Violations of Conditional Independence in Mixed-State Models of Visual Working
Memory

Samuel Winiger

University of Zurich

Henrik Singmann

University College London

David Kellen

Syracuse University

Abstract

Mixed-state models of visual working memory assume discrete mental states of knowing and guessing; information is either represented in memory or not at all. At the core of most conceptualizations of mixed-state models is the notion of conditional independence, that is, the probability of a process taking place when a given memory state is reached is independent of the probability of reaching said state. For example, the probability of committing a memory-based error being conditionally independent of the probability of a memory-based response being made. We test this assumption using an extended recognition-memory paradigm that allows for second choices, along with a task-difficulty manipulation. The second choices obtained through this paradigm provide testable predictions for conditional independence that are violated in the reported data. Overall, our results indicate that the probabilities of memory-based responses and of memory-based errors go hand in hand.

Word Count (Introduction, Results, and Discussion): Approx. 2850 words

Keywords: Visual Working Memory, Mixed-State Models, Conditional Independence

Introduction

The study of visual working memory is concerned with people’s ability to briefly retain visual representations of stimuli that were previously encountered. In recent years, a number of experimental findings have been used to argue for the presence of guessing (i.e., responses devoid of any memory information) in both delayed reproduction tasks (e.g., Adam, Vogel, & Awh, 2017), as well as change detection tasks (e.g., Nosofsky & Donkin, 2016). For example, Nosofsky and Donkin (2016) found that people would often judge test stimuli as being the same represented in working memory, despite a large difference between the previously encountered stimuli and its test counterpart. In other trials, the same participants were found to be able detect small differences between studied and tested stimuli. Based on these results, Nosofsky and Donkin argued that *mixed-state models* provide a natural account of this data, as they postulate the possibility that sometimes people are merely guessing. This view is reflected in the numerous computational models of visual working memory that conceptualize responses as either being memory-based or guessing-based (e.g., van den Berg, Awh, & Ma, 2014; Hardman, Vergauwe, & Ricker, 2017; Rouder et al., 2008).

At the core of most of these mixed-state models is the notion of *conditional independence* (Province & Rouder, 2012), according to which the events taking place in the context of a specific mental state are independent of the probability of entering that state. So far, this assumption has most frequently been investigated using confidence rating responses. In the case of confidence ratings, conditional independence implies that the *state-response mapping* – the mapping of mental states onto observable confidence ratings (Klauer & Kellen, 2010; Moran, 2016; Rouder & Morey, 2009) – is independent of the probability of entering a state. For example, manipulating the probability of being in a memory state (e.g., by changing the difficulty of a task) should not change the distribution of memory-based confidence ratings, but only alter the mixture of memory-based versus guessing-based responses (e.g., Ricker, Thiele, Swagman, & Rouder, 2016).

More recent mixed-state models of visual working memory relax the assumption of

perfect accuracy or precision in memory-based responses (e.g., Nosofsky & Gold, 2018, Winiger, Singmann, & Kellen, 2021) that was found in earlier variants (e.g., Rouder et al., 2008, Donkin, Nosofsky, Gold, & Shiffrin, 2013). The possibility of an inaccurate memory representation introduces the possibility of *memory-based errors*. For example, the color red, when presented during study, could be wrongly encoded or stored as orange. At test, the resulting mismatch would lead to a rejection of the same red color (i.e., it is judged as different from the presented color). In this context, conditional independence implies that the probability of memory-based errors occurring is independent of the probability of memory-based responses.

However, one could also conceptualize mixed-state models without this notion of conditional independence (e.g., Keshvari, van den Berg, & Ma, 2013; Devkar, Wright, & Ma, 2017; van den Berg et al., 2014). Models that take this alternative route assume that responses from memory states are not independent from the probability of entering the respective states. According to these models, the manipulation of task difficulty could affect the probability of giving a memory-based response, as well as the quality of the memory-based response, for example the distribution of confidence ratings, or the proportion of memory-based errors in binary decisions.

The two conceptualizations of visual working memory discussed above are illustrated in Figure 1, in the context of a visual working memory recognition task with binary responses (“old” vs. “new”). Memory representation of colors are assumed to be fallible, an idea that is captured by the Gaussian distributions around the presented/tested colors on a color scale. Placed on this color scale are two thresholds: If the representation of a color is within these thresholds, then it is deemed to be “similar enough” to the test color, which leads to a (correct) recognition. If the representation of a color is outside of the thresholds, the mismatch between it and the test color is too great and the test color is (erroneously) rejected – a memory-based error. If a color was not encoded in the first place, then no mnemonic information is available and the decision maker is forced to guess when requested to produce a response.

The goal of the present manuscript is to report a direct comparison between these

Figure 1. Graphical representation of the two conceptualizations of mixed-state models. m = probability of a memory-based response (i.e., a color being encoded). g = probability of guessing “old” from a guessing state. The actual color is represented as a point on a color scale. Memory representations are represented by distributions around the actual color. If a color representation is within the bounds of the two thresholds, it is recognized as the same color (i.e., similar enough), leading to a “old” response. If a color representation is outside the thresholds, it is deemed different from the actual color (i.e., dissimilar), leading to a “new” response. Left: Under the assumption of conditional dependent memory errors, stronger encoding is associated with a larger probability of the color being encoded (parameter m), as well as different distributions for strong and weak colors. Here, the distribution for strong colors is tighter around the actual color than that for the weak color, leading to fewer incorrect rejections (i.e., less area outside the similarity boundaries). Right: Under the assumption of conditional independent memory errors all encoded colors are equally strong represented in memory. Here, stronger encoding means larger probability of the color being encoded only (parameter m).

two conceptualizations of visual working memory, by directly testing the assumption of conditional independence (see Province & Rouder, 2012; Ricker et al., 2016). More specifically, we will show that the first conceptualization of mixed-state models discussed above, which assumes conditional independence, imposes testable constraints on *second-choice probabilities*, and to report data that show a clear violation of these constraints. In contrast, the second conceptualization discussed, which eschews conditional independence, does not impose such constraints and is able successfully account for the data at large.

Encoding-strength manipulations, such as varying the number of times each color

is presented, allow us to establish different classes of stimuli — *weak* and *strong* colors. Under the assumption of *conditionally-independent memory errors*, a strength manipulation can only affect the probability of a color being encoded into working memory (this probability is denoted by parameter m in Figure 1, right panel) — the state-response mapping remains unaffected. In contrast, the assumption of *conditionally-dependent memory errors* allows a successful strength manipulation to also decrease the probability of memory-based errors (e.g., by leading to a ‘tighter’ distribution of color representations; see Figure 1, left panel).

Testing Conditional Independence With Second Choices

Binary old-new judgments alone cannot distinguish between these two different accounts. In order to do so, we will need to rely on an extended paradigm in which *second choices* are elicited (see Kellen & Klauer, 2011; Kellen, Winiger, Dunn, & Singmann, in press). In the present work, we elicited these choices as follows (see Figure 2): In the proposed paradigm, each trial begins with a study phase in which participants see eight colored circles comprised of six different unique colors. Four of the colors — the *weak colors* — appear only once and two of the colors — the *strong colors* — appear twice.

The test phase begins with a subsetting task in which four different colors are shown. These four colors are comprised of a random combination of (old) colors that were shown in the study phase and (new) colors that were not previously shown. The participant is requested to select all the colors they believe to be old. These recognition judgments establish two subsets of colors — those that were recognized (*subset-old*) and those that were not (*subset-new*). Figure 2 provides an illustration of a subsetting trial, in which only one of the colors was selected as old.

In a proportion of the trials (roughly 50%), the initial selection made in the subsetting task was followed up by a request to provide a second choice. Participants are presented with two colors coming from one of the subsets and asked to choose for which of them they are most willing to change their initial judgment. Figure 2 provides an example: The participant is shown two of the colors that they did not select as old, and

Figure 2. Abbreviated illustration of a trial in the subsetting paradigm. Colors were presented in sequence, 200 ms each, in clockwise order. Black placeholder circles were visible throughout the learning phase. Four colors were presented once (weak) and two colors were presented twice (strong). In the test phase, participants were asked to select any number of colors they thought were old. In the follow-up trial (depicted here: subset-new with a weak and a strong color), a label would indicate what the task at hand was: e.g., “Follow-up: These are two colors you judged to be **NEW**. Select the one color for which you are most willing to change your answer from **NEW** to **OLD**.”).

requested to indicate which one they are most willing to recognize.

To help us understand how the assumption of conditionally independent memory-based errors imposes testable constraints on second choices, we first need to establish the four dimensions along which follow-up trials can vary (the first three being directly observable, the fourth not):

1. *Subset*: the items come from subset-new or subset-old.
2. *Study Status*: The colors taken from a subset can be i) both old, ii) both new, or iii) one old and one new.
3. *Study Strength*: Any old colors can either be strong or weak colors.
4. *Mental State*: The mental state in which each of the colors is in — memory or guessing.

Now, we need to establish what responses mixed-state models predict in a follow-up trial. These predictions are determined by the mental states the colors are in: If the two colors are in the same mental state (either memory or guessing), then each of them can be selected with probability $\frac{1}{2}$. If they are in different states, then the guessed

color will be selected (for discussions, see Erdfelder, Küpper-Tetzel, & Mattern, 2011; Kellen et al., in press; Luce, 1963).

The assumption of conditional independence imposes non-trivial constraints on the expected status of unrecognized colors (i.e., old colors in subset-new). Specifically, the relative proportion of memory-based errors is expected to increase along with the probability m of memory-based responses. In order better understand this prediction, consider a case in which m is very large: In this case, very few old colors will be unrecognized. But among those that are, a considerable portion of them (if not most) will consist memory-based errors, not guesses. In contrast, when m is relatively low, we can expect that most unrecognized colors were guessed. This relationship between memory-based responses and memory-based errors reflected in the choices probabilities that are expected across different follow-up trials:

1. The probability of selecting an unrecognized strong color (S) over an unrecognized weak color (W) is smaller than $\frac{1}{2}$; i.e., $P(S | W, S) \leq \frac{1}{2}$.
2. The probability of selecting an unrecognized strong color (S) over an unrecognized new color (N) is smaller than the probability of selecting an unrecognized weak color (W) over an unrecognized new color (N); i.e., $P(S | S, N) \leq P(W | W, N)$.

In order to better illustrate these predictions, we developed a multinomial processing tree (MPT; Batchelder & Riefer, 1999) implementation of the models shown in Figure 1. Without loss of generality, this MPT model describes choice probabilities in the subsetting trials as a function of the following probability parameters:

- m_S and m_W – the probabilities of strong and weak colors being encoded
- e_S and e_W – the probabilities of memory-based errors for strong and weak colors
- g – the probability of guessing “old”, when in a guessing state.

The assumption of conditional independence corresponds to the equality constraint $e = e_W = e_S$; i.e., the probability of memory-based errors is the same for weak and strong colors. Alternatively, conditional dependence is captured by the

Figure 3. Predicted responses for the subset-new follow-up questions for varying encoding probabilities m and memory-based error rates e . Top Row: Conditionally independent memory-based errors assumes the percentage of memory-based errors to be equal for weak and strong colors ($e = e_w = e_s$). Bottom Row: Conditionally dependent memory-based errors allow for unequal memory error rates, with lower error rates for strong colors ($e_w \geq e_s$).

inequality constraint $e_w \geq e_s$; i.e., memory-based errors are more probable for weak colors. Figure 3 shows the prediction of these two accounts for the follow-up trials to the subset-new across a selection of parameter values:¹ When conditional independence is assumed, we see that the above-stated predictions hold, with $P(S | W, S) \leq \frac{1}{2}$ and $P(S | S, N) \leq P(W | W, N)$. In contrast, when we assume that memory-based errors are

¹ Scripts implementing the MPT model and generating the predictions reported here can be found in the online supplementary materials.

conditionally dependent, no such constraints are imposed over the choice probabilities.

To sum up: the assumption of conditional independence implies specific constraints at the level of choice probabilities in follow-up trials. These constraints effectively delineate a “forbidden region” for choice probabilities. It follows that conditional independence can be directly tested by investigating whether or not the probability estimates obtained from people’s choices fall inside this region (see Dunn & Anderson, 2018).

Experiment

We conducted an extended color recognition experiment to test whether, in the context of a mixed-state model of visual working memory, the rate of memory-based errors varies with task difficulty. Task difficulty was manipulated by varying the number color presentations. The data as well the script implementing its analysis are available on the Open Science Framework:

Methods

Participants. Participants were recruited through **Prolific** (www.prolific.co), and received a fixed £1.3 reward in exchange for their participation. The experiment took on average 12 minutes to complete. Participants could report to not having taken the study seriously (labeled: “better throw my data away”) without any repercussions on their pay or prolific score. Data collection was stopped once 1500 eligible participants were recruited.

Stimuli and Procedure. A set of twelve dissimilar colors was used (light grey, dark grey, orange, yellow, pink, red, dark-brown-red, cyan, blue, lime, dark green, and violet). The full set of colors was presented to the participants during the initial instruction phase. Stimuli were presented at eight fixed positions — equally spaced circles (variable fill, black outline) on an imaginary circle around the middle of the screen (see Figure 2) on a white background.

After reading the instructions, participants completed 5 example trials, followed by 40 experimental trials, roughly 50% of which were followed by a follow-up question. Each trial started with a fixation cross, presented for 1000 ms. After an additional 500 ms blank-screen presentation, the learning phase started: Six random colors were selected, with four colors being presented once (*weak colors*) and two of them twice (*strong colors*). Each of the eight positions showed one random color for 200ms with no inter-stimulus interval.² Colors were always shown in clockwise order starting with the 12 o' clock position. During the whole of the study phase, the location circles were presented.

After the last color was presented and an additional 1000 ms blank retention interval, the test phase started. The test phase started with the subsetting task; Four colors were presented next to each other with a little checkbox beneath each color. This array of colors could consist of any ratio of old to new colors. Old colors were randomly sampled from the six colors shown in the learning phase. Participants were asked to select any number of colors they judged to be old before proceeding ("Select any colors you think are **OLD**"). After the subsetting task, there was roughly a 50% chance for a follow-up trial; two colors were shown again next to each other with check boxes beneath them. Participants had to select the one color for which they were more willing to change their judgment from the subsetting task. Invariant to whether or not a follow-up trial occurred, there was a 3000 ms blank after each trial before the next one started with the fixation cross.

At the end of the study, participants filled out a short demographic questionnaire. This questionnaire included the question on whether participants took the study seriously mentioned above. An additional question tested whether participants paid attention to the kind of follow-up trial they encountered (new-to-old or old-to-new). Participants who reported to not having taken the study seriously were excluded from all analysis and participants who reported to not having done the the follow-up

² The 200 ms presentation time along with no inter-stimulus interval was selected in order to guarantee enough time for encoding while simultaneously discouraging any kind of further 'elaborative' processing of the stimuli.

questions correctly were excluded from analyses of follow-up questions.

Results

In the subsetting trials, observed average recognition rates for all 1500 participants were 33.5%, 50.6%, and 64.3%, for new, weak, and strong colors, respectively. These results indicate a successful strength manipulation overall. Results and analyses presented on the follow-up tasks include only data from participants who showed an effect of strength manipulation in the subsetting task and reported to heaving paid attention to the labels in the follow-up tasks (select “old” or select “new”). This left us with data from 778 participants, who engaged on average in 19.34 follow-up trials from which 6.55 were old-to-new and 12.79 were new-to-old trials. For an adequate capture of sampling uncertainty, all data were analysed with a binomial Bayesian generalized linear mixed model³. See Figure 4 for an overview of all results.

Model estimates of the two informative conditions indicate that the restrictions imposed by conditionally independent memory-based errors do not hold in our data: In the direct comparison of strong vs. weak colors, $P(S | S, W)$, the 95% credibility interval of the group-level posterior mean excludes the critical value of 50% (CI: 50.9% to 55.3%). The 95% credibility interval of the difference between strong vs. new colors and weak vs. new colors, $P(S | S, N) - P(W | W, N)$, excludes the critical value of 0% (CI: 0.1% to 7.1%).

We further analyzed the the data coming from these conditions, aggregated over participants, with a Bayesian joint-binomial model (e.g., Lee & Wagenmakers, 2013). The probability parameters postulated by this model adhered to the restrictions imposed by conditionally independent memory-based errors. The discrepancy between the data and the model was assessed by means of a posterior-predictive test (see Gelman & Shalizi, 2013; Klauer, 2010). Small posterior-predictive p -values indicate a

³ Each model included a fixed effect for condition (either study strength condition or comparison condition), by-participant random intercepts, and by-participant random slopes for condition, as well as the correlation among random-effects parameters. Thus, each model implemented the maximal random-effects structure justified by the design and adequately accounted for the repeated-measures nature of the data. Models were estimated with the `stan_glmer` function from the R package `rstanarm` (Goodrich, Gabry, Ali, & Brilleman; 2020).

Figure 4. Subsetting: Faint grey dots represent individual participant means. All panels: Results show estimates from a binomial Bayesian generalized linear mixed model. Black dots represent the group-level posterior means, and the black lines represent the 95% credibility intervals. Grey shaded are the “forbidden regions” for the notion of conditionally independent memory-based errors.

poor fit of the model to the data. We obtained a Bayesian p -value smaller than .001.

Discussion

The present work tested the assumption of conditional independence in mixed-state models of visual working memory. According to conditional independence, the probability of certain events (e.g., an error) occurring within a given mental state

are independent from the probability of that state being reached (Kellen & Klauer, 2018; Kellen et al., in press; Province & Rouder, 2012; Ricker et al., 2016). In the context of a recognition paradigm, conditional independence constraints the probabilities of memory-based errors as well as guessing behaviour. We extended a binary “old” vs. “new” recognition paradigm to allow for second choices for which the assumption of conditionally independent memory-based errors constraints the range of predictions in testable ways. We observed violations of these constraints, violations that serve as direct evidence against the assumption of conditional independence in mixed-state models of visual working memory.

When building models of visual working memory, a theorist is confronted with a plethora of modelling choices. In mixed-state modelling, one of the central modelling choices that helps separating the models from continuous-resource models (e.g., Bays, Catalano, & Husain, 2009; Wilken & Ma, 2004) is the notion of fixed memory precision in memory states. In the language of ‘slots modelling’ (e.g., Cowan, 2001), this means that each memory slot has fixed precision, which implies conditional independence at the level of memory-based errors. For example, the Slot-Averaging model of Zhang and Luck (2008) assumes one parameter representing the probability of a color being in memory (P_m) and one parameter representing memory precision (SD). These parameters are usually estimated independently over a set of conditions, and there is a body of literature that found evidence for this form of conditional independence, although bundled with other modelling assumptions (e.g., Zhang & Luck 2008, 2009; Xie & Zhang 2017), as well as evidence against it (e.g., Keshvari et al., 2013). Other efforts went into testing multiple modeling choices at once in an exercise of building a separate model for each possible combination – a factorial design – and finding the best model by assessing penalized model fits (van den Berg et al., 2014; Oberauer, submitted). When analyzing a broad range of delayed reproduction data, van den Berg et al. (2014) found memory precision to vary, whereas Oberauer (submitted) found memory precision to be constant across many experimental manipulations. Complementary to these two methods, we singled out the assumption of conditional independence and tested it

directly without having to rely on parameter estimation or model-comparison exercises that excessively rely on ancillary assumptions (see Kellen et al., in press). Additionally, our test is based on a novel experimental paradigm, which should be helpful in avoiding mono-operation biases of any kind (Shadish, Cook, & Campbell, 2002, Chap. 3).

Our test of conditional independence complements similar efforts by Ricker et al. (2016), and overcomes some of the limitations of their study. Ricker et al.'s participants were presented with a single dot on a circle. At test, two dots were presented on the same circle, one of which (target) was at the same location as the first presented dot and one was slightly off (lure). Participants were then asked to find the target while simultaneously providing confidence ratings. Task difficulty was manipulated by varying the distance between target and lure. The idea was again that, under the assumption of conditional independence, task difficulty should not affect the distribution of confidence ratings. Unfortunately, their testing of conditional independence was compromised, given that participants were aware all times of the task-difficulty condition that they were in (as difficulty is captured by the distance between the target and the lure). As discussed by Kellen and Klauer (2015), conditional independence is only a reasonable assumption when no 'external features' inform the participant of the difficulty condition they are presently in. This limitation is avoided in the present study, as there are no external cues associated with specific levels of the task-difficulty manipulation.

Finally, one could argue that the present results could be explained by 'slot averaging' (Zhang & Luck, 2008; for a recent implementation, see Nosofsky & Gold, 2018). In a nutshell, the idea behind slot averaging is that, if the number of colors to be remembered is smaller than the overall working memory capacity, then the available memory slots can be used to encode additional representations of certain colors. At test, the availability of the multiple representations enables a more accurate recall (by relying on the average representation) than for a color that is only represented once. In the present case, slot averaging might be used to reduce memory-based errors among strong colors, which would result in a violation of conditional independence. We find this explanation implausible, for a number of reasons: First, slot averaging makes little

sense in the present paradigm, given that the participant is not being asked to judge the color presented at a given location — participants were only asked whether or not a given color was presented, *regardless of location*. Moreover, slot averaging has been proposed as a strategy that would take place only when the number of items presented is below capacity. Given that we always presented six unique colors, this explanation would require a working memory capacity of at least seven items, which is at odds with the estimates of 3 or 4 that have been reported in the literature at large (e.g., Adam et al., 2017; Taylor, Thomson, Sutton, & Donkin, 2017; Luck & Vogel, 1997). Perhaps slot averaging was applied strategically in cases where the number of items exceeded capacity. But besides the fact that this proposal is at odds with every application of slot averaging that we are aware of, it is not clear what the proposed strategy is trying to accomplish in the present paradigm.

Open Practices Statement

The data and analysis scripts as well as the experimental code using `lab.js` (Henninger et al., 2021) (which includes all materials) are available at

The experiment and analysis were not preregistered.

References

- Adam, K. C. S., Vogel, E. K., & Awh, E. (2017). Clear evidence for item limits in visual working memory. *Cognitive Psychology*, *97*, 79–97.
<https://doi.org/10.1016/j.cogpsych.2017.07.001>
- Batchelder, W. H., & Riefer, D. M. (1999). Theoretical and empirical review of multinomial process tree modeling. *Psychonomic Bulletin & Review*, *6*(1), 57–86.
<https://doi.org/10.3758/BF03210812>
- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 7–7.
<https://doi.org/10.1167/9.10.7>
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity, (4), 100.
- Devkar, D., Wright, A. A., & Ma, W. J. (2017). Monkeys and humans take local uncertainty into account when localizing a change [Publisher: The Association for Research in Vision and Ophthalmology]. *Journal of Vision*, *17*(11), 4–4.
<https://doi.org/10.1167/17.11.4>
- Donkin, C., Nosofsky, R. M., Gold, J. M., & Shiffrin, R. M. (2013). Discrete-slots models of visual working-memory response times. *Psychological Review*, *120*(4), 873–902. <https://doi.org/10.1037/a0034247>
- Change detection task für Projekt 2 Experiment 2
- Dunn, J. C., & Anderson, L. (2018). Signed difference analysis: Testing for structure under monotonicity. *Journal of Mathematical Psychology*, *85*, 36–54.
<https://doi.org/10.1016/j.jmp.2018.07.002>
- Erdfelder, E., Küpper-Tetzel, C. E., & Mattern, S. D. (2011). Threshold models of recognition and the recognition heuristic. *Judgment and Decision Making*, *6*(1), 16.
- Gelman, A., & Shalizi, C. R. (2013). Philosophy and the practice of Bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, *66*(1), 8–38.
<https://doi.org/10.1111/j.2044-8317.2011.02037.x>

- Goodrich, B., Gabry, J., Ali, I., & Brilleman, S. (2020). Rstanarm: Bayesian applied regression modeling via Stan. [R package version 2.21.1].
<https://mc-stan.org/rstanarm>
- Hardman, K. O., Vergauwe, E., & Ricker, T. J. (2017). Categorical Working Memory Representations are used in Delayed Estimation of Continuous Colors. *Journal of experimental psychology. Human perception and performance*, 43(1), 30–54.
<https://doi.org/10.1037/xhp0000290>
- Henninger, F., Shevchenko, Y., Mertens, U., Kieslich, P. J., & Hilbig, B. E. (2021). Lab.js: A free, open, online experiment builder.
<https://doi.org/10.5281/zenodo.4584705>
- Kellen, D., & Klauer, K. C. (2011). Evaluating models of recognition memory using first- and second-choice responses. *Journal of Mathematical Psychology*, 55(3), 251–266. <https://doi.org/10.1016/j.jmp.2010.11.004>
- Kellen, D., & Klauer, K. C. (2015). Signal detection and threshold modeling of confidence-rating ROCs: A critical test with minimal assumptions. *Psychological Review*, 122(3), 542–557. <https://doi.org/10.1037/a0039251>
- Kellen, D., & Klauer, K. C. (2018). Elementary Signal Detection and Threshold Theory, In *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*. American Cancer Society. <https://doi.org/10.1002/9781119170174.epcn505>
- Kellen, D., Winiger, S., Dunn, J. C., & Singmann, H. (in press). *Testing the Foundations of Signal Detection Theory in Recognition Memory* (preprint). PsyArXiv. <https://doi.org/10.31234/osf.io/p5rj9>
- Keshvari, S., van den Berg, R., & Ma, W. J. (2013). No Evidence for an Item Limit in Change Detection (L. T. Maloney, Ed.). *PLoS Computational Biology*, 9(2), e1002927. <https://doi.org/10.1371/journal.pcbi.1002927>
- Klauer, K. C. (2010). Hierarchical Multinomial Processing Tree Models: A Latent-Trait Approach. *Psychometrika*, 75(1), 70–98.
<https://doi.org/10.1007/s11336-009-9141-0>

- Klauser, K. C., & Kellen, D. (2010). Toward a complete decision model of item and source recognition: A discrete-state approach. *Psychonomic Bulletin & Review*, 17(4), 465–478. <https://doi.org/10.3758/PBR.17.4.465>
- Lee, M., & Wagenmakers, E.-J. (2013). Bayesian data analysis for cognitive science: A practical course. New York, NY: Cambridge University Press.
- Luce, R. D. (1963). A Threshold Theory for Simple Detection Experiments. *Psychological Review*, 70(1), 61–79. <https://doi.org/10.1037/h0039723>
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279–281. <https://doi.org/10.1038/36846>
- Moran, R. (2016). Thou shalt identify! The identifiability of two high-threshold models in confidence-rating recognition (and super-recognition) paradigms. *Journal of Mathematical Psychology*, 73, 1–11. <https://doi.org/10.1016/j.jmp.2016.03.002>
- Nosofsky, R. M., & Donkin, C. (2016). Response-time evidence for mixed memory states in a sequential-presentation change-detection task. *Cognitive Psychology*, 84, 31–62. <https://doi.org/10.1016/j.cogpsych.2015.11.001>
- Nosofsky, R. M., & Gold, J. M. (2018). Biased guessing in a complete-identification visual-working-memory task: Further evidence for mixed-state models [Publisher: American Psychological Association]. *Journal of Experimental Psychology: Human Perception and Performance*, 44(4), 603–625. <https://doi.org/10.1037/xhp0000482>
- Oberauer, K. (submitted). *Measurement Models for Visual Working Memory – A Factorial Model Comparison* (tech. rep.). PsyArXiv. <https://doi.org/10.31234/osf.io/9vwhx>
- Province, J. M., & Rouder, J. N. (2012). Evidence for discrete-state processing in recognition memory. *Proceedings of the National Academy of Sciences of the United States of America*, 109(36), 14357–14362. Retrieved February 3, 2017, from <http://www.jstor.org/stable/41706212>

- Ricker, T. J., Thiele, J. E., Swagman, A. R., & Rouder, J. N. (2016). Recognition Decisions From Visual Working Memory Are Mediated by Continuous Latent Strengths. *Cognitive Science*, n/a–n/a. <https://doi.org/10.1111/cogs.12436>
- Rouder, J. N., Morey, R. D., Cowan, N., Zwillling, C. E., Morey, C. C., & Pratte, M. S. (2008). An assessment of fixed-capacity models of visual working memory. *Proceedings of the National Academy of Sciences*, 105(16), 5975–5979. <https://doi.org/10.1073/pnas.0711295105>
- Rouder, J. N., & Morey, R. D. (2009). The nature of psychological thresholds [Publisher: American Psychological Association]. *Psychological Review*, 116(3), 655–660. <https://doi.org/10.1037/a0016413>
- Shadish, W., Cook, T., & Campbell, D. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Houghton Mifflin Company, Boston.
- Taylor, R., Thomson, H., Sutton, D., & Donkin, C. (2017). Does working memory have a single capacity limit? *Journal of Memory and Language*, 93, 67–81. <https://doi.org/10.1016/j.jml.2016.09.004>
- van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial comparison of working memory models. *Psychological Review*, 121(1), 124–149. <https://doi.org/10.1037/a0035234>
- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, 4(12), 11–11. <https://doi.org/10.1167/4.12.11>
- Winiger, S., Kellen, D., & Singmann, H. (in press). *Bias in confidence: A critical test for discrete-state models of change detection* (preprint). PsyArXiv. <https://doi.org/10.31234/osf.io/6cm5p>
- Xie, W., & Zhang, W. (2017). Dissociations of the number and precision of visual short-term memory representations in change detection. *Memory & Cognition*, 45(8), 1423–1437. <https://doi.org/10.3758/s13421-017-0739-7>
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory [Number: 7192 Publisher: Nature Publishing Group]. *Nature*, 453(7192), 233–235. <https://doi.org/10.1038/nature06860>

Zhang, W., & Luck, S. J. (2009). Sudden Death and Gradual Decay in Visual Working Memory. *Psychological Science*, *20*(4), 423–428.

<https://doi.org/10.1111/j.1467-9280.2009.02322.x>