

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Chunks are not “Content-Free”: Hierarchical Representations Preserve Perceptual Detail within Chunks

Permalink

<https://escholarship.org/uc/item/2b0846p3>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN

1069-7977

Authors

Allen, Michael G
Destefano, Isabella
Brady, Timothy F.

Publication Date

2021

Peer reviewed

Chunks are not “Content-Free”: Hierarchical Representations Preserve Perceptual Detail within Chunks

Michael G. Allen (miallen@ucsd.edu)

University of California San Diego, Department of Cognitive Science
9500 Gilman Dr., La Jolla, CA 92093 USA

Isabella C. DeStefano (idestefa@ucsd.edu)

University of California San Diego, Department of Psychology
9500 Gilman Dr., La Jolla, CA 92093 USA

Timothy F. Brady (timbrady@ucsd.edu)

University of California San Diego, Department of Psychology,
9500 Gilman Dr., La Jolla, CA 92093 USA

Abstract

Chunks allow us to use long-term knowledge to efficiently represent the world in working memory. Most views of chunking assume that when we use chunks, this results in the loss of specific perceptual details, since it is presumed the contents of chunks are decoded from long-term memory rather than reflecting the exact details of the item that was presented. However, in two experiments, we find that in situations where participants make use of chunks to improve visual working memory, access to instance-specific perceptual detail (that cannot be retrieved from long-term memory) increased, rather than decreased. This supports an alternative view: that chunks facilitate the encoding and retention into memory of perceptual details as part of structured, hierarchical memories, rather than serving as mere “content-free” pointers. It also provides a strong contrast to accounts in which working memory capacity is assumed to be exhaustively described by the number of chunks remembered.

Keywords: chunks; chunking; hierarchical representations; distributed memory; visual working memory; perception; objects; binding

Introduction

A core principle of learning and memory is that we compress frequently encountered information into “chunks”, and these chunks allow us to efficiently represent the world (e.g., Feigenbaum & Simon, 1984). Chunking has long been known to enhance our otherwise extremely limited working memory capacity by allowing us to use learned knowledge to structure working memory (Miller, 1956; Cowan, 2005). For example, if remembering the set of letters “C”, “G”, “Q”, there exists no long-term associations that will help facilitate performance; whereas when remembering the letters “F”, “B”, “I”, existing associations will help connect these items together and thus they can be remembered as a single “chunk”, improving working memory performance.

Despite a long history of work on chunking, key questions about the very nature of chunk storage and representation in working memory remain unaddressed. One particularly critical question is whether access to perceptual details of the particular instantiation is lost when you use a learned chunk to remember something in working memory. Consider the yellow circle inside of a white circle in Figure 1. If these were

two arbitrary colors, you would need to store them as two separate representations to remember the inner and outer color (Wheeler & Treisman, 2002). However, in trying to remember this particular set of two colors, you might notice that this color pair corresponds to an existing chunk in your mind – “egg” – and use this to efficiently represent the colors. When you do this, does your working memory representation consist solely of this abstraction (“egg”), with all item-specific information accessed only by decoding the chunk using long-term memories (Fig. 1, left)? If so, this would suggest specific ways the perceptual instantiation of this chunk in this situation would be largely lost (e.g., you would know it was white and yellow, but not *which exact* white and yellow) – a ‘content-free’ view of what is stored in working memory. Alternatively, it could be that the ‘egg’ chunk benefits memory by serving as an additional retrieval cue to structure memory – a distributed model of memory where you use the chunk pointer, like “egg”, to help you encode and maintain the specific colors present in this instance, rather than store it solely as an abstraction (Fig. 1, right). Such a view sees memory as distributed and hierarchical, with chunks *and* perceptual details maintained in parallel in working memory.

The extent to which chunks serve as content-free labels vs serve as retrieval cues to distributed memory representations is a critical question. The most common view within the study of working memory, is to see chunks in working memory as ‘content-free’ – single-level representations which rely on decoding from long-term memory to access their content. This is sometimes explicitly claimed (e.g., Huang & Awh, 2018), or sometimes implicit in the manner that chunks are understood. For example, chunks are often used to measure working memory capacity, and this is only a useful measure of capacity if the chunk “pointers” (Ngiam et al, 2019) alone make up the “capacity” of the system (e.g., Cowan, 2001).

Although less commonly considered, the distributed and hierarchical view of chunks is also plausible and broadly consistent with literature from visual working memory, where, for example, visual objects seem to serve as ‘chunks’ (Cowan, 2001) but do not result in all-or-none abstractions




Figure 1. The nature of chunking. When remembering a familiar chunk, do people remember solely the chunk ‘label’ or do they use this label to access the perceptual details?

(e.g., Brady et al. 2011). In fact, visual features ‘within’ an object tend to be loosely structured, and forgotten independently (Fougnie & Alvarez, 2011), more consistent with a distributed, hierarchical account of how objects structure memory. Chunks could function in a similar manner, providing structure but not serving as solely pointers. This view is aligned with “template theory” from studies of chess expertise. In this theory, templates are seen as structures that develop via chunk learning. Their “core” information consists of details that are regular across exposures and that are crucial for recognition and activation of the template. However, rather than being “content-free”, they also contain “slots” which, in the case of chess, are positions that may be filled in variable, instance-specific ways. Crucially, experts using such a template have a greater memory for the instance-specific pieces, as well as the regular, core properties of the learned template (Gobet & Simon, 1996).

With few exceptions (e.g., Brady et al., 2009, Huang & Awh, 2018), the existence of templates or chunks that can facilitate the encoding and retrieval of instance-specific, perceptual details has almost never been directly tested within the study of working memory, however. In part this is because chunking has traditionally been studied with discrete stimuli (e.g., Gobet & Charness, 2018; Cowan, Chen & Rouder, 2004). Where chess positions allow for some variability between pre-learned positions and current instantiations, verbal stimuli generally do not raise the issue of whether perceptual details remain accessible within chunks, because the chunk ‘pointer’ alone is sufficient to capture all of the verbal information (De Groot, Gobet and Jongman, 1996): if you chunk F, B and I into FBI, there are no remaining details not captured by the chunk. Simply knowing ‘FBI’ was the relevant chunk allows you to “decode” all the verbal information that was present.

The study of visual working memory necessitates a revision of theories of chunking to allow for continuous stimuli typical in real-life vision. Without understanding what chunks in visual working memory consist of, we cannot interpret studies in which chunking improves apparent working memory capacity. For example, Brady et al. (2009) showed subjects 8 colors in a memory display arranged into 4 concentric colored pairs (as in the ‘egg’ in Figure 1). In a ‘patterned’ condition, the same color pairings were highly likely to recur throughout the experiments, whereas in a

matched ‘uniform condition’, the configurations of colors were random on each trial. After learning occurred in the patterned condition, participants were able to remember twice as many colors as those in the uniform condition. The impact of this kind of result on fixed-capacity accounts of working memory rests on whether the chunks used to remember more colors retained instance-specific perceptual information of each display, or were abstractions which required decoding from long-term memory in order to respond at test. If chunking arises from pointers to long-term memory, this data is consistent with a single fixed capacity in terms of chunks (Cowan, 2001). However, if chunks are richer and more distributed, such learning reflects a clear increase in capacity, as participants would be holding more perceptual information in the chunked case.

Huang and Awh (2018) provided the most relevant test of this issue. Following up on Brady et al. (2009), they argued that chunk learning in this paradigm was in fact content-free; and thus suggested that fixed capacity models of working memory were accurate. They presented two experiments, one involving color-pairs as in Brady et al. (2009), and one with letter pairs, either forming well-known words or not. They found that reaction times for choosing targets at test were consistently longer for chunked displays as compared with displays with no opportunities for chunking, an effect which increased as learning increased in the color-pairs study. They suggested that this provided evidence for a content-free account, based on the indirect inference that ‘decoding’ from long-term memory might be slower than accessing items actively stored in working memory. However, this approach is quite indirect and does not address the core issue of what is stored (e.g., Figure 1). In order to pull apart the content-free position and the distributed, hierarchical view, it is necessary to test chunks in which long-term representations diverge somewhat from the particular instantiations on a given trial. While this is difficult for verbal stimuli, in the visual domain stimuli often vary continuously. In a case where you have formed a chunk of the concept “egg”, these views can be teased apart by asking whether you still know the exact yellow/orange of the yolk on a particular trial. If chunks are indeed content-free, then such instance-specific information should be lost when performance reflects the use of learned-chunks – i.e., there should be a trade-off between memory for information that is captured by regularities and memory for information that varies from trial to trial. If, however, chunks act as hooks onto instance information, with both activated as a hierarchical memory representation, then the varying instance information should still be available at test. In fact, if chunks actually facilitate the encoding and retrieval of the lower-level information, there would be no trade-off between instance specific information and chunk learning, and memory should actually improve for low-level details as chunks are formed.

In the current work, we provide the first direct test of this critical question about chunking. Experiment 1 tests memory for colored objects, and Experiment 2 tests memory for letters arranged in pairs. We manipulate chunking by introducing

regularities in the displays in Experiment 1 (objects frequently appear in the same color) and by using pre-learned chunks in Experiment 2 (letter pairs that form words). In order to tease apart content-free vs. distributed accounts, both experiments include instance-specific information that varies from trial to trial: In Experiment 1, the luminance of the color varies, and in Experiment 2, letters appear in different fonts unrelated to the long-term chunks. Thus, memory performance on instance-specific information will uniquely tell apart content-free accounts and the distributed accounts of chunking.

Experiment 1: Learning chunks

Participants had to remember 3 colored objects on each trial. Items varied in color, and their hues were predictable but not deterministic across trials (e.g., the backpack was red 80% of the time), allowing participants to learn chunks (e.g., ‘red backpack’). Items also varied in luminance in an unpredictable way across trials. Most trials probed for what object and hue was in a given location (“regular tests”). However, on a subset of trials we probed this luminance value rather than the color hue of the object (“detailed tests”), allowing us to examine how detailed trial-specific information is affected by chunking.

Methods

Participants. Our sample consisted of 75 US-based Prolific users between the ages of 18-45 who reported normal or corrected-to-normal vision. 15 additional participants were excluded based on a priori exclusion criteria. In particular, we excluded individual trials with RTs less than 200ms or greater than 5000ms, and excluded participants if they had more than 10% of trials excluded. We also excluded participants with poor accuracy. Since the test trials that focus on trial-specific information (luminance; “detailed tests”) would be expected to result in poor performance according to the content-free theory, we excluded participants based on accuracy only in the regular, hue tests, not detailed tests (<60%, chosen based on pilot data).

Stimuli and Procedure. We had 6 real-world objects, picked such that they had no single representative color (backpack; stapler; pot; dresser; garbage can; chair). They could each appear in 6 possible colors, spaced equally along a color wheel in CIELAB color space, and corresponding to roughly red; yellow; green; teal; blue; purple. For each participant, each of the objects was assigned a ‘high probability color’ which it appeared in 80% of the time. The other 20% of the time the object would appear in a different, randomly chosen color (see Fig. 2). To allow probing of trial-specific information, our main question of interest, the items varied in luminance as well as color: thus, even if the backpack was generally red and was red on the current trial, it could range from dark red to light red. Luminance values were sampled from a normal distribution in L of LAB space with a standard deviation of 10 units, with the exception of the fact that on



Figure 2. Participants had to remember the color and object at each of 3 locations. Colors and objects were paired reliably (80%) but not deterministically for each participant.

detailed probe trials, L was always set to exactly -10 or +10 from the mean, so that the shown luminance and foil luminance were always 20 units apart (2 stan. dev.) in L.

On each trial, participants saw 3 colored objects appear in separate spatial locations for 1500ms and then after a 700ms delay were probed on one of the items via a 4-alternative forced-choice (4-AFC). Most (75%) 4-AFC trials consisted of a 2x2 with the target object and another object that had not been present on the display, both presented in the color the target object had been vs. another unique color. However, 25% of trials instead consisted of a detailed color test: rather than a unique color, the alternative color was a different luminance of the same color (e.g., a dark vs. light yellow; see Figure 3). Participants did 5 blocks of 60 trials each, and we analyzed data separately by block to understand how learning over the course of the experiment affected memory for the objects and their colors, in both regular and detailed tests.

Results and Discussion

We found the expected performance improvement with block in the regular tests ($F(4,296)=7.81, p < 0.0001$; Fig. 3). This is consistent with participants forming and using chunks.

Note that in theory this increase could also arise due to a reliance solely on ‘guessing’, rather than either content-free or distributed chunking, to respond about the color: participants could simply remember which object was in each location while ignoring color, and then guess the probed objects’ color based solely on long-term memory (what Brady et al. 2009 called ‘post-perceptual inference’ in a similar paradigm). To examine this, we looked at the trials where items appeared in low probability colors. A solely guessing-based account predicts that performance should drop reliably on such trials as participants learn the chunks, even to below chance levels, as they treat all trials like high-probability pair trials. Because there were relatively few such trials, we compared the first and second half of the experiment. We found performance on low probability trials was similar in the halves (72.1% vs. 73.3%; $t(74)=0.54, p=0.590, d_z=0.06$), inconsistent with this guessing strategy. This suggests participants are generally aware of not only which object was in each location but also whether or not it




Figure 3. Participants improved at both regular and detailed memory as they learned the chunks.

is in its usual color, consistent with both chunking-based accounts.

Our major question was whether the improvement with learning — as people learn to encode the objects as ‘chunks’, like ‘red backpack’ or ‘yellow stapler’ — resulted in a cost to their detailed memory (for *what* yellow was present), as predicted by the content-free pointer account (e.g., Huang & Awh, 2018); or whether performance stayed stable or even increased in the detailed tests, as predicted by a distributed memory account. We found evidence favoring the distributed account. Performance in the detailed tests significantly increased over block, ($F(4,296)=4.78$, $p=0.0009$). Furthermore, detailed performance on only trials where chunking was most likely to have taken place, those (1) from the second half of the experiment, (2) where the item had appeared in its most frequent color, and (2) where the correct object was chosen (regardless of the color chosen), we found that selection of the correct luminance was far above chance (72.3% vs. 50% chance; $t(74)=15.56$, $p<0.001$, $d_z=1.80$). Within individual subjects, a strong correlation was also found between performance on the regular tests and performance on the detailed tests, across all blocks ($r(74)=0.68$, $p<.001$), further demonstrating that chunks benefit performance on both types of tests, rather than inducing a trade-off.

Considering only the regular tests, we also found, in contrast to Huang and Awh (2018), no difference in RT between tests in which the target was a high-probability pairing and tests in which the target was a novel pairing of object and color, even in the second half of the study ($t(74)=0.02$, $p=0.98$, $d_z=0.00$).

Experiment 2: Pre-learned chunks

Although learning was present in Experiment 1, performance did not increase dramatically from block 1 to block 5. It is therefore possible that the presence of instance-specific information in Experiment 1 was a result of the chunks of

object-color pairings not being fully learned. To test this idea, Experiment 2 relied on well-known words as chunks, probing memory for letters as well as the fonts the letters appeared in and manipulating whether words could be used as chunks to store the letters. The structure of the experiment was similar to Experiment 1, examining whether trial-specific information is lost when chunking via a 4-AFC task that combines trial-specific information with chunkable information. The prediction of the content-free account is that when participants use words to bolster memory performance (e.g., “TO”, “WE”), they store them simply as pointers to a long-term memory, losing instance-specific information (font type, which varies across trials). By contrast, the distributed memory account predicts that while chunking should improve memory for the object, it should come at no cost — and possibly a benefit — to memory for trial-specific details.

Methods

The participant number, exclusion rules and analyses were preregistered: <https://aspredicted.org/blind.php?x=cp4k8d>

Participants. Consistent with our preregistration, our final sample consisted of $N=75$ participants. They were US-based Prolific users between the ages of 18-45 who reported normal or corrected-to-normal vision. Eleven additional participants were excluded according to the preregistered criteria, which were similar to Experiment 1 (performance below chance overall; or more than 10% of trials with an $RT>5s$ or $<200ms$).

Stimuli and Procedure. Participants performed 160 trials of a working memory experiment. On each trial, participants saw an encoding display for 2000ms, with 4 spatially-separated sets of 2 letters and each letter randomly assigned one of two fonts (bold or italic). This was followed by a 700ms delay, and then they were probed on a single randomly chosen letter. The test was always a 4-AFC test where they had to indicate which letter and font was present by clicking the appropriate options (see Figure 4). The foil letter was always a letter not presented on the display at all, and both the target letter and foil letter were presented as options in both of the possible fonts. Participants clicked to choose what letter they had seen. The critical manipulation was that on 50% of trials, the letters were arranged so that each pair made




Figure 4. Participants were probed on the letters and their font in a 4-AFC task; the letters sometimes formed word chunks and sometimes were randomly paired.

up a two-letter word (WE; TO; AS; IF; UP; MY), allowing chunking to facilitate memory for letter identity. On the other trials, the 8 letters were chosen randomly from the same set of 12 letters.

Results and Discussion

When the items formed 2 letter words participants were more accurate at knowing which letter was in the cued location (regardless of font): $t(74)=12.16$, $p<0.001$, $d_z=1.40$; see Fig. 5. This suggests participants did indeed form higher-level ‘chunks’ of the letters to facilitate memory.

Our critical question was what happens to memory for the font when the items are chunked. The content-free account predicts that font memory is worse for chunked items; whereas the distributed memory account suggests memory is better. The data were consistent with the distributed memory account: considering font memory alone, independent of letter, participants were more, not less accurate in reporting the font of letters that were part of words ($t(74)=5.21$, $p<0.001$, $d_z=0.60$) even though the word chunks provided no information about the font the letter had appeared in. This benefit for font memory was also true only on trials where the letter they choose was correct: when reporting the correct letter, participants were more likely to report the correct font on chunkable trials than randomly-paired letter trials ($t(74)=4.28$, $p<0.001$, $d_z=0.49$).

Finally, we asked about ‘boundness’ — how dependent knowing the style was on knowing the letter. The content-free account predicts that chunking letter identity decouples letter identity from font, reducing boundness. Consistent with our preregistration, we operationalized boundedness as described by Brady et al. (2013). In short, the boundedness score captures how much more likely you are to know the font given you know the letter (vs. not), adjusting for overall percent correct at both tasks. This adjustment is critical because the difficulty of the two tasks is in many ways arbitrary: by making the letters or fonts more vs. less distinct, the forced-choice task could be made more or less difficult for either. Thus, to assess their dependence on each other, performance in both must be taken into account.

We found that the boundedness scores for chunked vs. random displays were statistically indistinguishable (23.5% vs. 22.7%; $t(74)=-0.20$, $p=0.840$, $d_z=0.02$), and both were greater than 0 ($t(74)=7.08$, $p<0.001$, $d_z=0.82$; $t(74)=7.67$, $p<0.001$, $d_z=0.89$). Thus, participants were more likely to know the font of letters they knew than ones they didn’t know, and this did not decrease with chunking.

Finally, contrary to the result found in a similar experiment in Huang and Awh (2018), we did not find a difference in reaction time between word and nonword displays ($t(74)=0.74$, $p=0.462$, $d_z=0.09$).

Overall, we find a benefit, rather than a cost, to trial-specific font memory when participants could encode letters as part of chunks (but see Zimmer and Fischer, 2020), in which chunking displays did not result in better memory for character type). This was true even though there were four




Figure 5. Chunking improves both letter knowledge and trial-specific font knowledge.

‘chunks’ of letters on the display, leaving — under many models (Cowan, 2001) — no extra storage capacity for independently storing additional font information in separate chunks. This suggests that the use of chunks enhanced, rather than removed, access to perceptual details.

General Discussion

Across two experiments, we tested the nature of memory representations that are present when chunking is used to scaffold performance in a working memory task. In Experiment 1, performance improved with increasing exposure to the high-probability object-color pairings. In Experiment 2, when letters formed words (pre-existing chunks), more individual letters were remembered from the display. Our key question for these studies was whether chunking in these tasks reflects the use of “content-free” chunks with no retention of instance-specific details, or whether the use of chunking does not imply losing such instance-specific detail, as in the proposed distributed memory account. We find support for the distributed memory account. In Experiment 1, the luminance of the objects’ colors varied throughout the trials. On 25% of tests, the foil color at test was the same as the target color but with a different luminance. On such tests, decoding from long-term memory would not help decide between the foil and target, since items were consistent only in their hue, not their luminance. Thus, the content-free view would predict a trade-off between regular and detailed tests as chunking increases. However, Experiment 1 showed the opposite result. Using chunks to scaffold performance on the regular tests actually improved performance on the detailed tests too, supporting the hypothesis that chunking facilitates the encoding and retrieval of instance-specific information, and rejecting the content-free view. The strong correlation between performance on regular and detailed tests within individual

participants also provided support against the trade-off predicted by the content-free view.

It is possible that the presence of instance-specific information in Experiment 1 was as a result of the chunks of object-color pairings not being fully learned. To account for this possibility, Experiment 2 relied on words as pre-existing, well learned chunks. The letter font varied trial to trial, and in cases where subjects could use words as chunks to improve memory, the content-free view would predict that they would consequently lose the instance-specific information of font type, which could not be decoded from long-term memory. Again, the opposite result was observed: participants were more accurate in reporting fonts when the letters were chunked into words. An analysis of the extent to which font was 'bound' with letter identity revealed that storage of the font detail was integrated with letter identity. These results support a view of working memory representations in which information is distributed across hierarchical levels; with learned chunks at one level facilitating efficient encoding of the memory display whilst maintaining access to instance-specific information at lower levels.

Chunk decoding and reaction times

Previous evidence for the content-free view came from differences in reaction times across two experiments of Huang and Awh (2018). They found that subjects took longer to respond on trials in which performance was boosted by chunking, inferring that this reflected the "extra step" of decoding the chunk from long-term memory. However, the experiments reported here failed to replicate this result. No reaction time difference was found in Experiment 2 between the words and non-words conditions. Similarly, no difference was found in Experiment 1 between high and low probability test targets. It is unclear what to make of these reaction time data. Reaction time results are likely influenced by multiple factors other than any potential 'decoding' time that may or may not be present. For example, both the strength of participants' memory as well as the complexity of the memory probe that must be processed influence reaction time (e.g., Baddeley & Ecob, 1973), so Huang and Awh's previous results may have simply reflected a confound of one of these factors. Generally speaking, as an indirect measure, reaction time does not provide strong evidence of chunks being content-free. By contrast, the data from Experiment 1 and 2 test the nature of the chunked representations directly and find them to contain rich instance-specific content.

Relation to other theories of chunking

The view of chunks as operating within levels of a hierarchical, distributed memory representation proposed here is supported by the "template theory" of chunking within chess expertise (Gobet & Simon, 1996). A template's "core" information consists of details necessary for activation of the template, i.e., details which are consistent across different instantiations and which discriminate the template from other templates. The template's "slots", on the other hand, are capable of incorporating new information by virtue of the

variability in information that can occupy those aspects of the template's structure across instantiations. These slots can thus accommodate, and improve memory for, instance-specific information, while simultaneously using learned representations (templates and chunks) to efficiently compress perceptual input. The difference between these views is that in the distributed, hierarchical view, the core and variable information may occupy the same aspect of the overall structure. The static and variable aspects of a chess template are separate pieces which can be simultaneously incorporated into a template and both encoded into memory. In chunks within the context of visual working memory, the regular, "core" information of learned abstractions such as "egg" will likely consist of generic shades of colors that the chunk is normally associated with, perhaps representing an integration across instances of chunk activation in perceptual contexts. The variable, instance-specific information in this case will also be a color shade of the egg-yolk – the specific shade of the current instantiation. Thus, the visual working memory case uniquely requires a hierarchical approach, with the pre-learned generic information at one level of the hierarchy and the instance-specific color information at a lower-level. Understanding how single pieces of information such as 'egg-yolk color' can be represented differently at multiple levels of the memory representation, and relatedly how chunks can form over continuous information, represent novel and potentially important areas of future research.

Working memory capacity

The proposed distributed view extends the historical view of chunking (Miller, 1956; Cowan, 2001) in ways which allow for a better understanding of the role of chunks in working memory and suggest a major rethinking of the nature of working memory capacity. Standard views of chunking generally suggest that while the number of items that can be remembered is increased by chunking, there is no genuine increase in working memory 'capacity' per se when chunks are used. For example, Cowan, Chen and Rouder (2004) argued that although teaching word-pairs to participants improved their recall of lists that included such pairs, they none-the-less had 'constant capacity'. This is because according to their model of working memory, the pairs were stored as content-free chunk pointers. Thus, both unpaired individual words and learned word pairs each took a single 'slot' in memory, and the same number of slots were used in all conditions.

By contrast, the data presented here suggest that when a chunk is employed in order to encode a portion of a memory display, this allows more of the lower-level information to be captured as well. The chunk, as an encoding "hook", facilitates the encoding and retrieval of this lower-level information into working memory, giving rise to a hierarchical memory representation (e.g., Brady et al. 2011). In this way, a single chunk allows for the encoding of a conjunction of features within the display, and the working memory representation consists of both the higher-level chunked "hook" and the lower-level information it is

encoding. Thus, our work suggests that working memory capacity is not exhaustively described at the level of how many chunks are encoded. Rather, capacity should be considered in the context of the full hierarchy. Under such a view, Cowan, Chen and Rouder (2004) — as well as many others — critically understate the improvement brought about by chunks: even if participants store the same number of ‘pointers’ in both learned and unlearned conditions, participants do not, in fact, store the same amount of *information actively* in working memory in both conditions, as chunks function to facilitate perceptual encoding as well as being stored as pointers themselves.

The question of whether chunks within visual working memory will *always* encode current-instance perceptual details is not answerable with the data presented here. It is possible that such details are available only because they are tested in the task and so known to be relevant. The ratio of regular to detailed tests in Experiment 1 (3:1) was designed to prioritize attention to and thus learning of the generic color-object pairings, but of course having so many detailed tests will quickly orient subjects to luminance information as well. Thus, the current work shows only that it is not a requirement for chunks to be content-free: in at least some cases, chunking can result in an increase in access to instance-specific information. Investigating when this does or does not occur will be a question for future research.

Conclusion

We showed that chunks allow for a more efficient representation of perceptual details in working memory. Access to perceptual details is not lost when items are chunked, as is commonly assumed, but is improved by chunking. This is consistent with an account of chunks as distributed and hierarchical representations, with perceptual information at a lower level of the hierarchy accessed via the chunks; and inconsistent with views that see working memory capacity as described largely by how many chunks can be encoded (Cowan, 2001).

References

Baddeley, A. D., & Ecob, R. J. (1973). Reaction time and short-term memory: Implications of repetition effects for the high-speed exhaustive scan hypothesis. *The Quarterly Journal of Experimental Psychology*, 25(2), 229-240.

Brady, T. F., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, 138(4), 487.

Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of vision*, 11(5), 4-4.

Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2013). Real-world objects are not represented as bound units: independent forgetting of different object details from

visual memory. *Journal of Experimental Psychology: General*, 142(3), 791.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and brain sciences*, 24(1), 87-114.

Cowan, N. (2005). *Working memory capacity*. Psychology Press.

Cowan, N., Chen, Z., & Rouder, J. N. (2004). Constant capacity in an immediate serial-recall task: A logical sequel to Miller (1956). *Psychological science*, 15(9), 634-640.

De Groot, A. D., Gobet, F., & Jongman, R. W. (1996). *Perception and memory in chess: Studies in the heuristics of the professional eye*. Van Gorcum & Co.

Gobet, F., & Charness, N. (2018). Expertise in chess.

Gobet, F., & Simon, H. A. (1996). Templates in chess memory: A mechanism for recalling several boards. *Cognitive psychology*, 31(1), 1-40.

Feigenbaum, E. A., & Simon, H. A. (1962, March). Generalization of an elementary perceiving and memorizing machine. In *IFIP Congress* (pp. 401-406).

Fougnie, D., & Alvarez, G. A. (2011). Object features fail independently in visual working memory: Evidence for a probabilistic feature-store model. *Journal of vision*, 11(12), 3-3.

Huang, L., & Awh, E. (2018). Chunking in working memory via content-free labels. *Scientific reports*, 8(1), 1-10.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2), 81.

Ngiam, W. X., Brissenden, J. A., & Awh, E. (2019). “Memory compression” effects in visual working memory are contingent on explicit long-term memory. *Journal of Experimental Psychology: General*, 148(8), 1373.

Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, 131(1), 48.

Zimmer, H. D., & Fischer, B. (2020). Visual working memory of Chinese characters and expertise: the expert’s memory advantage is based on long-term knowledge of visual word forms. *Frontiers in psychology*, 11, 516.