Précis of *Characterizing the limits of visual working memory*

Kirsten C. S. Adam

1. Background and overview

Working memory is a core facet of complex cognition

Common everyday activities require the complex interaction of many cognitive processes. For example, when driving to the store we must, among other things, perceive rapid changes to our surroundings, shift our attention between the road and mirrors, temporarily track and remember approaching cars, and retrieve directions from long-term memory. In such real-world situations, overlapping cognitive processes work in tandem and can be difficult to disentangle, but decades of cognitive science experiments have carefully dissected the constituent cognitive processes that contribute to complex behavior. Ongoing work seeks to further understand the separable contributions of processes like attention, working memory, and long-term memory.

Working memory, in particular, plays a central role in complex cognition. Akin to a "buffer" (in a computer analogy), working memory acts as an interface between perception, memory, and action. Consistent with working memory's role as the active workspace of complex cognition, individual differences in working memory performance predict individual differences in fluid intelligence (1, 2), academic attainment (3), and standardized test scores (4). In addition to predicting individual differences in young, healthy college students, working memory is likewise frequently disrupted in many clinical disorders, including schizophrenia (5–7), Parkinson's (8), and Attention-Deficit/Hyperactivity Disorder (9). In sum, decades of individual differences work have revealed a strong correlative relationship between the amount of information that may be held in mind and the ability to perform complex cognitive tasks.

Despite its importance for cognition, the amount of information that can be held in working memory is sharply limited. This *online storage limit* is thus a critical limiting factor in how much information we can juggle in mind. The experiments in this dissertation are motivated by two core questions relating to the broad impact of working memory on cognition, and to its underlying limits.

Are individual differences in working memory performance driven by capacity or consistency?

Why does working memory performance predict so many other behaviors and outcomes? The dominant interpretation has been that working memory *capacity* differs from person to person. Here, the term *capacity* implies that people differ in the ceiling – some people can store less information than others. If you can hold fewer items in mind, you then have fewer resources available for other tasks like problem-solving. Previous research on working memory has assumed that individuals differ in capacity; as such, visual working memory performance metrics (e.g. hit and false alarm rates) are typically converted into a single capacity score, "K" (10–12). In part because capacity scores appear to explain individual differences, this assumption has been largely untested. In Chapter 2, I test an alternative model in which the *consistency* of performance explains individual differences. In this alternative model, most participants share the same upper-bound on performance, but differ in how often they achieve their capacity (i.e., they frequently store 0 or 1 items even though they are capable of storing 3). Support for a consistency account would invite major reinterpretation of much prior work on individual differences in working memory, and would suggest new avenues for interventions aimed at improving working memory function.

Are information limits due to random guesses or to low-resolution representations?

Although there is broad consensus that working memory can only hold a limited amount of information, there is stark disagreement as to how this information is distributed across remembered items. Some classes of models propose that the capacity of working memory is limited (13–15); if presented with more items than one can hold in mind, one will store only a sub-set of the items. Other classes of models propose that the resolution of remembered items is limited, rather than the number of items *per se* (13, 16–18). In resolution-limited models, people will store all items (though very imprecisely) when presented with a very large number of items. Much prior work has been devoted to this question, but has reached an empirical stalemate, in part because refined versions of both classes of model provide strong fits to observed data and, thus, fight to a tie (13). To break this stalemate, experiments in Chapters 3 and 4 use new behavioral task designs and neural decoding approaches to provide more diagnostic evidence for discriminating between these models. New experimental traction on the question of capacity limits is important for building robust cognitive and neural models of working memory storage, with potential implications for understanding how and why working memory is disrupted in clinical disorders.

The work in this dissertation uses a complementary blend of novel task variants, behavioral modeling, and electroencephalogram (EEG) measures to interrogate the nature of information limits in visual working memory and to provide new insights into the limitations of this core cognitive system.

2. Chapter 2: The contribution of attentional lapses to individual differences in visual working memory capacity

Published as Adam, Mance, Fukuda & Vogel (2015), Journal of Cognitive Neuroscience

Chapter 2 used new approaches to ask whether individual differences in working memory performance are best explained by *capacity differences* or *consistency differences*.

Typical working memory measures are inadequate for measuring trial-to-trial variation in performance because they only probe a single item on each trial. To quantify the consistency of working memory performance, we designed a discrete whole-report task (Figure 1). In this task, participants report all items from the memory array. In Experiment 2-1, we examined performance across many set sizes (1-6) and verified that this new task yielded similar behavioral performance as a typical change detection task. In Experiments 2-2 and 2-3, we held set size constant at 6 items. Holding the set size constant allowed us to examine endogenous fluctuations in performance while holding external task demands constant.

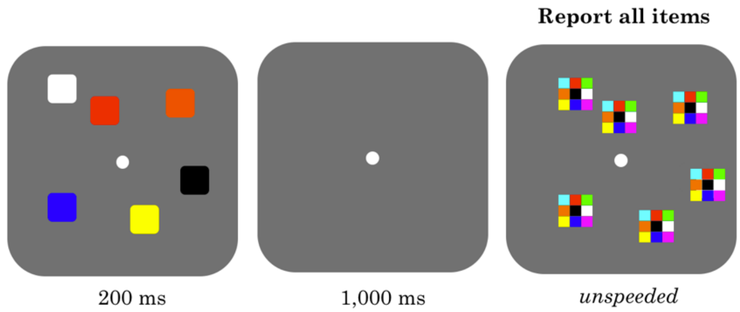

Figure 1. Discrete whole report task (Chapter 2). Participants remember the color of each square. At response, they use the mouse to click a color for each "response grid" at each location. They can report the items in any order, but must respond to all items.

Working memory performance fluctuates from moment to moment

Each trial of the discrete whole-report task yields an accuracy score that ranges from 0 items correct to the maximum number (e.g. 6 items), allowing precise tracking of trial-to-trial variations in performance (Figure 2). Both high- and low-performing participants experience "failures", where they perform no better than random chance (they get only 0 or 1 item correct out of 6). However, the frequency of these failures varies drastically across participants. The best performing participant had 0 failure trials throughout the session, whereas the worst performing participant had >35% failure trials.

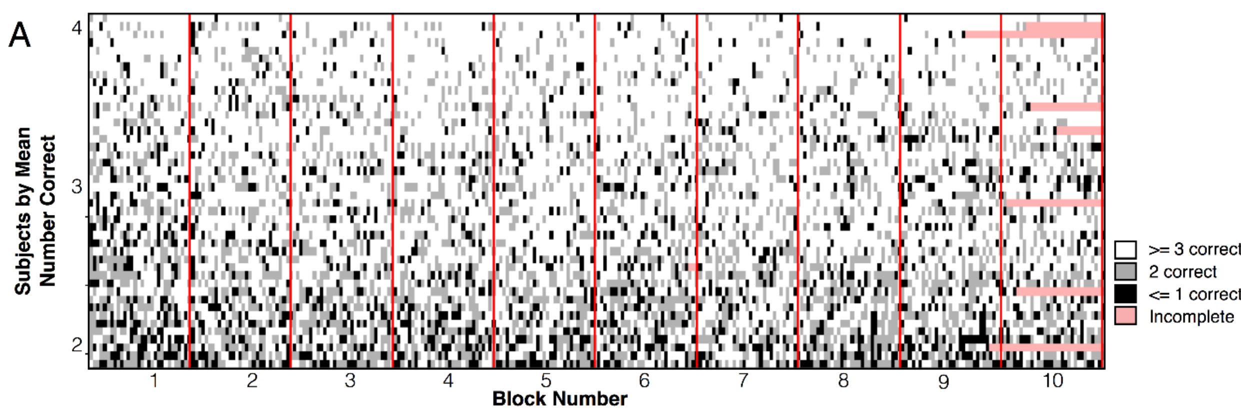

Figure 2. Single trial performance fluctuations in Experiment 2-2. Participants are sorted on the y-axis by average memory performance. Performance fluctuates across trials, but each individual's fluctuations are relatively consistent.

Lapse and attentional control models of attention's influence on working memory performance

We developed computational models with parameters for both capacity and consistency. The capacity parameter of each model corresponded to the maximum number of items ever remembered by the participant. The model also allowed that participants can sometimes score higher than capacity because of random guesses (e.g., if you guess, you will sometimes be correct by chance). We tested two alternative models of how working memory capacity fluctuates over time: all-or-none lapses or graded loss of attentional control. All model-fitting procedures were blind to the full distribution of single-trial accuracy values. Instead, each model tried to predict mean accuracy; fit of each model was then assessed by comparing its predictions to the full performance distribution for each participant.

The "all-or-none lapse model" was Bernoulli-distributed; this model assumed that either participants store their full capacity, K, or they store nothing and guess randomly for all items. This model was based on previous work that added a lapse parameter to improve fits to models of change detection performance (19). The lapse model fit the data poorly, as it predicted a bimodal distribution of accuracy values: one peak at 0 correct (participant lapsed) and another peak at capacity. In contrast, the true performance distributions were unimodal. As shown in Figure 3, we also found that the vast majority of participants shared the same performance mode of 3 items correct.

In contrast, the "attentional control" model assumed that participants experienced graded fluctuations of attention that caused them to under-achieve capacity. These fluctuations were modeled by a beta distribution (α parameter could vary from zero to six in steps of .01; the β parameter was constrained to one). Critically, this beta distribution dictated that performance could fluctuate only *below* capacity. To simulate performance for a single working memory trial, a value would be pulled from the beta distribution (range: 0 to 1), multiplied by the capacity parameter (K) and rounded to the nearest whole item (*n*); semi-random guesses were added for remaining responses (6-*n*), which sometimes led to trial accuracy above K. This model resulted in a unimodal distribution that closely matched participants' true performance, even though initial model selection was blind to the underlying distribution. Most importantly, this model also revealed that the vast majority of participants (84%) were best fit by a capacity of 3 items. The remaining participants were best fit by a capacity of 4. In other words – the worst-performing participants were still *capable* of maintaining 3 items in mind; they just did so much less consistently.

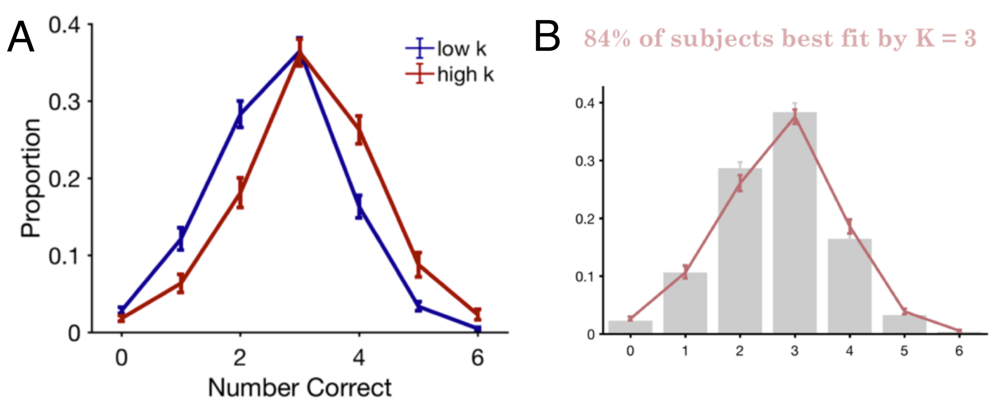

Figure 3. Performance distributions in Experiment 2-2. (a) Median split visualization of performance for the top- and bottom-half of participants. (b) Average behavioral data (pink line) and model (grey bars).

Trial-by-trial working memory performance can be predicted by frontal theta power and posterior alpha power

Finally, we examined neural (EEG) correlates of working memory fluctuations. We hypothesized that performance fluctuates because people actually store fewer items, and we further predicted that fluctuations of working memory storage are related to momentary failures of executive attention. However, there are several other possible sources of behavioral errors, and not all of these are memory-related. For example, poor-performing participants may simply be task non-compliant – they may blink or move their eyes more often during the brief presentation of the memory array.

As in Experiment 2-2, in Experiment 2-3 participants performed many trials of set size 6 so that we could examine endogenous fluctuations of performance. We split trials into "good performance" trials (>3 correct) and "poor performance" trials (<3 correct). We found evidence that people actually stored fewer items for poor performance relative to good performance trials, as indexed by alpha power suppression (20, 21). Frontal theta power also predicted behavioral performance, supporting the interpretation that working memory failures arise from failures of executive attention as frontal theta covaries with cognitive control and is anti-correlated with default mode network activity (22–24). Intriguingly, frontal theta power began to predict eventual behavioral performance ~700 ms before the memory array had appeared. This pre-trial signal highlights a potential target for performance improvements via real-time neural feedback.

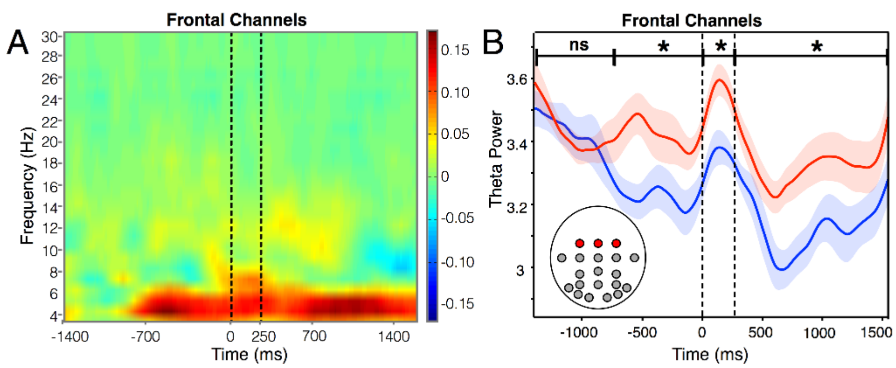

Figure 4. Frontal theta power predicts working memory success. (a) Time-frequency plot of power difference for "good trials" (>3 correct) minus "poor trials" (<3 correct) at frontal electrodes (F1, F2, and Fz). (b) Theta band power (4–7 Hz) for good (red) and poor (blue) performance trials

3. Chapter 3: Clear evidence for item limits in visual working memory

Published as Adam, Vogel & Awh (2017), Cognitive Psychology

The experiments in Chapter 2 revealed that "capacity" differences between participants are better described as differences in the consistency of visual working memory storage. Yet, these experiments could not speak directly to debates about the underlying cause of this mostly-common information limit (capacity vs. resolution). The experiments in Chapter 2 used a relatively coarse measure of memory fidelity (participants chose 1 of 9 highly discriminable colors). In Chapter 3, we developed a novel variant of the continuous report task to better distinguish between capacity and resolution-limited models.

Continuous whole-report task measures the fidelity of all items

Rather than probing a single item, we instead asked participants to recall all items in a continuous whole-report task (Figure 5). This yields a separate response error distribution for each response within each set size condition (rather than one error distribution per set size). In Experiments 3-1 and 3-1b, participants freely recalled the color of all items in any order that they wished. In Experiment 3-2a and 3-2b, participants had to respond to the item the computer randomly selected for them. This version controlled for the effects of output interference and time delays on response errors.

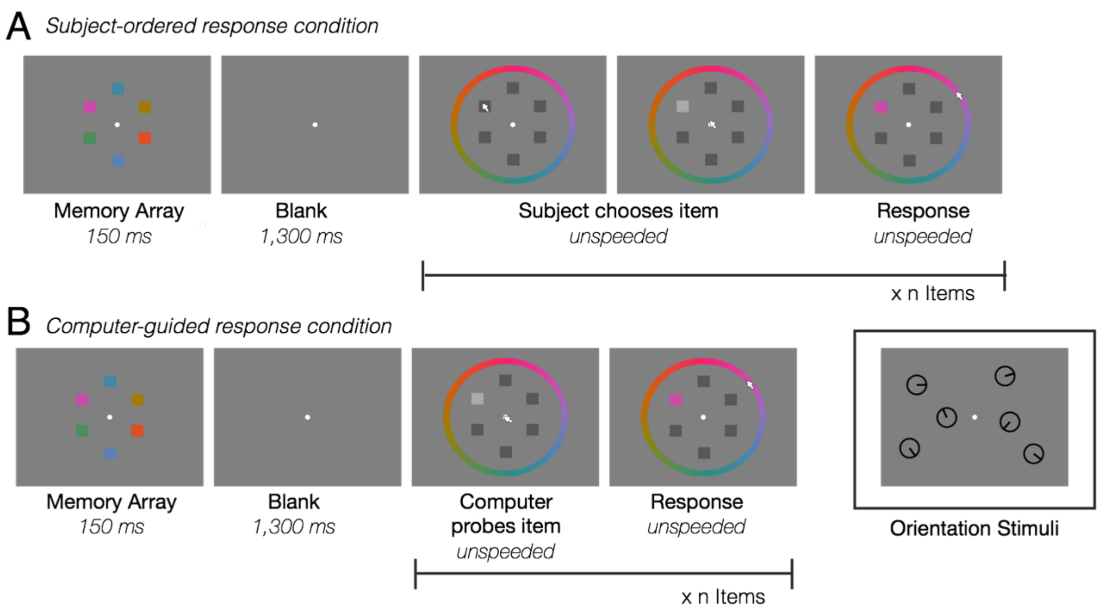

Figure 5. Continuous-whole report task (Chapter 3) (a) Experiment 3-1a: Participants freely recall the color of all remembered items in any order. (b) Experiment 3-2a: The computer probes memory for all items in a random order, controlling for the effects of output interference and time decay. (inset) Orientation stimuli (Experiments 3-1b and 3-2b).

Participants report no information about a subset of items from large arrays

A key distinguishing feature of two competitive classes of models is whether or not people randomly guess. When fitting models to typical single-probe error distributions, both models do extremely well (Figure 6). Mixture models can achieve excellent fits to the average distribution by assuming that some items are remembered (with varying precision) and others are not (uniform distribution). Continuous resource models can achieve excellent fits by including many low-precision representations. Thus, we tested whether random guesses are necessary to explain performance.

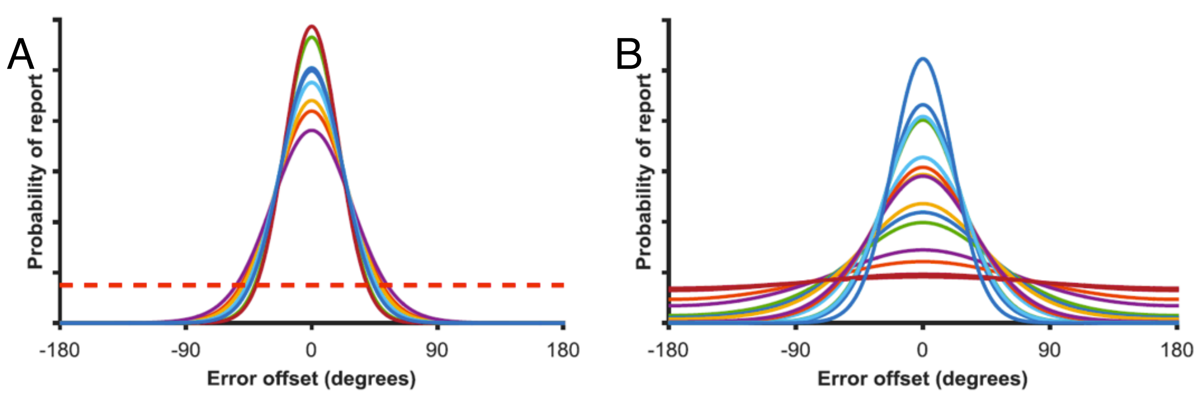

Figure 6. Hypothetical mixture of distributions that could yield an identical set-size 6 error distribution in the standard continuous report task. (a) A mixture model in which some items are stored (with varying

fidelity) and other items are forgotten (dashed line). (b) A variable precision model in which all items are stored with varying fidelity. Representations are allowed to be extremely imprecise.

Prior work established that participants have robust metaknowledge of the quality of remembered items (25, 26). In the free recall task, participants used this knowledge and tended to report the items in order of decreasing fidelity. Capacity-limited models predict that, when capacity is exceeded, the remaining responses for each participant will be uniformly-distributed. In contrast, resource models predict that *all* responses will contain useful information (i.e. non-uniform) about the memory item.

We ran a model comparison for each response distribution of each participant; this model comparison tested whether each response distribution was better fit by a zero-parameter uniform distribution or by any of several competing models (e.g., standard mixture model, variable precision model). All participants showed at least 1 set-size 6 response that was best described as uniform (Figure 7). As predicted, these uniform distributions fell at the end of the response period, and the vast majority of subjects had 2-3 uniform responses (consistent with a modal capacity limit of 3-4 items). Note, however, this method does not precisely estimate participants' true underlying capacity as in Chapter 2, as they might sometimes report well-remembered items last (thus inflating capacity estimates).

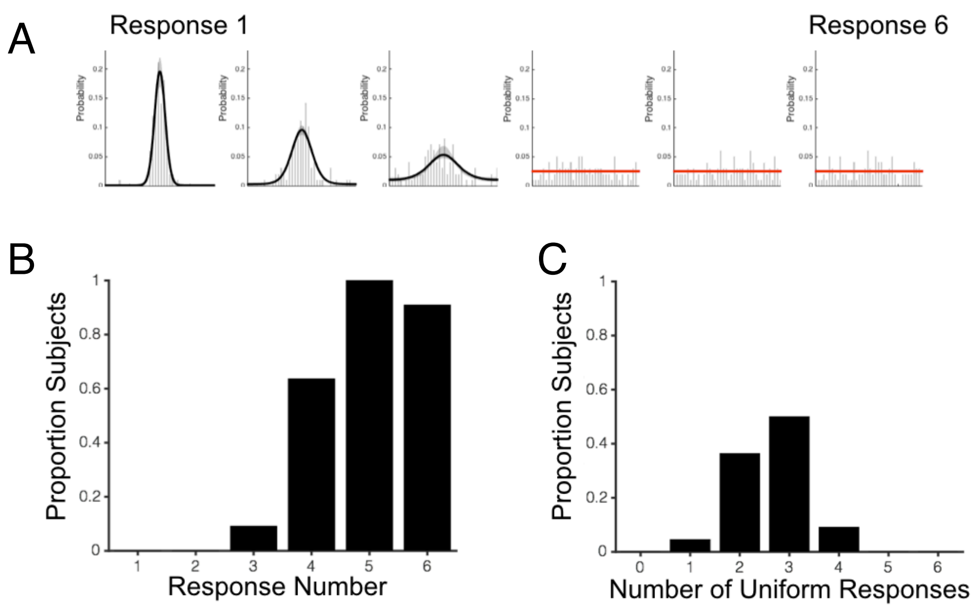

Figure 7. Most participants have no information beyond the 3rd or 4th response for set size 6 trials. (a) Set-size 6 responses for an example participant in the free-recall condition. Red fit lines indicate that a zero-parameter uniform distribution won the model competition. (b) Proportion of participants whose response was best fit by a uniform distribution. (c) Histogram of the number of uniform distributions (out of 6) for participants in the free recall experiment.

Importantly, control experiments (Experiments 3-2a and 3-2b) revealed that uniform responses are not due to output interference or temporal decay. In fact, fidelity across responses decreased very little in this control experiment; output interference could explain only 16% of the decline observed in the free recall experiment. Additional analyses revealed that participants have excellent metaknowledge of random guesses;

self-reported guess rates correlated with mixture model estimates of guess rates with a slope near 1.

Hypothesized low-resolution responses mimic random guesses.

Variable precision models assume that an aggregate distribution is made up of a mixture of many different widths of von Mises (circular normal) distributions, and some of these distributions are extremely imprecise. We used a simulation approach to test whether these very imprecise von Mises distributions may mimic random guesses.

First, we quantified the widths of von Mises distributions used by the variable precision model to explain actual set size 6 data (Figure 8). Then, we tested how many trials would be needed to discriminate imprecise von Mises distributions from a uniform distribution. This analysis revealed that even with 1,000,000 noise-free trials per subject, we would be unable to distinguish a significant percentage of the hypothesized von Mises distributions from uniform. This supports the notion that a significant proportion of responses are random guesses, and that these random guesses can be mimicked by extremely imprecise distributions.

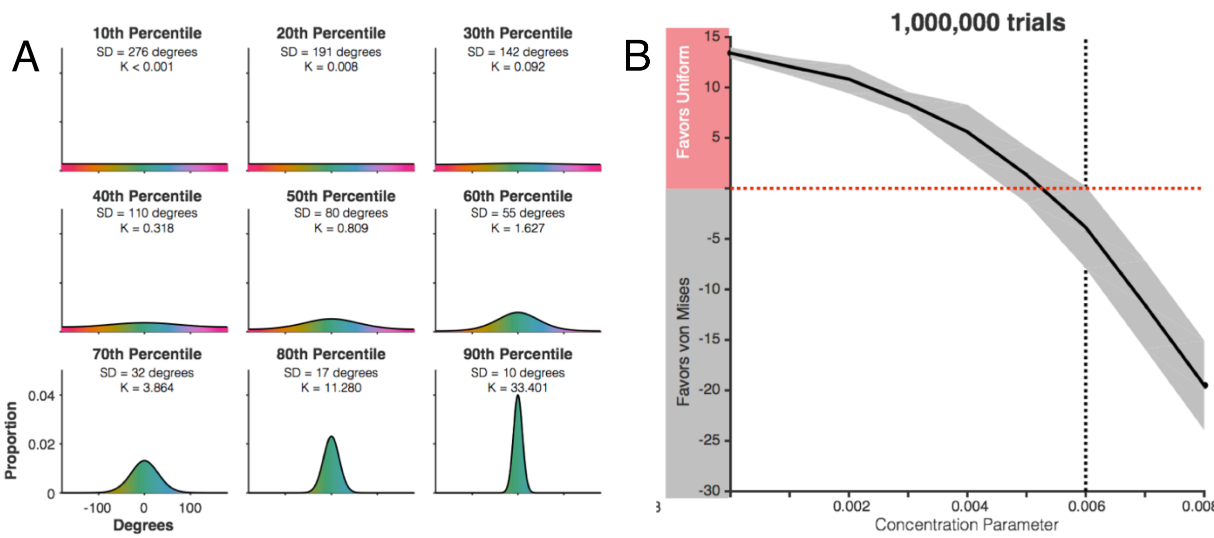

Figure 8. Illustration of the potential for variable precision models to mimic uniform responses. (a) Decile split for the von Mises distributions used by the variable precision model for Set Size 6 aggregate data. (b) Simulation demonstrating that even 1 million noise-free trials per participant would be inadequate to detect the information contained in such an imprecise distribution.

4. Chapter 4: Decoding the limits of simultaneous storage

The experiments in Chapter 3 found strong supporting evidence that a substantial proportion of participants' responses are random guesses. However, one weakness of these experiments is that they relied on purely behavioral measures that can probe only the final outcome of the trial. Thus, behavioral measures are somewhat insensitive to intermediate cognitive processes. For example, perhaps participants held some information about all items in mind for the majority of the delay period, but failed to retrieve this information at test. By acquiring neural evidence for whether each item was represented during the delay period, we could determine whether some items failed to be stored.

Univariate evidence for capacity limits in visual working memory

Prior work has shown that univariate neural measures of working memory storage increase as a function of set size (suggesting that more items are stored), but appear to reach a plateau once capacity is filled (27–29). This has frequently been interpreted as evidence for capacity limits – the signal is thought to saturate when capacity is reached. Yet, interpretation of these univariate measures is fraught, as they measure the aggregate activity to an entire array, rather than item-level storage. Further, recent work has questioned the validity of interpreting this "plateau" in neural activation (30). Rather than examining an average univariate signal for the entire array, we need methods for reading out the representation of each individual item while it is being held in mind.

The topography of alpha-band power tracks remembered locations

Prior work has shown that the location of a single item held in memory can be tracked by the topography of alpha-band power across the scalp (31) using an inverted encoding model (IEM) approach (32–34). The experiments in Chapter 4 take advantage of this IEM technique (Figure 9) to examine whether there was an active neural representation of each item throughout the delay interval.

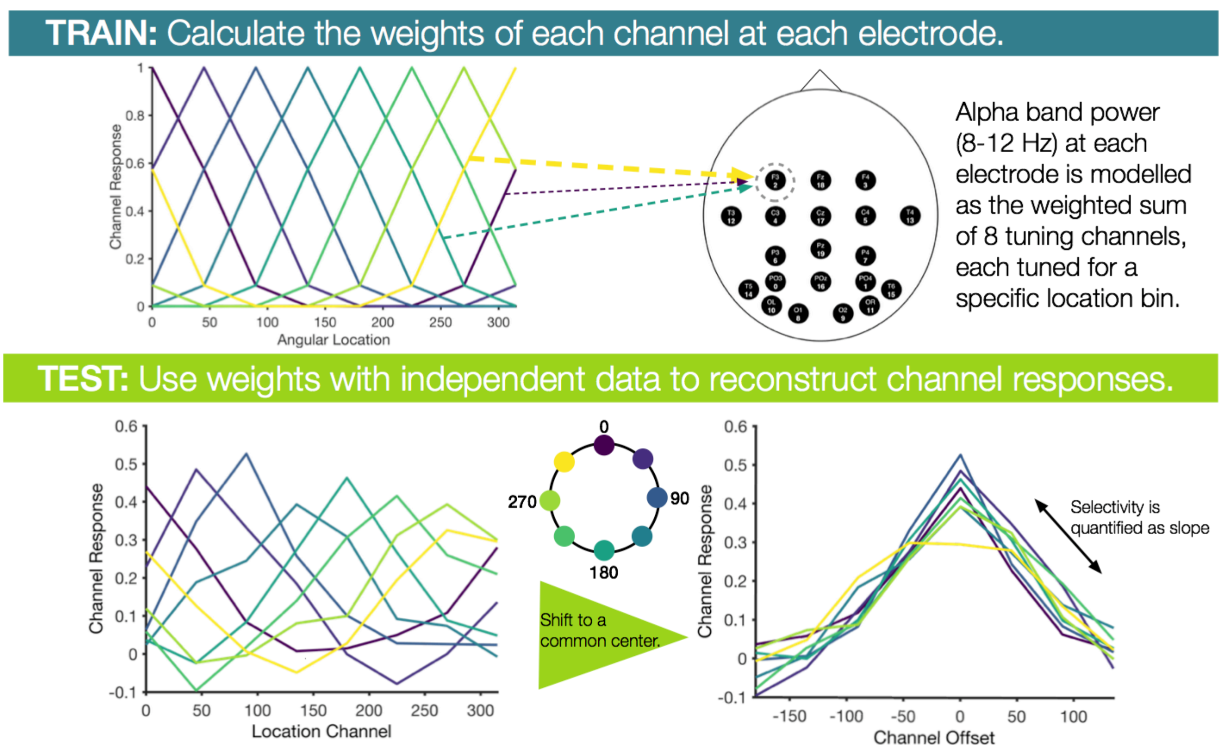

Figure 9. Inverted encoding model approach. The IEM assumes that the power at each electrode reflects the weighted sum of spatially selective information channels which reflect the activity of underlying neuronal populations tuned to each location.

For each participant, we partition data into independent training and test sets and apply an IEM routine at each time-point. Training data B_1 (m electrodes × n_1 observations) are used to estimate weights W (m electrodes x k channels) that approximate the relative contribution of eight equally-spaced spatial channels to the observed response C_1 (k

channels × n_1 measurements). The relationship between B_1, C_1, and W can be described by a general linear model:

$$B_1 = WC_1$$

We obtain the weight matrix through least-squares estimation:

$$\widehat{W} = B_1 C_1^{T} \left(C_1 C_1^{T}\right)^{-1}$$

In the test stage, we invert the model to transform the observed test data B_2 (m electrodes × n_2 observations) into estimated channel responses, C_2 (k channels × n_2 measurements), using the estimated weight matrix, $\widehat{W}$, that we obtained in the training phase:

$$\widehat{C_2} = \left(\widehat{W}^{T}\widehat{W}\right)^{-1}\widehat{W}^{T}B_2$$

Each estimated channel response function is circularly shifted to a common center (0 represents the remembered location). Here, the selectivity of tuning is quantified as the slope of the tuning function folded in half (larger numbers indicate more robust reconstruction).

Reconstructing multiple remembered stimulus locations using alpha-band power

First, in Experiment 4-1 we tested the ability of the model to generalize from 1 item to small multi-item arrays (2 and 3 items). Consistent with prior work (33, 35, 36), we found decreased fidelity of decoded representations for larger set sizes. That is, an average reconstruction of an item from a set size 3 array was more imprecise than an average item from a set size 2 array. In addition, by training the model on set size 1 data and testing on each item as a function of response order, we could examine the quality of the reconstruction for each item.

In Experiment 4-1, we could decode all items from the small arrays, but the critical question is whether we can decode all items from supra-capacity arrays. In Experiment 4-2, we examined whether we could reconstruct the spatial location of all items from a supra-capacity array (set size 6). As before, participants freely recalled all items, and tended to report their best items first. Thus, we predicted that we should be unable to reconstruct the spatial locations of items that are not stored (responses 4-6).

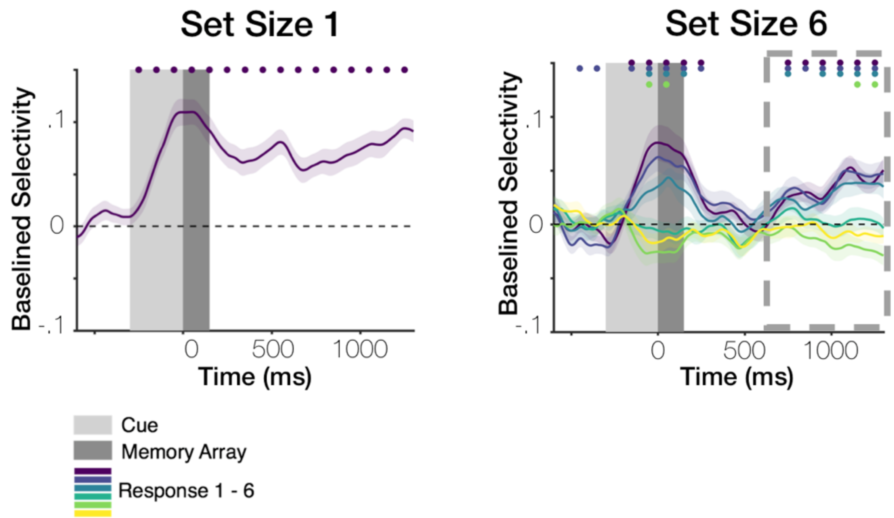

Figure 10. Selectivity of decoded representations for set sizes 1 and 6 in Experiment 4-2. An Inverted Encoding Model was trained on set size 1 data then tested on each response from set size 6 data. During the delay period (gray dotted box), we could significantly decode only 3 out of 6 items.

As predicted, we found that only a subset of items are encoded by the topographic distribution of alpha power, which tracks the active maintenance of spatial information in working memory (31, 37, 38) (Figure 10). These results provide converging evidence with the behavioral estimates of capacity limits, and demonstrate the feasibility of using encoding models to measure the contents of working memory in an item-specific fashion. This method has exciting potential for testing competing models of the neural codes supporting working memory (39–42).

5. Conclusions

Rethinking individual differences in working memory "capacity"

Prior individual differences work has assumed that working memory *capacity* explains individual variation in performance, but Chapter 2 revealed that the *consistency* of performance is key. This finding has important implications for understanding working memory deficits in developmental and clinical population, and potential applications for interventions targeted toward improving working memory performance. For example, many working memory training interventions have the explicit goal of raising individuals' capacity ceiling (43–46). My work suggests that, instead, interventions should focus on reducing the frequency of abject failures (47, 48).

New avenues for neural correlates of working memory performance

In Experiment 2-3, we showed that frontal theta power predicts behavioral performance even before the memory array has appeared. We have since replicated this finding, and this signal is a promising target for real-time neural feedback. In Experiments 3-1 and 3-2, we demonstrated that we can use an inverted encoding model to examine how the representational quality of a neural code varies across individual items within an array. This work thus represents an exciting step forward from prior work that has examined the average representation of all items from a multi-item array (33, 35, 36).

Harnessing the power of trial-to-trial and item-to-item variability

A key approach in these experiments was the development of measures that are sensitive to trial-by-trial and item-by-item variation in memory success. These methods allowed us to more directly examine the variation in mnemonic quality across items, and allowed us to generate new predictions for old models. Future work examining fluctuations of performance in other aspects of cognition may likewise reveal an important role of temporal and attention-related dynamics (e.g., successfully discriminating faint targets, trial-to-trial success in abstract reasoning).

References

1. Engle RW, Kane MJ, Tuholski SW (1999) Individual Differences in Working Memory Capacity and What They Tell Us About Controlled Attention, General Fluid Intelligence, and Functions of the Prefrontal Cortex. *Models of Working Memory*, eds Miyake A, Shah P (Cambridge University Press, Cambridge), pp 102–134.

2. Unsworth N, Fukuda K, Awh E, Vogel EK (2014) Working memory and fluid intelligence: Capacity, attention control, and secondary memory retrieval. *Cognitive Psychology* 71:1–26.

3. Alloway TP, Alloway RG (2010) Investigating the predictive roles of working memory and IQ in academic attainment. *Journal of Experimental Child Psychology* 106(1):20–29.

4. Daneman M, Carpenter PA (1980) Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior* 19(4):450–466.

5. Erickson MA, et al. (2015) Impaired Working Memory Capacity Is Not Caused by Failures of Selective Attention in Schizophrenia. *Schizophrenia Bulletin* 41(2):366–373.

6. Gold JM, Wilk CM, McMahon RP, Buchanan RW, Luck SJ (2003) Working memory for visual features and conjunctions in schizophrenia. *Journal of Abnormal Psychology* 112(1):61–71.

7. Hahn B, Robinson BM, Leonard CJ, Luck SJ, Gold JM (2018) Posterior Parietal Cortex Dysfunction Is Central to Working Memory Storage and Broad Cognitive Deficits in Schizophrenia. *The Journal of Neuroscience* 38(39):8378–8387.

8. Lee E-Y, et al. (2010) Visual working memory deficits in patients with Parkinson's disease are due to both reduced storage capacity and impaired ability to filter out irrelevant information. *Brain* 133(9):2677–2689.

9. Martinussen R, Hayden J, Hogg-Johnson S, Tannock R (2005) A Meta-Analysis of Working Memory Impairments in Children With Attention-Deficit/Hyperactivity Disorder. *Journal of the American Academy of Child & Adolescent Psychiatry* 44(4):377–384.

10. Cowan N (2001) The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav Brain Sci* 24(1):87–114; discussion 114-185.

11. Pashler H (1988) Familiarity and visual change detection. *Perception & Psychophysics* 44(4):369–378.

12. Rouder JN, Morey RD, Morey CC, Cowan N (2011) How to measure working memory capacity in the change detection paradigm. *Psychonomic Bulletin & Review* 18(2):324–330.

13. van den Berg R, Awh E, Ma WJ (2014) Factorial comparison of working memory models. *Psychological Review* 121(1):124–149.

14. Zhang W, Luck SJ (2008) Discrete fixed-resolution representations in visual working memory. *Nature* 453(7192):233–235.

15. Zhang W, Luck SJ (2011) The Number and Quality of Representations in Working Memory. *Psychological Science* 22(11):1434–1441.

16. Bays PM, Catalao RFG, Husain M (2009) The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision* 9(10):7.

17. van den Berg R, Shin H, Chou W-C, George R, Ma WJ (2012) Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences* 109(22):8780–8785.

18. Wilken P, Ma WJ (2004) A detection theory account of change detection. *J Vis* 4(12):1120–1135.

19. Rouder JN, et al. (2008) An assessment of fixed-capacity models of visual working memory. *Proc Natl Acad Sci USA* 105(16):5975–5979.

20. Fukuda K, Kang M-S, Woodman GF (2016) Distinct neural mechanisms for spatially lateralized and spatially global visual working memory representations. *Journal of Neurophysiology* 116(4):1715–1727.

21. Fukuda K, Mance I, Vogel EK (2015) Alpha Power Modulation and Event-Related Slow Wave Provide Dissociable Correlates of Visual Working Memory. *Journal of Neuroscience* 35(41):14009–14016.

22. Cavanagh JF, Frank MJ (2014) Frontal theta as a mechanism for cognitive control. *Trends in Cognitive Sciences* 18(8):414–421.

23. Itthipuripat S, Wessel JR, Aron AR (2013) Frontal theta is a signature of successful working memory manipulation. *Exp Brain Res* 224(2):255–262.

24. Scheeringa R, et al. (2008) Frontal theta EEG activity correlates negatively with the default mode network in resting state. *Int J Psychophysiol* 67(3):242–251.

25. Fougnie D, Suchow JW, Alvarez GA (2012) Variability in the quality of visual working memory. *Nature Communications* 3:1229.

26. Rademaker RL, Tredway CH, Tong F (2012) Introspective judgments predict the precision and likelihood of successful maintenance of visual working memory. *Journal of Vision* 12(13):21–21.

27. Todd JJ, Marois R (2004) Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature* 428(6984):751–754.

28. Vogel EK, Machizawa MG (2004) Neural activity predicts individual differences in visual working memory capacity. *Nature* 428(6984):748–751.

29. Vogel EK, McCollough AW, Machizawa MG (2005) Neural measures reveal individual differences in controlling access to working memory. *Nature* 438(7067):500–503.

30. Bays PM (2018) Reassessing the Evidence for Capacity Limits in Neural Signals Related to Working Memory. *Cerebral Cortex* 28(4):1432–1438.

31. Foster JJ, Sutterer DW, Serences JT, Vogel EK, Awh E (2016) The topography of alpha-band activity tracks the content of spatial working memory. *Journal of Neurophysiology* 115(1):168–177.

32. Brouwer GJ, Heeger DJ (2009) Decoding and Reconstructing Color from Responses in Human Visual Cortex. *Journal of Neuroscience* 29(44):13992–14003.

33. Sprague TC, Ester EF, Serences JT (2014) Reconstructions of Information in Visual Spatial Working Memory Degrade with Memory Load. *Current Biology* 24(18):2174–2180.

34. Sprague TC, Serences JT (2013) Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience* 16(12):1879–1887.

35. Buschman TJ, Siegel M, Roy JE, Miller EK (2011) Neural substrates of cognitive capacity limitations. *Proceedings of the National Academy of Sciences* 108(27):11252–11255.

36. Emrich SM, Riggall AC, LaRocque JJ, Postle BR (2013) Distributed Patterns of Activity in Sensory Cortex Reflect the Precision of Multiple Items Maintained in Visual Short-Term Memory. *Journal of Neuroscience* 33(15):6516–6523.

37. Foster JJ, Bsales EM, Jaffe RJ, Awh E (2017) Alpha-Band Activity Reveals Spontaneous Representations of Spatial Position in Visual Working Memory. *Current Biology* 27(20):3216-3223.e6.

38. Sutterer DW, Foster JJ, Adam KCS, Vogel EK, Awh E (in prep) Delay-period activity encodes multiple items stored in working memory.

39. Bays PM (2015) Spikes not slots: noise in neural populations limits working memory. *Trends in Cognitive Sciences* 19(8):431–438.

40. Bouchacourt F, Buschman TJ (2018) A Flexible Model of Working Memory. *bioRxiv*. doi:10.1101/407700.

41. Standage D, Paré M (2018) Slot-like capacity and resource-like coding in a neural model of multiple-item working memory. *Journal of Neurophysiology* 120(4):1945–1961.

42. Wei Z, Wang X-J, Wang D-H (2012) From Distributed Resources to Limited Slots in Multiple-Item Working Memory: A Spiking Network Model with Normalization. *Journal of Neuroscience* 32(33):11228–11240.

43. Jaeggi SM, Buschkuehl M, Jonides J, Perrig WJ (2008) Improving fluid intelligence with training on working memory. *Proceedings of the National Academy of Sciences* 105(19):6829–6833.

44. Buschkuehl M, Jaeggi SM, Mueller ST, Shah P, Jonides J (2017) Training Change Detection Leads to Substantial Task-Specific Improvement. *Journal of Cognitive Enhancement* 1(4):419–433.

45. Melby-Lervåg M, Hulme C (2013) Is working memory training effective? A meta-analytic review. *Developmental Psychology* 49(2):270–291.

46. Shipstead Z, Redick TS, Engle RW (2012) Is working memory training effective? *Psychological Bulletin* 138(4):628–654.

47. Adam KCS, Vogel EK (2016) Reducing failures of working memory with performance feedback. *Psychonomic Bulletin & Review* 23(5):1520–1527.

48. Adam KCS, Vogel EK (2018) Improvements to visual working memory performance with practice and feedback. *PLOS ONE* 13(8):e0203279.

49. Adam KCS, Robison MK, Vogel EK (2018) Contralateral Delay Activity Tracks Fluctuations in Working Memory Performance. *Journal of Cognitive Neuroscience*:1–12.