants were required to correctly answer several comprehension questions, or view the instructions again, before proceeding to the experiment. They were then shown a short video involving a target randomly selected from the sixteen-member pool, followed by a visual search distractor task. Participants were shown pre-lineup instructions corresponding to those in the U.S. National Department of Justice (1999) guidelines, including instructions that the procedure would terminate as soon as they made an identification. They were then shown a sequentially presented target present or target absent lineup, with the target (if any) and the appropriate number of foils (five for target present, six for target absent) randomly selected from the remaining fifteen members of the stimulus pool. The positions of the target (if any) and foils were selected at random.

Participants saw each lineup item individually with an option either to identify it as the target or to reject it and move onto the next item (if any remained). If an item was identified, the procedure was terminated and the participant provided a typed confidence estimate for their identification on a 0-100 scale and a written confidence estimate in words. If all lineup items were rejected, participants were informed that the lineup was exhausted, indicating a rejection decision, and they then provided a typed confidence rating of this

decision on a 0-100 scale as well as a written confidence estimate in words. Participants also provided a written justification for their identification or rejection decision. They then answered follow-up questions about instruction clarity and task difficulty and given the opportunity to provide feedback.

Analyses

The ISL model was fit using the *fmincon* function in Matlab[®]. Model fit was tested using a χ^2 test that compared observed counts against those predicted by parameter estimates derived from the model.

We employed likelihood ratio tests to test for differences in parameters. We compared the χ^2 value of an unconstrained model to the χ^2 value of a model with one or more parameters constrained either to be equal to each another or to a fixed value. Each constraint added a degree of freedom to the nested model. The resulting difference in χ^2 values is distributed as χ^2 with degrees of freedom given by the number of constraints.

Results and Discussion

Table 2 presents the observed frequency of identifications in the format used by the ISL model. The ISL model fit these data well, $\chi^2(25) = 31.5$, $p = 0.17$. In order to determine whether it is necessary to distinguish the pre-target F_i and post-target F_i^* parameters, we fit the model under the constraint that these should be equal across identification positions 2 to 6 (as F_i^* is not defined at position 1). The fit of this constrained model was poor, $\chi^2(30) = 88.2$, $p < 0.001$, and the corresponding likelihood ratio test was also highly significant, $\chi^2(5) = 56.7$, $p < 0.001$. This tells us that pre-target false alarm rates differed from corresponding post-target false alarm rates.

Figure 1a plots identification rates, $1 - G_i$, $1 - F_i$, and $1 - F_i^*$, as functions of identification position, i . Figures 1b and 1c plot corresponding parameter estimates from our re-analysis Wilson et al. (2019) which we describe in the following section. For our experiment (Figure 1a), visual examination reveals a robust target effect – the probability of identifying the target is greater than the probability of identifying a foil at all identification positions. There is also little apparent change in the probability of a target identification across identification position and the corresponding likelihood ratio test is not significant, $\chi^2(5) = 7.18$, $p = 0.21$. Although subject to floor effects, the probability of identifying a pre-target foil, $1 - F_i$, appears to decrease across identification position, especially between positions 1 and 2. This is confirmed by a significant likelihood ratio test, $\chi^2(5) = 50.5$, $p < 0.001$. However, if these values are constrained to be equal only across positions 2 to 6, the resulting test is not significant, $\chi^2(4) = 4.72$, $p = 0.32$, confirming the impression that only position 1 differs from the others. In contrast to the decreasing trend for pre-target foils, the probability of identifying a post-target foil appears to increase with identification position although the relevant likelihood ratio test was not significant, $\chi^2(4) = 7.66$, $p = 0.105$.

To explore this effect further and to examine possible changes in underlying discriminability, we fit the ISL model to data categorized according to confidence. The higher false alarm rate in position 1 compared to positions 2 to 6 suggests a shift in response bias from relatively lenient in position 1 to more conservative at later positions.

Participants rated their confidence in an identification on a 0-100 scale with 100 representing maximal confidence that the identified item is the target. We pooled these responses and formed five approximately equal-sized confidence categories. The minimum

confidence rating of each of these categories were 0, 60, 75, 85, and 95, corresponding to $k = 1, 2, 3, 4$, and 5, respectively, in the model.

The ISL model fit the confidence data well, $\chi^2(125) = 118.4, p = 0.65$. Figure 2 shows the estimated values of the hit rates and false alarm rates in the form of an ROC plot (Figure 2a) and a zROC plot (Figure 2b). In the zROC plot, each of the axes of the ROC plot has been transformed by the inverse normal cdf, Φ^{-1} . Each point corresponds to a target position and confidence category, with points corresponding to target position 1 explicitly labelled. The shape and fill of the points indicate the confidence category (indexed by k) as described in the accompanying legend. Figure 2a suggests that the points corresponding to target position 1 may lie on a separate ROC curve to those corresponding to target positions 2 to 6. We have highlighted this by fitting points corresponding to position 1 by one theoretical curve (the dashed line) and fitting the points corresponding to the remaining target positions by a separate curve (the solid line). These curves are based on the zROC plot shown in Figure 2b. If Equation (7) holds and if $d_i = d$ and $s_i = s$ for a given set of points, indexed by i , then these points should lie on a straight line with a slope equal to $1/s$ and an intercept equal to d/s . As shown in Figure 2b, points corresponding to positions 2 to 6 appear to fall on one line, with estimated $d_{2-6} = 1.84$ and $s_{2-6} = 0.76$, while points corresponding to position 1 appear to fall on a separate line, with estimated $d_1 = 1.55$ and $s_1 = 0.63$. Plugging these two sets of values into Equation (7) generated the theoretical curves shown Figure 2a.

In order to confirm the impressions suggested by the ROC and zROC plots in Figure 2, we fit the SDT-ISL model parameterized according to Equation (7) to the same confidence data. This model, like the ISL model, fit these data well, $\chi^2(143) = 131.6, p = 0.74$. Figure 3a plots selected parameter values as a function of identification position. Parameters c_i and c_i^*

correspond to the most lenient response criteria ($k = 1$) that separate identifications from non-identifications. Because possible changes in the variance of the target distribution, s_i^2 , across identification position, the values of d_i are not directly comparable. For this reason, we chose to plot $d_{a,i}$, a measure of underlying discriminability scaled according to the variance of both target and foil distributions (Macmillan & Creelman, 2005). That is,

$$d_{a,i} = \frac{d_i}{\sqrt{\frac{1}{2}(1 + s_i^2)}}.$$

As suggested by Figure 3a, and consistent with the previous ROC and zROC analyses, underlying discriminability changes with position, particularly between position 1 and the remaining positions. In order to investigate this further, we tested a sub-model in which d_i and s_i were constrained to be constant over target position i . This model fit more poorly, $\chi^2(153) = 194.3$, $p = 0.013$, and the likelihood ratio test, comparing this model to the unconstrained model, was highly significant, $\chi^2(10) = 62.4$, $p < 0.001$. We also fit a model in which d_i and s_i were constrained to be constant over target positions 2 to 6. This model fit the data well, $\chi^2(151) = 141.2$, $p = 0.71$, and the likelihood ratio test was not significant, $\chi^2(8) = 9.29$, $p = 0.32$. Furthermore, the estimated values of d and s were also very close to those estimated from the zROC curve. These were, $d_{2-6} = 1.82$, $s_{2-6} = 0.71$, and $d_1 = 1.54$, $s_1 = 0.64$.

We also tested a sub-model in which the values of $d_{a,i}$ were constrained to be constant, a less strict requirement than holding both d_i and s_i constant. While this sub-model fit the confidence data well, $\chi^2(148) = 159.7$, $p = 0.24$, when compared to the fit of the unconstrained model, the likelihood ratio test was highly significant, $\chi^2(5) = 27.8$, $p < 0.001$.

In addition, when $d_{a,i}$ was constrained to be constant only over positions 2 to 6, the resulting likelihood ratio test was not significant, $\chi^2(4) = 0.80, p = 0.94$.

Reanalysis of Wilson et al. (2019)

We re-analyzed the data from the two lineup experiments conducted by Wilson et al. (2019) using the ISL and SDT-ISL models in order to determine the extent to which our results generalize to their data and the effects, if any, of an explicit stopping rule.

Data Considerations and Model Fit Performance

As Wilson et al. (2019) employed a 1:1 ratio of target present to target absent lineups, there were relatively fewer target present responses in both experiment one ($N_{TP} = 3,196$, $N_{TA} = 3,186$)⁶ and the lineup condition of experiment two ($N_{TP} = 994$, $N_{TA} = 967$)⁷ compared to our experiment ($N_{TP} = 5,784$, $N_{TA} = 1,420$). Recall that participants in Wilson et al. (2019) provided an identification decision and a confidence rating for all items. Following their approach, we imposed a virtual stopping rule by counting only the first identification response from each participant. For example, if a participant identified both an item at position two and an item at position five, they were treated as having identified only the item at position two. Table 3 and Table 4 show the observed frequencies from Wilson et al. (2019) experiment one and from the lineup condition of experiment two in ISL model format.

The ISL model fit the identification data from both experiment one, $\chi^2(25) = 26.6$, $p = 0.38$, and experiment two, $\chi^2(25) = 28.1$, $p = 0.31$. Figures 1b and 1c show hit rates, and pre- and post-target false alarm rates, plotted as functions of identification position for

⁶ These totals exclude 62 TP participants and 86 TA participants who had missing data.

⁷ These totals exclude 5 TP participants who had missing data.

experiment one (Figure 2b) and for experiment two (Figure 2c). As we found in our experiment, pre-target false alarm rates were significantly greater than post-target false alarm rates in both experiments from Wilson et al.; for experiment one, $\chi^2(5) = 45.8, p < 0.001$, and for experiment two, $\chi^2(5) = 21.0, p < 0.001$. There was a significant effect of position on hit rates for experiment one, $\chi^2(5) = 13.3, p = 0.021$, but not for experiment two, $\chi^2(5) = 2.23, p = 0.816$, and there was also a significant effect of position on pre-target false alarm rates in experiment one, $\chi^2(5) = 42.8, p < 0.001$, and (marginally) in experiment two, $\chi^2(5) = 11.3, p = 0.047$. The effect of position on post-target false alarm rate was significant for experiment one, $\chi^2(4) = 15.8, p = 0.003$, but not for experiment two, $\chi^2(4) = 8.54, p = 0.074$.

As we did for our experiment, we created five approximately equal-sized confidence categories for each of the two experiments conducted by Wilson et al, based on the distribution of confidence ratings. For experiment one, the minimum confidence ratings were 0, 40, 61, 75, and 90, while for experiment two they were 0, 51, 70, 80, and 91.

The ISL model fit the confidence data well for both experiments; for experiment one, $\chi^2(125) = 125.1, p = 0.48$, and for experiment two, $\chi^2(125) = 105.3, p = 0.899$. Figure 4 shows the corresponding zROC plots. As in Figure 2, we have shown separate estimates of the best-fitting straight lines separately for position 1 and for positions 2 to 6. For experiment one, the corresponding parameter estimates were $d_1 = 1.50, s_1 = 0.72, d_{2-6} = 1.57, s_{2-6} = 0.82$. For experiment two the estimates were, $d_1 = 1.56, s_1 = 0.73, d_{2-6} = 1.58, s_{2-6} = 0.77$. It is apparent that, in comparison to our experiment, there is a greatly reduced difference between position one and the remaining positions. To confirm this impression, we also fit the SDT-ISL model to these data.

The SDT-ISL model fit the confidence data well for both experiments; for experiment one, $\chi^2(143) = 147.5$, $p = 0.38$, and for experiment two, $\chi^2(143) = 122.2$, $p = 0.90$. Figures 3b and 3c plot the best fitting parameter estimates as a function of position for experiment one and experiment two, respectively. To examine any effect of position on underlying discriminability, we undertook the same analyses we conducted on our data. We first fit an equal discriminability sub-model in which d_i and s_i were constrained to be constant across all positions, i . The fit of this model was not significantly different from the fit of the unconstrained model; for experiment one, $\chi^2(10) = 9.32$, $p = 0.50$, and for experiment 2, $\chi^2(10) = 5.80$, $p = 0.83$. Similar results were obtained when we fit the sub-model in which $d_{a,i}$ was constrained to be constant across position. The relevant likelihood ratio tests were not significant; for experiment one, $\chi^2(5) = 3.67$, $p = 0.60$, and for experiment two, $\chi^2(5) = 1.07$, $p = 0.96$. These results confirm that there are no significant differences in underlying discriminability for either experiment one or experiment two. It obviously follows that there are also no differences between position 1 and later positions.

Within-Participant Analysis

A striking observation in our experiment as well as in the two experiments conducted by Wilson et al. is that the post-target false alarm rate is considerably less than the pre-target false alarm rate. There is clearly something about the presentation of a target item that leads to a change in responding. To account for this effect, two possibilities present themselves. The first takes a “within-participant” perspective, hypothesizing that the difference is due to a change in the behaviour of participants across lineup position. On this view, prior to presentation of the target, all participants adopt the same (or broadly similar) response criteria, consistent with the assumptions of the ISL and SDT-ISL models. However, after a

target is presented, these criteria increase and hence the post-target false alarm rate decreases. The second possibility takes a “between-participant” perspective, hypothesizing that the difference between pre-target and post-target false alarm rates is due to a selection bias. On this view, different participants adopt different response criteria which they hold constant throughout the lineup. The more conservative the criteria, the fewer identifications are made at each lineup position. This also holds true when the target is presented. As a result, conservative participants will tend to form a larger proportion of the group whose lineup has not been terminated by a successful target identification. Consequently, the overall post-target false alarm rate is necessarily lower post-target than pre-target.

To distinguish between the within- and between-participant hypotheses, we took advantage of the design used by Wilson et al. which required participants to respond to every item in succession. This allowed us to examine responses for all participants throughout the lineup sequence. If the within-participant hypothesis is correct then post-target false alarm rates should be less than pre-target false alarm rates, just as was found in the previous analyses where a stopping rule was imposed. On the other hand, if the between-participant hypothesis is correct, there should be no change in false alarm rates from pre-target to post-target because each group of participants stays the same. In order to undertake this test, we developed and fit a “within-participant” version of the ISL model. Using the terminology introduced in Equation (1), this model is defined as

$$P_{ID}(k)_{ij} = 1 - Q_{ijk},$$

$$Q_{ijk} = \begin{cases} F_{ik}, & i < j \\ G_{ik}, & i = j \\ F_{ik}^*, & i > j, \end{cases} \quad (8)$$

where $P_{ID}(k)_{ij}$ is the conditional probability of identifying item i at confidence level k given that item j is the target. This model treats the sequential lineup as effectively equivalent to a sequence of n showups.

We fit the within-participant ISL model to data from both experiments conducted by Wilson et al. The results are shown in Figure 5, plotted in the same format as Figure 1. The first thing to note is that there is a difference between pre-target and post-target false alarm rates at each target position. Bearing in mind that each point in each panel of Figure 5 is based on the same set of participants⁸, the corresponding likelihood ratio tests are highly significant – for experiment one, $\chi^2(5) = 581.4$, $p < 0.001$, and for experiment two, $\chi^2(5) = 154.4$, $p < 0.001$. This is consistent with the within-participant hypothesis, indicating that participants adjusted their response criteria to become more conservative following the target. In fact, this adjustment occurred whether or not a participant actually identified the target. For participants who correctly identified the target, the pre-target false alarm rate was 0.14 and the post-target false alarm rate was 0.04. For participants who failed to identify the target, the pre-target false alarm rate was 0.24 and the post-target false alarm rate was 0.09. Interestingly, contra the between-participant hypothesis, participants who failed to identify the target made *more* false alarms (both pre- and post-target) suggesting that their response criteria were less rather than more conservative.

A second feature of the within-participant analysis is that we observed significant changes in both pre-target and post-target false alarm rates across target position. In both experiments, pre-target false alarm rates tended to decrease – for experiment one, $\chi^2(5) =$

⁸ Thereby violating independence.

71.0, $p < 0.001$, and for experiment two, $\chi^2(5) = 13.7$, $p = 0.018$. In contrast, post-target false alarm rates tended to increase – for experiment one, $\chi^2(4) = 23.5$, $p < 0.001$, and for experiment two, $\chi^2(4) = 12.3$, $p = 0.013$. Of perhaps greater interest is the observation that hit rates declined substantially across target position. For experiment one, $\chi^2(5) = 63.8$, $p < 0.001$, and for experiment two, $\chi^2(5) = 23.5$, $p < 0.001$. To determine if this is attributable to a decrease in underlying discriminability, we fit the within-participant version of the SDT-ISL model to these data. The results are shown in Figure 6 with discriminability plotted in terms of $d_{a,i}$. In both experiments, $d_{a,i}$ appears to decline across position. This effect is significant for experiment one – the likelihood ratio test of the sub-model in which d_i and s_i are held constant is significant, $\chi^2(10) = 29.5$, $p = 0.001$, as was the test of the sub-model in which $d_{a,i}$ is held constant, $\chi^2(5) = 22.4$, $p < 0.001$. Neither of these tests is significant in experiment two.

General Discussion

The aim of this study was to examine the effects of target position on underlying discriminability and response bias in a sequential lineup with a stopping rule. To achieve this, we developed and applied a model of the sequential lineup task that accounted for the joint effects of target and identification position on identification rates. In its more general form, which we termed the ISL model, it enabled us to estimate hit and false alarm rates at each lineup position. We fit the ISL model to data from our experiment which had an explicit stopping rule and compared this to a similar analysis of data from the two experiments conducted by Wilson et al. (2019), in which participants were permitted to make multiple identifications. While our results broadly replicate those found in the earlier study, there were also some differences. We found a small increase in underlying discriminability from

identification position 1 to position 2 in our experiment, but underlying discriminability did not increase beyond identification position two. For the Wilson et al. experiments, we did not find any differences in underlying discriminability across identification position. This is contrary to the results found by Wilson et al., based on fits of the SDT-SEQ model. As we discussed previously, because of the assumptions made by this model, parameter estimates derived from this model may be misleading in the current circumstances. Nevertheless, it is somewhat ironic that a key result reported by Wilson et al – an increase in underlying discriminability between target positions 1 and 2 – while it did not survive our re-analysis of their data, was nevertheless observed in our data.

In the following sections, we discuss the implications of our results, and those found by Wilson et al., for a theoretical understanding the effects of sequential lineup position on underlying discriminability and response criteria. As well as providing important information concerning the features of the sequential lineup, these results have implications for theories of recognition memory.

Underlying discriminability

In our experiment, we observed a small increase in underlying discriminability between positions 1 and 2. On the face of it, this is consistent with the DFDT. According to this theory, participants learn to identify diagnostic features that support correct target identifications as a consequence of successive presentations of lineup items. In principle, the more foils that are presented prior to the target, the more learning should take place. The fact that the improvement is apparent early in the lineup – essentially complete after the first foil – may be due to how foils are selected and hence the nature of the diagnostic features. Re-

calibration may be rapid as participants appreciate that the foils are highly similar rather than randomly similar to the target.

However, the DFDT interpretation of the change in underlying discriminability is a “within-participant” account – each participant learns something about the foils and adjusts their behaviour accordingly. If this is the case then the “within-participant” analysis of experiment one of Wilson et al (2019) should also have revealed an increase in underlying discriminability across lineup position – or at least between positions 1 and 2. In fact, this analysis revealed that underlying discriminability *decreased* across the lineup.

Why is there a “within-participant” decline in underlying discriminability? To address this question, we note that the sequential lineup can be viewed as a special case of the yes-no recognition memory task. In this task, participants study a list of items (usually words) and are subsequently tested on a set of items consisting of both old (i.e., studied) and new (unstudied) items. In the sequential lineup, there is only one old item (the target) and $n - 1$ new items (the foils) which are presented in a random order to participants who must decide if any is the target. Similarly, participants in a yes-no recognition task are presented with a random sequence of targets and foils and instructed to identify each as either old or new. Apart from the presence of multiple targets, there is no stopping rule. Research has shown that in old-new recognition, the ability to discriminate between old and new items declines across test trials. This has been called the *test position effect* (Osth, Jansson, Dennis & Heathcote, 2018) or *output interference* (Criss, Malmberg & Shiffrin, 2011).

One explanation for the test position effect is that it is due to differences in the retention interval which is less for targets presented in early test positions compared to later positions. However, by using appropriate controls, Averell, Prince and Heathcote (2016)

showed that this could not account for the test position effect. At the present time, there are two alternative explanations for this effect. Both presuppose that a joint representation of the target and the *study context* is stored in memory. For each test item, memory is probed by the joint representation of that item and the current *test context*. The greater the similarity between the probe and target memory, the greater the discriminability. The first account of the test position effect was proposed by Criss et al (2011) who suggested that items presented at test are encoded into memory and associated with the retrieved study context. This means that for if a target is subsequently presented, there are additional items (foils) competing with it in the retrieval process. This impairs discriminability by adding *item-noise* to the decision. The second account was proposed by Osth and Dennis (2015) who suggested that as more items are presented at test, the test context itself becomes increasingly dissimilar to the study context, leading to increasingly poor retrieval of targets. This is called the *context-drift* account. Recent attempts to distinguish these two accounts have found evidence that both may have a part to play in different circumstances (Fox, Dennis & Osth, 2020; Osth et al., 2018)

The observation of target position effects in the sequential lineup reveals that it occurs when there is only one target and as few as six test items. While context-drift may play a role, we speculate that item-noise may be more important. This is because it is difficult to see how the relatively few items presented in the lineup could have a major impact on the test context. On the other hand, even a few items may have a strong effect on item-noise given a study list of only one item⁹. The application of memory models such as REM (Shiffrin & Steyvers,

⁹ The authors are grateful to Adam Osth (private communication) for this observation.

1997) or that proposed by Osth and Dennis (2015) to these data may help to resolve this question.

Target position effects may also account for the observed detrimental effect of multiple laps of the lineup on identification performance (Stebly, Dietrich, Ryan, Raczynski, & James, 2011; Horry, Brewer, Weber & Palmer, 2015; Seale-Carlisle, Wetmore, Flowe & Mickes, 2019). In frequent applications of the sequential lineup, the rule that the lineup is terminated on identification is waived and participants may be permitted or even required to view all the lineup items prior to making a decision – completing a so-called *lap* of the lineup. For example, in the UK witnesses are required to view each member of a 9-item lineup twice before making a decision (Seale-Carlisle et al., 2019). In their study, Seale-Carlisle et al. found that underlying discriminability was less when this procedure was used compared to an equivalent simultaneous lineup. They also found a numerical decrease in underlying discriminability in a sequential lineup consisting of one mandated lap, two mandated laps, and two optional laps, respectively. Similar results were reported by Stebly et al. (2011) and Horry et al. (2015) who also found an increase in false alarm rates from lap 1 to lap 2.

In contrast to the detrimental effect of multiple laps, the application of a stopping rule does not appear to impair performance in a sequential lineup. In a recent review of the literature and meta-analysis, Kaesler et al. (2020) found little or no difference between simultaneous and sequential presentation of lineup items, as long as the stopping rule was in place. This is consistent with the present results. When a stopping rule is applied, we find no decline in underlying discriminability with increasing target position – in fact, in our experiment we observed an *increase* in discriminability from position 1 to position 2. The conclusion we draw is that within-participant target position effects are ameliorated at the

group level by the stopping rule which we attribute to a selection effect. We assume that a participant's response criterion is sensitive to their underlying discriminability (we discuss this assumption in the next section). This means that participants with relatively low discriminability will tend to have lower hit rates and higher false alarm rates. Accordingly, they are more likely to false alarm at early positions in the lineup thereby removing themselves from the cohort of responders. This cohort will therefore increasingly consist of participants with relatively strong memories, offsetting an accompanying within-participant decline corresponding to the target position effect.

Response criteria

Previous research has found that false alarms tend to increase over the course of the sequential lineup (Carlson et al., 2016; Carlson et al., 2008; Clark & Davey, 2005; Horry et al., 2012; Lindsay & Wells, 1985). This is consistent with a progressively more lenient response criterion. However, our analyses reveal that this forms only part of the story. A critical distinction needs to be made between pre-target and post-target false alarms. In both our between-participant and within-participant analyses, we found that pre-target false alarm rates remained relatively constant across lineup position and, if anything, tended to decrease, at least from position 1 to position 2 (see Figures 1 and 5). In contrast, post-target false alarm rates, although relatively low compared to pre-target rates, increased significantly across lineup position. These results indicate that prior to experiencing the target, participants become progressively *less* lenient in their responding. In contrast, once they experience the target, there is a shift to an even more conservative criterion which, in turn, becomes increasingly *more* lenient across subsequent lineup positions. As noted earlier, the difference between pre-target and post-target false alarm rates is a “within-participant” effect. The

criterion shift is also greater earlier in the lineup (e.g., after viewing the target in position 1) than later in the lineup (e.g., after viewing the target in position 5).

That presentation of the target should have such a pronounced effect on false alarm rates was unexpected. While previous research indicates that people can readily shift their response criteria between different sets of items in yes-no recognition tasks, item-to-item shifts have been more difficult to observe (for a review, see Hockley, 2011). For example, in a now classic study, Stretch and Wixted (1998) gave participants a study list consisting of *weak* words presented once and *strong* words presented five times. At test, strong and weak words were presented in different colors. Non-studied lures were similarly presented in the same colors. By taking advantage of the color cue, a participant could optimize their performance by shifting their response criterion to be more lenient for weak items and less lenient for strong items. This would result in a “mirror effect” – higher hit rates and lower false alarm rates for strong items compared to weak items. However, the results showed no difference in the false alarm rates, suggesting that participants did not (or could not) shift their response criteria. Subsequently, Starns and Olchowski (2015) showed that participants could be induced to make the requisite criterion shift if forced by task demands to respond to the strength cue.

In contrast to research suggesting that people are largely unwilling to make item-to-item shifts in response criteria, our results reveal a substantial shift from immediately before to immediately after viewing the target. While it is clear that a shift has occurred, the reason for it is less obvious. In developing such an account, we first note that the observed shift is unlikely to have anything to do with the target per se¹⁰. Participants do not know that it is

¹⁰ We thank an anonymous reviewer for this insight.

target that is presented and the “between-participant” analyses show that the shift occurs even when participants fail to identify it correctly as the target. More likely, participants are responding to an item with relatively high memory strength, even if it is not high enough to trigger a response. While other items may also have high memory strength, and participants may respond in the same way, it is impossible to identify these items for any one participant – only the target can be singled out in this way.

Why does an item with relatively high memory strength induce a criterion shift? If participants rely on absolute judgments (as per Lindsay & Wells, 1985) they would maintain a fixed response criterion applicable to all items and there would be no requirement for it to change from item to item. The fact that a participant experiences an item with relatively high memory strength should make no difference. If it exceeds the criterion, it is identified, otherwise not. On the other hand, if participants rely on relative judgments then response criteria would necessarily change. For example, a simple relative judgment would be to compare the memory strength of the current item to that of the previous item. If this difference exceeds a specified value then the item is identified. Formally, let x_i , be the memory strength of item i , $i > 1$. The rule would be if $x_i - x_{i-1} > c$ for some constant c then identify item i as the target. We can also write this rule as $x_i > c_i$, $c_i = c + x_{i-1}$. It follows from this that the current criterion c_i depends on the memory strength of the previous item. If this strength is relatively high then so is c_i .

Why should participants rely on relative judgments? One possibility is that they are learning about the distribution of the memory strength of foils. Following suggestions by Treisman and Williams (1984), participants may at the beginning of the lineup apply a response criterion based on their prior experience and expectations concerning the relative

similarity of the target and foils in the lineup. We observed a decrease in (pre-target) false alarm rates from position 1 to position 2. This suggests that each participant had set a value for their initial criterion that was, on average, too lenient. Based on the relative judgment strategy, this is consistent with participants having an expectation that the first lineup item would be more similar to the target than it turned out to be. Once the initial adjustment has been made, each participant thereafter adjusts their criterion in response to the memory strengths of the items they subsequently experience. Again following Treisman and Williams (1984), participants may be trying to maintain a running estimate of the average memory strength of the foils. If so, they would be employing a decision strategy broadly similar to the *ensemble model* proposed by Wixted, Vul, Mickes and Wilson (2018). In this model, the memory strength of the current item is compared against the *average* of the memory strengths of the previously presented items.

The ensemble model is consistent with three features of our results. First, because the average memory strength of foils is zero (by definition), pre-target false alarm rates should remain constant across lineup positions 2 to 6. This is the case for the “between-participant” analyses but not for the “within-participant” analysis of the data from Wilson et al. which showed a small decrease. However, this decrease is attributable to a small change from position 1 to position 2 as there on after there is little systematic change (see Figure 5). Direct comparison of false alarm rates in positions 2 and 6 reinforce this conclusion. Second, following presentation of a high memory strength item (such as the target), false alarm rates show a marked decrease. Third and most critically, the ensemble model is also consistent with the observed increase in post-target false alarms over lineup position. This arises directly from the averaging process. Let x_t be the target memory strength. Then $c_2 = c + x_t$,

$c_3 = c + \frac{1}{2} x_t$, $c_4 = c + \frac{1}{3} x_t$, and so on. Because the post-target response criterion declines, false alarm rates increase.

Our analysis of criterion shifts in the sequential lineup raises the question of why similar shifts are so difficult to observe in yes-no recognition tasks. One possibility is that there is only one study item in the sequential lineup. By learning about the distribution of foil memory strength at test, the participant is able to optimize their decision making with respect to this one target. In contrast, in a typical yes-no recognition task, there are multiple study items and potentially multiple foil distributions. In this environment, even if it was clear what to do, adjusting response criteria from item-to-item may be too demanding a task.

Like all models, the ISL model makes several simplifying assumptions. However, the fact that it provides a good fit to three sets of data suggests that these assumptions are not too far from the mark. Of these, the most critical is the independence assumption – that decisions at each identification position are independent of any previous decision. This is a strong assumption but one that, at present, cannot be rejected. However, given the advent of large-scale studies with many thousands of participants, it is likely that fine-grained deviations from the model will eventually be observed. This is also likely to be true of the SDT-ISL model which adopts the simplifying assumption that memory strength is normally distributed. It is highly probable that this assumption may prove to be only an approximation but, in the meantime, both models continue to provide a set of analytical tools to explore the effects of different factors on sequential lineup performance. Finally, the ISL model also assumes that post-target false alarm rates are independent of target position. Further research may reveal that this assumption is not strictly true and the nature of observed differences may reveal more information concerning the relative judgment strategies that may be used.

Concluding comments

Our investigations have revealed some of the complexities of the sequential lineup task. It matters a lot where the target is placed – people respond more leniently prior to the target's occurrence than afterwards. It also matters a lot whether or not a stopping rule is applied. Without a stopping rule, underlying discriminability declines across the lineup and, by extension, across subsequent re-presentations or laps. With a stopping rule, this does not happen because, we suggest, it selectively removes witnesses with relatively poor memories. Lastly, although the sequential lineup was proposed by Lindsay and Wells (1985) to encourage identifications based on absolute rather than relative judgments, our results suggest the opposite – that people base their identifications on a comparison of the current item to those they have seen before.

Context

This paper was written in large part as a response to the paper by Wilson, Donnelly, Christenfeld and Wixted (2019) which aimed to clarify the effect of target position on eyewitness identification using a sequential lineup. Because the sequential lineup is very frequently used across multiple countries, if there are position effects, it is important that we know about them. Our previous work has focused on the construction of mathematical models of the lineup identification task based on a signal detection framework. We believe it is only through the development and rigorous testing of such models that it will be possible to gain a clear understanding of the factors affecting eyewitness identification. The model we present clarifies a number of effects found in our data and the data reported by Wilson et al. Importantly, we find that when there is no stopping rule and witnesses are permitted or required to view all the lineup items before making a decision, memory is significantly impaired, similar to target position effects found in yes-no recognition tasks. This is a novel finding and raises doubt concerning the efficacy of the widespread use of this practice in actual lineups. We also observed a significant increase in false alarm rates following presentation of a target. Both of these effects have been shown for the first time and have important implications for current theories of recognition memory.

References

- Andersen, S. M., Carlson, C. A., Carlson, M. A., & Gronlund, S. D. (2014). Individual differences predict eyewitness identification performance. *Personality and Individual Differences, 60*, 36-40.
- Averell, L., Prince, M. & Heathcote, A. Fundamental causes of systematic and random variability in recognition memory. *Journal of Memory and Language, 88*, 51–69.
- Carlson, C. A., & Carlson, M. A. (2014). An evaluation of lineup presentation, weapon presence, and a distinctive feature using ROC analysis. *Journal of Applied Research in Memory and Cognition, 3*(2), 45-53.
- Carlson, C. A., Carlson, M. A., Weatherford, D. R., Tucker, A., & Bednarz, J. (2016). The effect of backloading instructions on eyewitness identification from simultaneous and sequential lineups. *Applied Cognitive Psychology, 30*(6), 1005-1013.
- Carlson, C. A., Gronlund, S. D., & Clark, S. E. (2008). Lineup composition, suspect position, and the sequential lineup advantage. *Journal of Experimental Psychology: Applied, 14*(2), 118-128.
- Clark, S. E., & Davey, S. L. (2005). The target-to-foils shift in simultaneous and sequential lineups. *Law and Human Behavior, 29*(2), 151-172.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, N.J: L. Erlbaum Associates.
- Criss, A. H., Malmberg, K. J. & Shiffrin, R. M. (2011). Output interference in recognition memory. *Journal of Memory and Language, 64*, 316–326.

- 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.
- Eng, H. Y., Chen, D., & Jiang, Y. (2005). Visual working memory for simple and complex visual stimuli. *Psychonomic Bulletin & Review*, 12(6), 1127-1133.
- Fox, J., Dennis, S., & Osth, A. F. (2020). Accounting for the build-up of proactive interference across lists in a list length paradigm reveals a dominance of item-noise in recognition memory. *Journal of Memory and Language*, 110, 104065.
- Gronlund, S. D., Carlson, C. A., Dailey, S. B., & Goodsell, C. A. (2009). Robustness of the sequential lineup advantage. *Journal of Experimental Psychology: Applied*, 15(2), 140-152.
- Gronlund, S. D., Carlson, C. A., Neuschatz, J. S., Goodsell, C. A., Wetmore, S. A., Wooten, A., & Graham, M. (2012). Showups versus lineups: An evaluation using ROC analysis. *Journal of Applied Research in Memory and Cognition*, 1(4), 221-228.
- Hockley, W. E. (2011). Criterion changes: How flexible are recognition decision processes? In P. A. Higham & J. P. Leboe (Eds). *Constructions of remembering and metacognition: Essays in honour of Bruce Whittlesea* (pp. 155-166). Palgrave Macmillan.
- Horry, R., Brewer, N., Weber, N., & Palmer, M. A. (2015). The effects of allowing a second sequential lineup lap on choosing and probative value. *Psychology, Public Policy, and Law*, 21(2), 121-133.

- Horry, R., Palmer, M. A., & Brewer, N. (2012). Backloading in the sequential lineup prevents within-lineup criterion shifts that undermine eyewitness identification performance. *Journal of Experimental Psychology: Applied*, 18(4), 346-360.
- Jackson, M. C., & Raymond, J. E. (2008). Familiarity enhances visual working memory for faces. *Journal of Experimental Psychology: Human Perception & Performance*, 34(3), 556-568.
- Kaesler, M., Dunn, J. C., Ransom, K., & Semmler, C. (2020). Do sequential lineups impair underlying discriminability? *Cognitive Research: Principals and Implications*, 5, 35. <https://doi.org/10.1186/s41235-020-00234-5>
- Lindsay, R., & Wells, G. L. (1985). Improving eyewitness identifications from lineups: Simultaneous versus sequential lineup presentation. *Journal of Applied Psychology*, 70(3), 556-564.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279-281.
- Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: From psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences*, 17(8), 391-400.
- Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide.*, 2nd ed. Mahwah, NJ: Lawrence Erlbaum Associates Publishers; US.
- Meisters, J., Diedenhofen, B., & Musch, J. (2018). Eyewitness identification in simultaneous and sequential lineups: An investigation of position effects using receiver operating characteristics. *Memory*, 26(9), 1297-1309.

- Memon, A., & Gabbert, F. (2003). Unravelling the effects of sequential presentation in culprit-present lineups. *Applied Cognitive Psychology*, 17(6), 703-714.
- 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.
- 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.
- Neuschatz, J. S., Wetmore, S. A., Key, K. N., Cash, D. K., Gronlund, S. D., & Goodsell, C. A. (2016). A comprehensive evaluation of showups. In M. K. Miller & B. H. Bornstein (Eds.), *Advances in Psychology and Law: Volume 1* (pp. 43-69). Cham: Springer International Publishing.
- Osth, A. F., & Dennis, S. (2015). Sources of interference in item and associative recognition memory. *Psychological Review*, 122(2), 260–311.
- Osth, A. F., Jansson, A., Dennis, S. & Heathcote, A. (2018). Modeling the dynamics of recognition memory testing with an integrated model of retrieval and decision making. *Cognitive Psychology*, 104, 106–142
- Rotello, C. M., & Chen, T. (2016). ROC curve analyses of eyewitness identification decisions: An analysis of the recent debate. *Cognitive Research: Principles and Implications*, 1(1), 1-12.

- 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.
- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM: Retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4, 145–166.
- Sporer, S. L. (1993). Eyewitness identification accuracy, confidence, and decision times in simultaneous and sequential lineups. *Journal of Applied Psychology*, 78(1), 22-33.
- Starns, J. J. & Olchowski, J. E. (2015). Shifting the criterion is not the difficult part of trial-by-trial criterion shifts in recognition memory. *Memory & Cognition*, 43, 49–59.
- Stebay, N. K., Dietrich, H. L., Ryan, S. L., Raczynski, J. L., & James, K. A. (2011). Sequential lineup laps and eyewitness accuracy. *Law and Human Behavior*, 35(4), 262-274.
- Treisman, M. & Williams, T. C. (1984). A theory of criterion setting with an application to sequential dependencies. *Psychological Review*, 91(1), 68–111.
- U.S. National Department of Justice. (1999). *Eyewitness Evidence: A Guide for Law Enforcement*. Retrieved from <https://www.ncjrs.gov/pdffiles1/nij/178240.pdf>
- Wells, G. L., Stebay, N. K., & Dysart, J. E. (2015). Double-blind photo lineups using actual eyewitnesses: An experimental test of a sequential versus simultaneous lineup procedure. *Law and Human Behavior*, 39(1), 1-14.

- Wilson, B. M., Donnelly, K., Christenfeld, N., & Wixted, J. T. (2019). Making sense of sequential lineups: An experimental and theoretical analysis of position effects. *Journal of Memory and Language*, 104, 108-125.
- Wixted, J. T., & Mickes, L. (2014). A signal-detection-based diagnostic-feature-detection model of eyewitness identification. *Psychological Review*, 121(2), 262-276.
- Wixted, J. T., & Mickes, L. (2018). Theoretical vs. empirical discriminability: The application of ROC methods to eyewitness identification. *Cognitive Research: Principles and Implications*, 3(1), 1-22.
- Wixted, J. T., Vul, E., Mickes, L. & Wilson, B. M. (2018). Models of lineup memory. *Cognitive Psychology*, 105, 81–114

Table 1. The structure of the ISL model for an $n = 6$ sequential lineup.

ID Position	Target Position						
	1	2	3	4	5	6	None
1	$(1 - G_{1k})$	$(1 - F_{1k})$	$(1 - F_{1k})$	$(1 - F_{1k})$	$(1 - F_{1k})$	$(1 - F_{1k})$	$(1 - F_{1k})$
2	$G_1(1 - F_{2k}^*)$	$F_1(1 - G_{2k})$	$F_1(1 - F_{2k})$	$F_1(1 - F_{2k})$	$F_1(1 - F_{2k})$	$F_1(1 - F_{2k})$	$F_1(1 - F_{2k})$
3	$G_1F_2^*(1 - F_{3k}^*)$	$F_1G_2(1 - F_{3k}^*)$	$F_1F_2(1 - G_{3k})$	$F_1F_2(1 - F_{3k})$	$F_1F_2(1 - F_{3k})$	$F_1F_2(1 - F_{3k})$	$F_1F_2(1 - F_{3k})$
4	$G_1F_2^*F_3^*(1 - F_{4k}^*)$	$F_1G_2F_3^*(1 - F_{4k}^*)$	$F_1F_2G_3(1 - F_{4k}^*)$	$F_1F_2F_3(1 - G_{4k})$	$F_1F_2F_3(1 - F_{4k})$	$F_1F_2F_3(1 - F_{4k})$	$F_1F_2F_3(1 - F_{4k})$
5	$G_1F_2^*F_3^*F_4^*(1 - F_{5k}^*)$	$F_1G_2F_3^*F_4^*(1 - F_{5k}^*)$	$F_1F_2G_3F_4^*(1 - F_{5k}^*)$	$F_1F_2F_3G_4(1 - F_{5k}^*)$	$F_1F_2F_3F_4(1 - G_{5k})$	$F_1F_2F_3F_4(1 - F_{5k})$	$F_1F_2F_3F_4(1 - F_{5k})$
6	$G_1F_2^*F_3^*F_4^*F_5^*(1 - F_{6k}^*)$	$F_1G_2F_3^*F_4^*F_5^*(1 - F_{6k}^*)$	$F_1F_2G_3F_4^*F_5^*(1 - F_{6k}^*)$	$F_1F_2F_3G_4F_5^*(1 - F_{6k}^*)$	$F_1F_2F_3F_4G_5(1 - F_{6k}^*)$	$F_1F_2F_3F_4F_5(1 - G_{6k})$	$F_1F_2F_3F_4F_5(1 - F_{6k})$
None	$G_1F_2^*F_3^*F_4^*F_5^*F_6^*$	$F_1G_2F_3^*F_4^*F_5^*F_6^*$	$F_1F_2G_3F_4^*F_5^*F_6^*$	$F_1F_2F_3G_4F_5^*F_6^*$	$F_1F_2F_3F_4G_5F_6^*$	$F_1F_2F_3F_4F_5G_6$	$F_1F_2F_3F_4F_5F_6$

Table 2. Observed counts for each target and identification position in experiment one.

ID Position	Target Position						
	1	2	3	4	5	6	None
1	701	119	104	122	90	106	185
2	8	626	68	75	68	62	94
3	7	10	631	72	47	91	99
4	11	10	12	532	63	63	86
5	17	12	6	6	485	53	87
6	15	12	11	9	10	421	79
None	224	185	175	154	148	143	790

Table 3. Observed counts for each target and identification position in Wilson et al (2019, experiment one).

ID Position	Target Position						
	1	2	3	4	5	6	None
1	444	103	103	104	107	98	602
2	2	340	77	61	73	69	414
3	6	7	291	77	83	76	439
4	12	13	9	246	45	48	375
5	16	5	3	2	175	41	239
6	7	4	5	6	3	152	170
None	52	57	40	33	51	50	947

Table 4. Observed counts for each target and identification position in Wilson et al (2019, lineup condition, experiment two).

ID Position	Target Position						
	1	2	3	4	5	6	None
1	135	34	24	19	25	27	164
2	0	110	21	21	18	24	125
3	3	0	100	31	26	13	115
4	2	2	0	85	22	14	113
5	5	2	2	0	54	12	66
6	2	2	4	2	2	62	62
None	23	17	12	12	12	15	320

Figure Captions

Figure 1. Best-fitting parameters of the ISL model as a function of identification position. (a) Experiment one; (b) Wilson et al. (2019, Experiment 1); (c) Wilson et al. (2019, Experiment 2).

Figure 2. ROC analysis of data from experiment one. (a) ROC plot of hit rates $(1 - G_{ik})$ plotted against pre-target false alarm rates $(1 - F_{ik})$ for each position and confidence level. The dashed line corresponds to the best-fitting normal curve for position one (corresponding points labelled). The solid line corresponds to the best-fitting normal curve for positions 2-6. The legend indicates the confidence level from most lenient ($k = 1$) to most conservative ($k = 5$). (b) zROC plot of inverse normal hit rates plotted against inverse normal false alarm rates. The dashed line corresponds to the best-fitting straight line for position one. The solid line corresponds to the best-fitting straight line for positions 2-6.

Figure 3. Best-fitting parameters of the SDT-ISL model as a function of identification position. (a) Experiment one; (b) Wilson et al. (2019, Experiment 1); (c) Wilson et al. (2019, Experiment 2).

Figure 4. zROC plots from fitting the ISL model to data from Wilson et al. (2019). The dashed line corresponds to the best-fitting straight line for position one. The solid line corresponds to the best-fitting straight line for positions 2-6. The legend indicates the confidence level from most lenient ($k = 1$) to most conservative ($k = 5$). (a) Experiment one. (b) Experiment two.

Figure 5. Best-fitting parameters of the within-participant version of the ISL model as a function of identification position. (a) Wilson et al. (2019, Experiment 1); (b) Wilson et al. (2019, Experiment 2).

Figure 6. Best-fitting parameters of the within-participant version of the SDT-ISL model as a function of identification position. (a) Wilson et al. (2019, Experiment 1); (b) Wilson et al. (2019, Experiment 2).

Figure 1

Figure 2

Figure 3

Figure 4

Figure 5

Figure 6