

Noisy and hierarchical visual memory across timescales

Timothy F. Brady  , Maria M. Robinson  & Jamal R. Williams

Abstract

Both in everyday life and in memory research, people tend to think that items are ‘held’ in mind, in the same way that a real-world object can be held in one’s hand. Inspired by this metaphor, traditional work on visual working memory and visual long-term memory focuses on understanding how many objects are remembered or forgotten, or held or lost, in particular circumstances. By contrast, newer computational and empirical work on visual memory focuses on the role of noise in memory representations – in which memories are thought to vary continually in ‘strength’ or ‘precision’ – as well as the role of the visual hierarchy and priors in structuring memory. In this Review, we merge these contemporary theories and evidence. We describe how fundamentally noisy memory representations are instantiated at different levels of the visual hierarchy and support both visual working memory and long-term memory. We also discuss how thinking of memory in this way can direct further research and illuminate the nature of cognitive function more broadly.

Sections

Introduction

Visual memory across timescales

Memories are noisy and vary in strength

Capacity limits

Influence of knowledge on visual memory

Summary and future directions

Introduction

Consider what your dinner plate looked like last night, including the colour, size and positions of food on the plate. When you try to bring these images to mind, you are trying to retrieve the details of a particular visual experience, using visual memory. Visual memory refers to information that can be successfully accessed from past visual experiences, and the structure of that information.

Visual memory is of broad interest because of its impressive capabilities. These can be appreciated in demonstrations of the method of loci, widely used to turn arbitrary words or playing cards into rich visual images by the best memorizers in the world¹. Visual memory also provides key insights into the nature of memory representations. Traditional memory research has often focused on identifying separate memory systems and on establishing a taxonomy of memory systems and their functions^{2,3} and has focused less on examining the nature of stored memory representations. By contrast, research on visual memory places a major emphasis on the structure and format of the representations of objects and scenes in memory and how properties of these representations impose limits on what information can be remembered accurately^{4,5}. Because the perceptual representations underlying visual processing are (comparatively) well understood relative to many other domains^{6–8}, research at the intersection between visual perception and memory can provide unique insights into memory processes by building on what is known about perceptual representations and examining how they change when held in memory.

A key insight is that visual memory representations are noisy – people can never remember perfectly accurately and the accuracy of memory tends to vary in a continuous manner. This insight derives in part from the emphasis in visual memory research on visual features that vary continuously, such as colour and orientation, and the metrics of performance used to assess memory of these features. However, the finding that visual memory representations are noisy conflicts with the tendency to use a physical metaphor to describe and understand memory. People tend to think of an object they are trying to remember as either in mind or not in mind, and to talk about items as being ‘held’ in mind to describe working memory, just as a real object can be held in the hand⁹. This metaphor often serves as a core mental model for how memory has traditionally been thought of: as all-or-none, discrete and operating over entire objects or chunks of information.

Given their continuous and noisy nature, however, the idea that memory representations are like physical objects that can be held or defined as simply existing or not existing, is not a useful metaphor. Thus, instead of using this mental model for memory, we turn to an analogy from the nature of processing in the visual system (Fig. 1). In particular, the visual system is usually conceived of in terms of population codes, with many neurons representing a given feature and a given location – and each doing so with substantial noise¹⁰. Furthermore, the visual system contains a hierarchy of representations, in which representations are richer and more complex as one moves from primary visual cortex up to more anterior visual regions¹¹.

In this Review, we detail how visual memory can be conceived of as a noisy and hierarchical system. We begin by reviewing work that suggests that representations for objects are strongly related across visual working memory and visual long-term memory, and in both cases are best conceived of as noisy and variable in strength. We then describe how such noisy memories are stored in terms of hierarchical memory representations. Finally, we discuss how these ideas lead to a reconceptualization of the limits of visual memory at different timescales, and the role of visual memory in other cognitive capacities.

Visual memory across timescales

We begin with the question of whether it makes sense to consider visual memory as a single construct, or whether the more traditional view, in which long-term memories are considered fundamentally distinct from working memories, is most appropriate. In traditional memory research, visual working memory and visual long-term memory are often studied separately, using different paradigms, by different researcher groups, and are often thought of as distinct systems. Although we acknowledge that a systems taxonomy can have heuristic value, this view is grounded in theories that posit that visual working memory and visual long-term memory differ in their mental representations and intrinsic structure. By contrast, contemporary evidence suggests that visual working memory and visual long-term memory share mental representations, even though varying task demands can lead to the differential recruitment of additional cognitive processes that operate over those representations.

Mainstream theories of memory often postulate that visual working and long-term memory representations differ from each other in key ways. One reason is distinctive neural correlates. For example, neuropsychological research suggests a distinction between visual long-term and working memory based on purported evidence of selective and distinct impairments to visual long-term and visual working memory^{12,13}. It has also been suggested that short-term memory reflects transitory electrical activity, whereas long-term memory reflects durable neurochemical changes in the brain¹⁴. Broadly in line with this view, some evidence indicates that there are distinct neural signatures associated with visual working memory and long-term memory, with the sustained and active firing of neurons involved particularly in working memory maintenance^{15,16} but not in long-term memory. In humans, such working-memory-specific activity can be tracked with electrophysiology, such as by the contralateral delay activity^{17–20}, a robust component thought to track the amount of information ‘actively’ stored in visual working memory²¹ that disappears when items are available in long-term memory²². Similarly, information in working memory can be observed in parietal and frontal region activity⁹ and decoded from ongoing activity during the maintenance period using neuroimaging^{23–25}.

However, these neural coding differences do not necessarily imply differences in the fundamental representations or intrinsic structure underlying working memory and long-term memory. For example, the neuropsychology work must be interpreted with caution based on the observation that participants from clinical populations suffer from a wide range of processing deficits that make claims of selective deficits to a specific cognitive system difficult²⁶. Similarly, evidence of distinct neural correlates for working memory, involving active maintenance processes, does not necessarily indicate that distinct representations are used in working memory compared to long-term memory. Instead, the distinction between typical working memory and long-term memory tasks can reflect the amount of activation of a particular representation^{27–30} or the engagement of attentional maintenance processes to maintain such activation, with the representations themselves being fundamentally the same. For example, the classic modal model of memory posited that short-term or working memory exists as a buffer that maintains information in a readily accessible state^{31,32}. It is therefore often posited that working memory is capacity-limited, such that only a small amount of information can be maintained in an active state at once, perhaps owing to limitations in attentional maintenance processes³³. Thus, rather than requiring two systems, the neural evidence is consistent with the view that representations in working

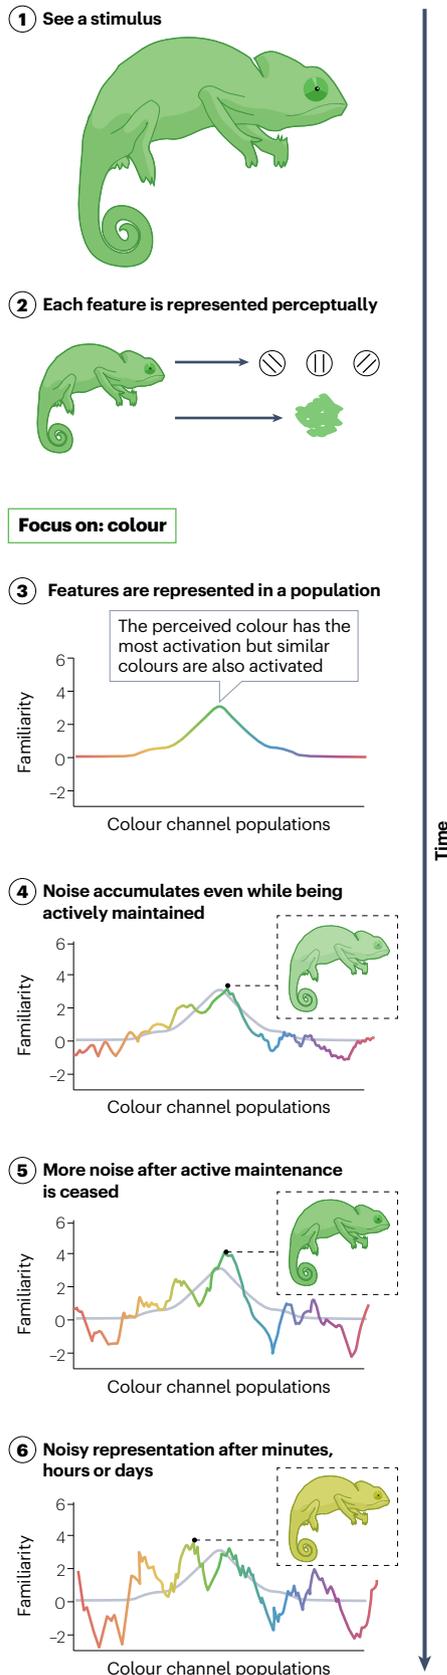


Fig. 1 | Visual memory at different timescales. Features of a visual object (1) are extracted and processed, eliciting a noisy population of activity over feature values, such as edges and colour (2). In this case, the colour of the object is relatively arbitrary. Perception of the relevant colour is noisy, and multiple colours near the colour shown (green) will all feel somewhat familiar after encoding the item (3). Actively focusing on the information in mind using attention can slow, but not stop, the accumulation of noise in that item's representation (noisy representation in colour, original distribution in grey) (4). Sampling the same memory representation at different time points reveals that the accumulated noise corrupts and alters the qualities of the originally encoded information, resulting in a noisy and imprecise sense of the colour of the original object ((5) and (6)).

memory and long-term memory are fundamentally the same, but that long-term memory is mostly the passive use of these representations and that working memory is a form of attention-demanding, activated memory, involving control systems that manipulate and maintain the same underlying memory representations. According to this account, once activated in working memory, a subset of information related to pertinent representations can be continuously protected from noise accumulation by attentional processes, but the representations remain fundamentally similar to those used in long-term memory.

Behavioural evidence is consistent with this account. Despite a long history of evidence suggesting that long-term memory is often less precise than working memory (such as theories suggesting that only the gist survives in long-term memory, causing false memories³⁴) – newer evidence shows that there is no fundamental difference in the range of levels of precision or memory strength that are possible in visual-working and long-term memory. Distributions of memory errors obtained from continuous reproduction visual-working and long-term memory tasks³⁵ fit with just one parameter model that captures memory strength across tasks³⁶ (Fig. 2). This evidence suggests that as memory gets weaker in the two systems, it degrades in an identical fashion. Of course, typical working memory and long-term memory tasks can differ in key ways, even with the same representations underlying them. For example, online attentional selection might be more often necessary in working memory tasks because many items are often presented at once in such tasks, and there is often less spatiotemporal context in long-term memory tasks compared to working memory tasks, with many items all presented in the same spatial location being relevant. Thus, although long-term memories on average might be weaker than working memories in everyday-life situations, evidence suggests that this difference in strength does not necessarily reflect a fundamental difference in memory systems or their representations but rather in the way they tend to be used and investigated.

Thus, the evidence is consistent with a view that visual working memory is a form of activated long-term memory, with similar representations that differ in their degree of activation or in the processes being used to manipulate and maintain them. Seen through this lens, attention 'selects' and upweights goal-relevant information, and provides protection against the accumulation of noise across populations of neurons throughout the visual hierarchy³⁷, and memories protected in this way are working memories.

Prior work has criticized the unidimensional view of visual memory in part because the idea that memories are active in working memory (and less active in long-term memory) has not been consistently operationalized. For instance, activation has been used to describe both the strength of mental representations, as well as dynamics in

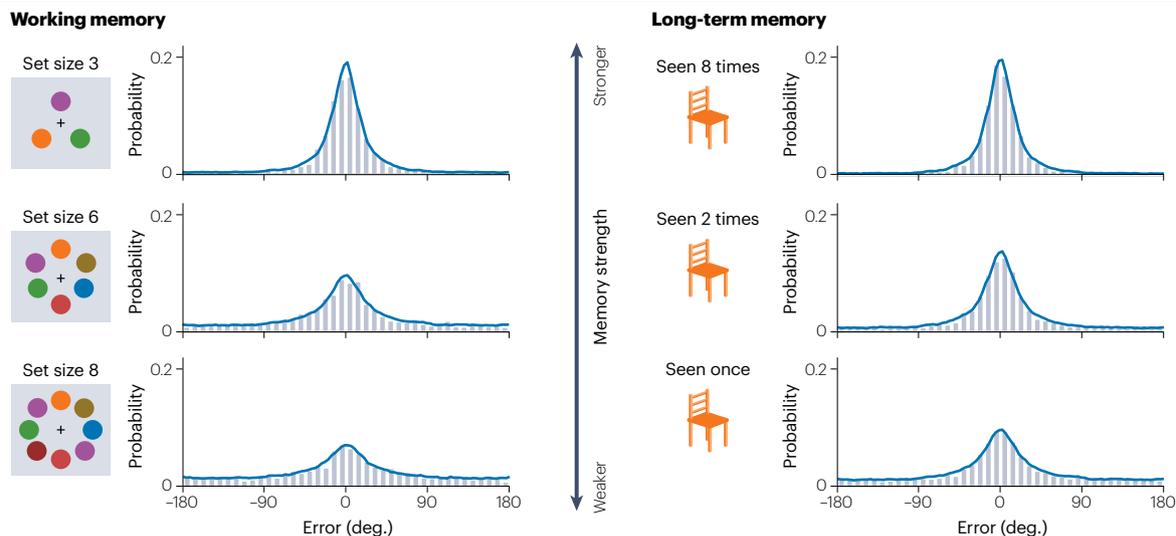


Fig. 2 | Memory strength across time. Distributions of memory errors (grey) obtained from continuous reproduction visual-working and long-term memory tasks³⁵ fitted with a model (blue line³⁶) that requires only a single parameter to capture memory strength. When the same stimulus features are used and are

probed in the same way, representations across the two timescales can be strong and precise or weak and noisy and degrade in an identical fashion as memory strength weakens.

neural processing. Furthermore, the conception that activation alone differs between working memory and long-term memory is generally insufficient to explain the differences across working memory and long-term memory tasks. For example, many theories treat activation levels as varying even among different long-term memories, and some work suggests that working memory can make use of inactive states^{38,39} in addition to active ones. This notion of ‘activation’ leads to theories arguing that up to a hundred or more recently seen items can be active in long-term memory⁴⁰, as compared to ‘inactive’ items that are seen equally often but are not relevant in the current context⁴¹. Reconciling all the notions of activation across both neural and cognitive meanings of the term and across both working memory and long-term memory remains a difficult and important task for unidimensional theories of memory, and one that has been the source of multiple proposals about the architecture of working memory^{33,42–45}.

In summary, within this unidimensional view of visual memory, visual working and visual long-term memory are not distinct systems: instead, both involve computations over the same types of memory representation. Importantly, the nature of these computations is constrained by the demands of the task, such that some visual memory tasks elicit processing constraints that others do not. Below, we develop this unidimensional framework by integrating findings from the visual working and long-term memory literature on the nature of visual memory representations. We then detail how differential demands can elicit processing bottlenecks in visual memory tasks. Finally, we discuss how the unidimensional view fits parsimoniously with more ecological theories that take into account the role of prior knowledge, stimulus complexity and meaning in visual memory across timescales.

Memories are noisy and vary in strength

A major dichotomy in the visual memory literature has been between views in which memories are conceptualized as all-or-none – with bound objects that are either remembered with complete accuracy or fully lost^{46,47} – and views in which memory representations are noisy,

hierarchically structured and distributed – with object features that degrade continuously and relatively independently^{48–50}. This question of how to construe the fundamental properties of visual working memory representations has dominated the literature on visual working memory for the past couple of decades, and physical analogies based on holding items in mind have long pushed researchers toward more discrete views of memory.

In the following sections, we describe behavioural, neural and computational evidence that visual memory representations are noisy, which is synonymous with the view that they vary in strength. We also discuss how noise accumulates over short and long timescales and how memory representations can vary at different levels of abstraction: from individual features to entire objects, as well as how the graded nature of memory representations emphasizes the role of decision-making in mainstream memory tasks.

Graded memory strength in working memory tasks

The question of whether memory is noisy or all-or-none has been most often examined in the domain of short-term, working memory tasks. All-or-none theories of memory in which items are either present (and perfectly represented) or absent entirely have traditionally dominated the analysis of visual working memory data from change detection tasks (Fig. 3a). In a typical change detection task, people are shown an array of simple stimuli that they have to remember after a brief delay. After this retention interval, people are shown a single stimulus and asked to make a discrete judgement regarding whether it changed or did not change compared to the stimulus shown in that location in the original memory array. The discrete nature of ‘change’ or ‘no change’ responses provide only a coarse measure of memory errors and often incorrectly leads to inferences such as that items are either present or absent in memory⁵¹.

By contrast, most contemporary work uses tasks that yield a more fine-grained measure of visual memory representations. Continuous reproduction tasks (Fig. 3b) provide a way of assessing gradations in memory representations and occur often in visual working memory.

Review article

In such tasks, instead of providing a binary response^{52–54}, people are asked to reproduce their memory of a probed feature or object using a continuous response wheel, which provides fine-grained information about both the magnitude and direction of memory errors. In these tasks, responses generate a distribution of memory errors in a single, continuous feature dimension (such as colour or orientation). For instance, in the change detection task, people are shown a memory array with simple stimuli and instructed to remember these stimuli and their locations over a brief retention interval. However, after the delay, people are shown a spatial probe along with a continuous wheel that shows all possible values in the continuous feature space. Accordingly, people are instructed to select the feature value they think most closely matches the value of the probed item. Thus, unlike change detection tasks, the continuous reproduction wheel provides a fine-grained measure of memory. Similar fine-grained performance metrics can

also be arrived at when using forced-choice recognition tests that vary the similarity of the seen and unseen items^{36,55}.

A first major insight from continuous reproduction tasks is that the distribution of errors in visual working and long-term memory tasks are incompatible with fully all-or-none theories of memory (Fig. 3c). These tasks reveal that the noise (variance) of memory errors increases with memory load⁵⁰ and delay^{36,49}, and decreases with more encoding time³⁶, as well as repeated study, and do so in a way that is not solely due to all-or-none failures³⁵. These results challenge an all-or-none interpretation because they indicate that simple task manipulations can expose the graded changes that exist in memory strength and accuracy.

A second major insight from tasks that allow continuous measurement of memory accuracy is that some task conditions introduce subtle memory biases, such that similar representations sometimes repulse away or attract towards one another^{56–61}. For instance, when people are

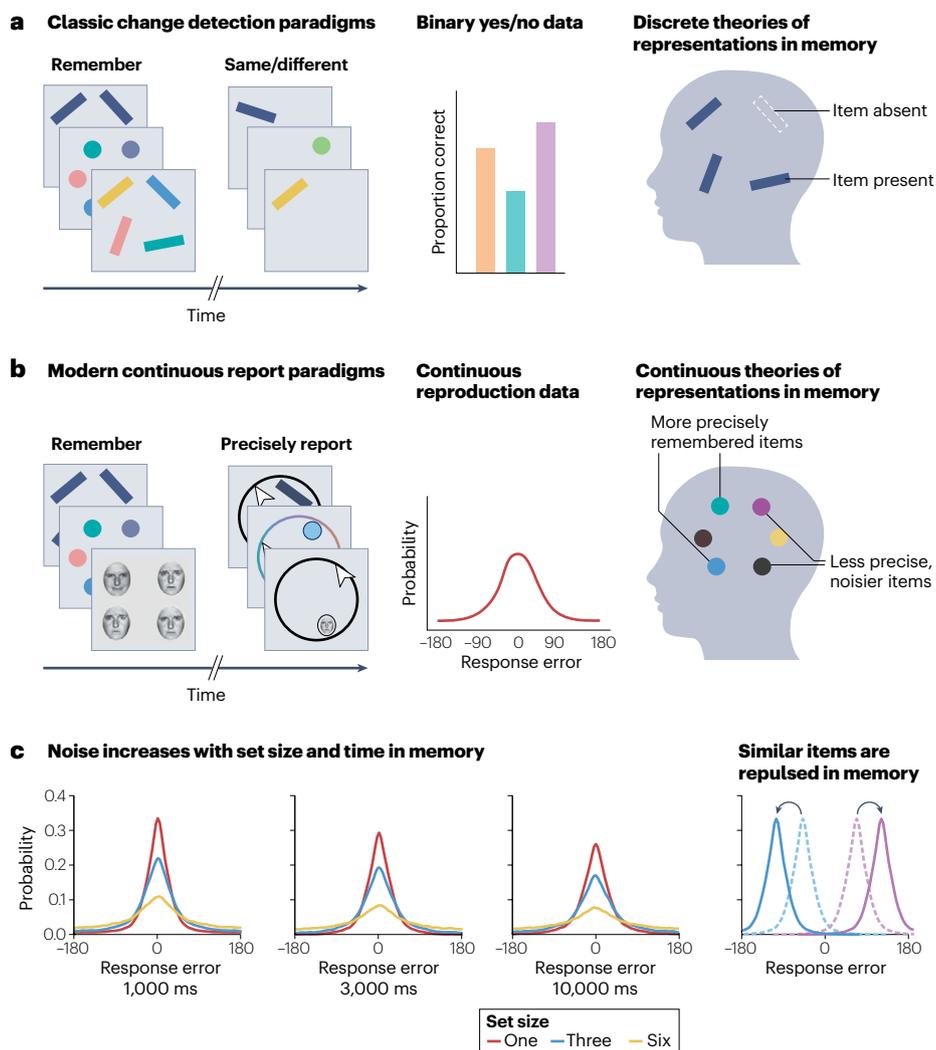


Fig. 3 | Categorical and continuous reproduction of memory. **a**, Early methods for measuring memory performance used binary tasks, such as old–new for visual working memory and change detection tasks for long-term memory. Data from binary tasks have usually been interpreted as providing support for discrete memory representations^{46,51}. **b**, Continuous reproduction tasks are memory versions of the psychophysics method of adjustment. These tasks provide a

fine-grained measure of the magnitude and direction of memory errors, which reveal that representations are noisy and are rarely completely lost in a discrete fashion. **c**, The major insights are that the precision of memory representations declines monotonically with memory load, declines with increasing delay, and that memory representations can be biased (for example, repulsed, or pushed away from one another) depending on the relationship between currently active items.

presented with many red items that slightly vary in hue, the representations of these items can be pulled or attracted towards their average hue. Alternatively, when people are presented with just a few similar red items, their representations can become more distinct and repel one another (Fig. 3c). These results imply that items are not represented

Box 1

The utility of random guesses

It is often proposed that memory performance includes many responses that are completely independent of the information held in working memory, with no relationship between the response and the encoded or remembered information. In such a view, quantifying ‘pure guesses’ is critical to understanding working memory^{68,69} and long-term memory^{214–216}. However, the utility of this construct is uncertain under the view that memory is both hierarchical and continuous in strength.

Consider a visual working memory task in which one sees several different apples presented on a display and must remember their sizes for a memory probe a few seconds later. The majority of evidence suggests that although noise will accumulate for each item, this noise cannot be characterized as instantaneous and infinite, as is required for an account that argues for complete forgetting (and thus pure guessing). Instead, most data suggest that information about every item is probably represented to some extent, even if a very large degree of noise is present^{36,51,65,66,70}. However, even if items are so noisy that they might functionally be completely gone from memory, continuous models do not conceive of this situation as a distinct state of ‘guessing’. Such models instead have the core assumption that the same populations that give rise to perceptual and memory experiences also give rise to the noise that is present in such memories, and thus the noise is never content-free (even when items are extremely poorly represented).

Even if, in contrast to this view⁶⁸, we accept that representations can be fully lost from memory, it is still unclear whether classifying a response to that item as a guess is useful. For example, people sometimes use the word ‘guess’ even when, without any specific information about a given item, one has substantial information about the items presented in a more general sense²¹⁷. In the example with a display of apples, one will probably know what size the items were in general on the display⁵⁷, will probably know that all the items present were apples, will have strong expectations about the expected size of an apple in general (versus other fruit such as a watermelon^{192,218}), and will know the expected size of each item on a display with just three apples on it. All of this information was acquired during the episode and informs the response about the size of the particular item, even if the individual representation of the actual size of a single apple is indistinguishable from noise at test, thus making the concept of guessing extremely incomplete as a description of both people’s behaviour and people’s internal representations. Thus, rather than a focus on whether participants are ‘guessing’, researchers should focus on capturing the true and complete contents of memory, taking into account the fact that memory is contextual, noisy, reconstructive in nature¹⁴⁰ and depends on integration across multiple levels of representational abstraction^{57,198}.

and recalled independently of each other but are represented in a way that causes subtle interactions between them. Both insights are broadly inconsistent with the typical physical analogy of holding discrete items in mind and therefore with all-or-none theories of memory representation.

Thus, all contemporary theories of visual working memory have moved beyond the concept of complete discreteness in memory, and broadly agree that information in memory is represented imperfectly and that this imprecision varies with a number of factors such as working memory load, encoding time, delay and several others. Most neural and cognitive theories broadly endorse the view that there is internal noise in perceptual, attentional and memory processes, as well as sources of external noise from the environment. These sources of noise can be affected by how attention is distributed during encoding⁶², memory load⁶³ and retention intervals⁶⁴. All of these sources influence the extent to which fluctuations in internal and external noise affect the fidelity of memory representations.

Whereas nearly all models agree that memories must be considered noisy and imprecise, an increasingly large number of theories of visual working memory now subscribe to the view that noise accumulation among items alone is sufficient to explain the limits on performance and the pattern of errors obtained in such tasks without any additional factor of all-or-none loss of items^{36,65,66}. However, some hybrid ‘mixture models’ include additional factors that postulate that continuously graded memories exist, but also that all-or-none loss of items can occur under some conditions^{67,68}. These models posit that some items that were seen and processed by the visual system while they were visually present are nonetheless completely unavailable and that observers have no remaining information about them, even less than a second later, such that when people are probed on those items, they are forced to make completely information-less guesses^{67,68} (Box 1). Although such hybrid models were initially influential because they were unique in providing an account of the distribution of errors in reproduction tasks⁶⁹, later work showed that continuous models can parsimoniously account for these distributions of errors just as well⁷⁰. For example, the finding that people tend to have little information about the fifth or sixth items when asked to report a set of six briefly shown items follows naturally from models that posit that noise accrues over items and time, without all-or-none failures⁷⁰. Thus, modern models without any added assumptions about all-or-none failures generally account for the data as well or better than such hybrid models^{36,65,66}. Moreover, hybrid models that attempt to predict how often pure, information-less guesses arise when an item is not in memory, across different set sizes (such as how many putative guesses there are when shown four versus six items)⁶⁹, generally fail to accurately predict the full pattern of errors as memory load increases⁷¹. Thus, modern hybrid models generally cannot sufficiently explain performance across set size, because they do not tend to provide any quantitative account of how often all-or-none losses are expected to occur across different conditions or set sizes, which was initially considered a major strength of hybrid models (for further discussion, see ref. 72). Overall, current findings about how visual working memory varies as a function of different variables, such as memory load and delay, can be parsimoniously explained with a unidimensional view according to which memory representations vary in noisiness, and do not require the additional assumptions about all-or-none coding that are necessary in hybrid models.

Noise across time

If working memory and long-term memory representations are fundamentally similar, one would expect models that focus on long-term

memory rather than working memory to also be focused on understanding memory noise and imprecision. However, one notable difference between traditional theories of visual working and long-term memory is how they conceptualize noisy memory representations. Theories of visual long-term memory typically conceive of noise as affecting the ‘strength’ of memory representations rather than the noisiness or precision of those memories. This construct of memory strength follows from a long tradition of signal-detection-theory approaches for measuring memory^{73,74}. Within the signal detection framework, previously seen and unseen items give rise to two different distributions of familiarity signals. When asked to make a memory judgement (such as which of two items you saw), familiarity signals from comparing memories to stimuli along different dimensions are collapsed into a unidimensional continuum of signal strength, with seen items having higher average signal strength than unseen items⁷⁵. Accordingly, the standardized distance between the seen and unseen distributions (traditionally referred to as the ‘noise’ and ‘signal + noise’ distributions and often captured by the signal detection measure d') is assumed to reflect the relative ‘strength’ of remembered items. Thus ‘memory strength’ is how much more familiar, on average, previously seen items seem compared to previously unseen items.

By contrast, most visual working memory models, as described in the previous section, conceive of noise in memory representations in terms of precision (such as how tight the errors are around the true value when asked to reproduce an item) rather than strength^{67,68,76}. The construct of visual memory precision follows from characteristics of the continuous reproduction task. In this task, the variance of the error distribution around the true item value is typically interpreted as a measure of how precisely a given feature or object was remembered⁶⁹.

The constructs of memory strength and precision ultimately rest on the same fundamental assumption: that memory representations are noisy and therefore vary in their match to the information that was actually seen. Additionally, the two constructs can be formally linked by signal detection theory with an entire population of signals³⁶ (Box 2). Importantly, how confident someone is in their memory tracks both the precision and strength of memory representations, which provides another empirical bridge between the two constructs. In the long-term memory domain, it is generally found that confidence tracks memory performance extremely well, and that such confidence judgements are generally unaffected by demand characteristics⁷⁷. Similarly, in working memory, confidence also tracks performance accurately even in continuous feature spaces^{51,78,79}, and this relationship can be naturally accounted for by models that relate precision^{65,70} to confidence or those that use signal-detection-based measures of strength to predict confidence³⁶.

Other measures that serve as an index of confidence – for example, asking participants to report the full range of feature values they think an object had rather than a single value^{79,80} – suggest that participants have an internal sense of uncertainty that tracks with memory performance in working memory tasks. This relationship is similar to the relationship between confidence and accuracy that is fundamental to measures of performance in long-term memory tasks. Taken together, there might be a direct mapping between people’s confidence judgements and a latent, continuous memory signal that is used on both visual working and long-term memory tasks. Confidence judgements and other judgements of uncertainty^{51,81} can be seen as providing convergent support for the view that precision and strength ultimately describe the same fundamental properties of visual memory representations: that visual memory representations are

noisy and vary continuously in both working memory and long-term memory tasks.

Noise across multiple features

The view that visual memory representations for single features (like colour or orientation) are noisy raises the question of how to think about items that consist of a conjunction of multiple features such as shape, orientation and colour. Although this ‘binding problem’⁸² has traditionally been framed in discrete terms (such as ‘are both features remembered or is one forgotten?’), in the context of noisy memory representations the question is whether features accrue noise independently or jointly or both. This question continues to be actively debated, but the latest evidence suggests that features seem to accrue noise largely independently⁸³. For instance, memory errors in a continuous reproduction task are largely uncorrelated across feature dimensions^{84,85}. Likewise, independent-feature resource models, which postulate that a separate set of resources are used to support memory for different features, such as colour and orientation, outperform shared resource models, in which colour and orientation compete for the same pool of resources, in a version of the change detection paradigm⁸⁶.

Evidence for independent noise accrual of features has also been reported for real-world objects. People commit more binding errors (swapping features such as colour across objects) when state (object configuration, like a full or empty mug of coffee) and exemplar (category, like mug) are jointly manipulated, suggesting that the two features can be forgotten or misbound separately from each other⁴⁸. This finding provides convergent support for the view that object features are stored at least somewhat independently and not as unified, all-or-none units that accrue noise holistically and are forgotten holistically. Similarly, in long-term memory tasks, different object features also seem to be represented at least to some degree independently⁸⁷.

Despite this evidence for independent noise accrual across stimulus dimensions, visual working memory studies tend to find a clear memory benefit when visual features are part of the same object⁸⁸. For example, it is easier to remember two features (such as colour and orientation) when they are bound to the same object, compared to separate objects⁸⁶. This benefit of objecthood might reflect the benefits of attending to a single as opposed to multiple spatial locations⁸⁹. However, classic work on visual attention also reveals evidence for benefits of object-based attention mechanisms, even when spatial distance is controlled⁹⁰ (for a critical analysis of object-based attention effects, see ref. 91). Collectively, the joint finding for independent noise accrual and spatial and object-based benefits for memory for features is consistent with the view that different features generally accumulate noise independently. However, spatial and object-based attentional mechanisms can introduce additional, correlated sources of noise accrual during encoding or modulate external and internal noise accrual across feature dimensions.

Decision-making under noise

A critical implication of the view that memories are always noisy and imprecise is that people must make probabilistic inferences based on noisy evidence to decide how to respond in memory tasks^{92,93}. Thus, unlike traditional all-or-none views, according to which people make memory judgements in a straightforward way by reading out whether an item is or is not in memory, within a continuous memory framework people have to use additional criteria for reporting whether they do or do not have enough information to endorse a particular response.

Box 2

Memory as a point or a population

Conceiving of memory representations as noisy helps to bridge the gap between cognitive and neural models of visual working and long-term memory. Much neural evidence suggests that cognitive representations are instantiated in distributed patterns of activity across neural populations and that activity across and within these populations is corrupted by noise^{70,219,220}. By contrast, many cognitive models of memory — especially long-term memory — tend to treat memory for an object as being effectively a single unitized familiarity signal, perhaps aided by other context-based sources of memory when determining whether an item has been seen before^{75,221}. In the past few years, work has suggested that these two views can be naturally reconciled in working memory and long-term memory by taking the signal-detection-based approach common in understanding familiarity and turning it into a population of familiarity signals modulated by psychophysical similarity³⁶. On the neural level, psychophysical similarity can arise from the tuning functions in feature-selective populations of neurons⁶⁶ and memory retrieval involves decoding these neural patterns of activations of stored features. Approaches like these therefore provide a common framework for thinking about noise across both cognitive decision models and neural models, bridging different levels of processing²²².

Conceptions of memory that are based on a population of signals for remembering a single colour are quite different from views in which what is stored is just a single point representation. An actively debated question deeply connected to the population coding view of visual memory²²³ is whether people truly represent perceptual and visual memories as point estimates or as a probability distribution over feature values²²⁴. This question is difficult to address because versions of point-estimate models can mimic models that postulate probabilistic cognitive representations²²⁵. Thus, addressing this question might depend critically on understanding the linking function between neural activity and cognitive representations. For instance, although population coding neural models are strongly compatible with the idea that people represent visual memories probabilistically (in early visual cortex), it is possible that this activity is inaccessible to higher-level processing and therefore cannot be used to make decisions in memory tasks. Nevertheless, people's subjective uncertainty in their memories tracks qualitative and quantitative properties of their memory errors^{79,80}, indicating that at a minimum people do have access to and read out information regarding how noisy their memory representations are, which suggests the possibility that visual memories are truly populations, not points.

Accordingly, researchers must always consider the decision-making processes that underpin memory tasks. This point and related ones have often been made previously in the context of recognition memory tasks^{78,94–96}, such as old–new and change detection tasks in long-term and working memory research, respectively. These tasks are frequently misused and misinterpreted because the decision process is not carefully considered⁷⁸. Decision-making processes are integral in such tasks because each stimulus elicits a noisy familiarity signal, so that people must set a decision criterion for responding whether an item is old or new, or whether a change occurred or did not occur⁸¹. Thus, it is critical to consider that variation in people's criterion setting can lead to large differences in measures like overall accuracy without reflecting any underlying change in the memory representation itself.

The study of decision-making processes in visual memory tasks is especially central in more applied settings, like eyewitness memory tasks in which people must select the face of the guilty suspect amongst a lineup of faces or report that the guilty suspect is not part of the lineup⁹⁷. Much work suggests that in such tasks people do not simply process each face independently, but might use higher-level decision-making strategies, such as discounting features that are common to all faces⁹⁸. The development of computational models to capture such strategies is extremely important to help to dictate how to construct lineups that increase the likelihood that the guilty suspect is correctly identified and to decrease the likelihood that an innocent suspect is incorrectly identified.

Decision models have been rigorously applied in many visual memory tasks beyond eyewitness memory^{99–102}. For instance, in one study researchers applied sequential sampling models to examine the decision processes that underpin performance in continuous

reproduction tasks¹⁰³. This modelling work provided insight into how memory precision relates to a noisy process of evidence accumulation in perception and memory that underpins memory-based decisions.

In sum, the view that visual memory representations are fundamentally noisy highlights that both short-term and long-term memory representations are continuous and share the same fundamental properties (Fig. 2). These assumptions apply to memory representations at different levels, from features to objects, and fits parsimoniously with neuro-computational population coding theories of perception and memory. Finally, the continuous representation view highlights that people must have ways of using noisy evidence to make memory-based judgements and actions, which highlights the fundamental role of decision-making processes in laboratory and real-world memory tasks.

Capacity limits

Given evidence that memory representations are continuous, a critical research goal has been to characterize limits in visual memory through a continuous framework. The working memory and long-term memory literatures have taken different approaches to this question. Whereas in long-term memory much work has accepted that no single fixed capacity limit can be found because the effect of interference between items and the role of retrieval cues will be complex and stimulus-dependent, the working memory literature has been much more focused on attempts to find a fixed resource limit that explains performance in simplified settings.

One major difference between visual working and long-term memory is often framed as differences in their putative capacities. Unlike visual long-term memory, which is thought to be virtually unlimited, visual working memory is thought to be extremely limited,

with performance dropping quickly as more information must be maintained^{54,104}. Although there are undoubtedly constraints on how much information can be actively maintained in working memory, especially given limits on attention or maintenance processes that work to reduce noise in task-relevant representations, most tasks designed to determine working memory limits also artificially constrain memory capacity by intentionally employing situations in which the most useful memory cues are limited in their availability (due to interference or cue overload). Despite their wide use and historical significance, such tasks produce a greater accumulation of memory-based noise than in tasks with reduced attentional demands, in which there is less interference and more diagnostic retrieval cues. As attentional mechanisms improve the quality of actively maintained information, these paradigms are informative about the capacity of attention and working memory maintenance processes, which is what they are designed to probe. However, they might not be very reflective of more genuine, real-world uses of working memory, which rarely occur under conditions designed to reflect solely the limits of attentional maintenance processes. Instead, many working memory tasks reveal primarily how accurately people can maintain information over short intervals only in quite unfavourable and unnatural circumstances. For example, typical visual working memory tasks involve the simultaneous flashing of many meaningless, single-feature items for less than a second⁵⁴, whereas typical visual long-term memory tasks involve the sequential presentation of meaningful items for several seconds each⁵⁵. As we explain in detail in the following sections, the contrast between these tasks makes it unsurprising that performance in the former would be worse than in the latter. Many purported differences between memory systems can instead be interpreted through a unidimensional view, in which different visual memory tasks place differential demands on attention and provide different amounts of task-relevant information that can be used to retrieve memories with higher fidelity.

Capacity in the laboratory versus the world

Differences in presentation format and stimuli across many working memory and long-term memory tasks introduce differential effects of spatial and temporal context that do not always reflect the true differences in everyday uses of shorter-term and longer-term visual memory. However, there are some differences between typical working memory and long-term memory paradigms that might reflect genuine differences in the real-world use of memory at short and long durations. For instance, spatiotemporal cues about an object – such as the spatial location in which it was shown and the context of other items around it – might be readily available in the environment in shorter-term but not longer-term memory tasks¹⁰⁵ because such cues tend to be available and stable only for short durations in the real-world. For example, while pouring tea, one's mug will probably stay in the same place, but over the course of a day it will probably change position quite often. These effects of context might differentially aid retrieval through external cues^{105,106} and introduce biases into memory representations^{58,60}.

However, other differences between laboratory-based working memory and long-term memory tasks are not necessarily typical of shorter-term and longer-term visual memory in the world. As an example, differences in encoding demands between typical working memory and long-term memory tasks are introduced if stimuli are presented simultaneously in several spatial positions or sequentially in the same position. For instance, a typical visual working memory task involves the simultaneous presentation of many objects, and people must distribute visual attention broadly among multiple items at once with

no priority difference between the items and little time for elaborative encoding. By contrast, sequential tasks in which objects are presented one at a time, like those typically used to assess long-term memory, might be more representative of real-world situations in which selective priority and elaborative encoding can take place. These and other differences in presentation format can drastically change estimates of working memory capacity for meaningful stimuli; people seem to have a higher capacity for meaningful objects presented serially¹⁰⁷ and under some conditions seem to have no upper bound in how many objects they can recognize after serial presentation¹⁰⁸. Indeed, in contrast to prominent claims^{52,54,109,110} that visual working memory tasks are not encoding-limited, but only limited by maintenance capacity, additional encoding time significantly enhances performance in working memory tasks even for objects with a single feature^{36,111,112}. Critically, in many realistic situations people might largely use working memory by sequentially focusing on a small number of items^{113,114}, implying that visual working memory tasks in which many items of equal relevance are simultaneously presented at once for a brief duration might be less characteristic of real-world demands. In summary, it is important to consider the role of such attentional limits and differential encoding demands in visual working and visual long-term memory tasks before attributing performance differences in these tasks to true differences in capacity or to core differences between distinct memory systems.

Visual working memory resources

In the working memory domain, many attempts have been made to formalize a single, limited capacity. This limited-capacity view was popularized by the famous report that people seem to maintain only about seven 'chunks' of information¹¹⁵. The idea that working memory capacity is set in terms of discrete chunks or items has dominated the traditional view of visual working memory capacity through the lens of 'slot' theories. According to these theories, the architecture of working memory is composed of a discrete number of slots that store objects composed of single or bound features and these slots are directly responsible for the limits of working memory capacity^{46,47,69,116–118}. This framework provides an intuitive way of thinking about capacity limits as a discrete number of items that can be held in memory. Despite this intuitive appeal, strong versions of this view also construe memory representations as all-or-none, which cannot account for the extensive evidence that memory representations vary in strength or noise.

However, although there is overwhelming evidence that memory representations are continuous, it has proved conceptually challenging to provide a rigorous characterization of working capacity in terms of resources^{119,120}. Attempts to measure and define a single resource limit often hold many of the factors that would be expected to influence performance in working memory tasks fixed (such as encoding time, delay, contextual cues, and/or presentation format). Then researchers manipulate and measure how many simultaneously presented simple items can be remembered. These tasks therefore investigate a highly limited subset of what would be required to truly understand what constrains performance in tasks in which attention is used to protect items in memory over short delays. It is therefore unclear whether any putatively fixed resource limits are meaningful in more real-world contexts or across even small variations in task. Within these confines, substantial computational modelling work has been devoted to developing concrete definitions of resources. For instance, some researchers define resources as 'neural real estate'^{76,121} and suggest that visual memory representations are distributed across populations of neurons in the visual cortex. In such models, the main source

of resource limits arises from the fact that a fixed firing rate is used to represent both a few and many items in memory (the basic concept of divisive normalization¹²²).

At the level of cognitive representations, researchers use principles from signal detection and sampling theory^{103,123,124} to define resources as a limited number of 'evidence samples'¹²⁵, which determine performance across variations in memory load. Although sampling theory makes more principled predictions than theories that simply postulate generic resources, the construct of samples is also poorly defined. It will therefore be important to define samples more rigorously, as in a new proposal that retrieval of representations from visual working memory reflect draws of samples from noisy spiking activity in a population of neurons⁷⁰. It will also be key to expand the understanding of resources beyond an understanding of how well people represent different numbers of simple objects shown briefly.

There has also been a resurgence of work using information-theory modelling approaches that has yielded insights into the properties of resource-limited memory systems. These approaches frame memory systems in terms of concepts like compression¹²⁶ or rate distortion theory¹²⁷. For instance, rate distortion theory, which aims to explain how a capacity-limited system should store information, can be coupled with prior knowledge and task relevance within deep neural network modelling¹²⁸ to simulate many fundamental aspects of people's memory errors, such as biases towards the ensemble of remembered items⁵⁷. Broadly, this modelling approach falls in line with resource rational models^{129,130} according to which people select computations that optimize outcomes but minimize the resource cost of implementing these computations. Such models have promise for understanding resources in a more general way, and across a broader set of tasks including visual working and long-term memory tasks.

As noted above, many of these computational modelling frameworks of capacity were developed using laboratory tasks that require short-term memorization of simple stimuli presented simultaneously. Such a simultaneous presentation format does not necessarily reflect how information is encoded over the short term in real-world scenarios, when each item might generally be encoded serially and one might often re-sample items rather than push the limits of working memory capacity^{113,131–133}. Indeed, estimates of visual working memory capacity when people encode meaningful (but not meaningless) items serially rather than simultaneously increase substantially, probably because this format provides time to connect online memory representations to existing knowledge¹⁰⁷. Even for simple stimuli, claims that encoding time was not a limit on performance^{52,54} have not held up in experiments using more psychophysical methods, which reveal that performance increases relatively smoothly with increasing encoding time^{36,111,112}. These findings underscore that there are many experimental levers that can, and should, be used to investigate and build ecological models of visual memory processes across time (Box 3).

Attention and encoding limits

Another major source of capacity limits in both visual working and long-term memory tasks are due to attentional processes during the initial encoding of items rather than maintenance processes. This insight initially stems from research on change blindness, a phenomenon wherein people fail to detect changes in an image after a brief transient, such as an eye movement or a flicker of the display¹³⁴. Importantly, these changes can be large and appear in salient parts of the display, meaning that change blindness does not simply reflect failures of the perceptual system¹³⁵. Change blindness effects have been used to motivate

theories of consciousness^{136,137} and working memory^{138,139} because they suggest that people remember only a small portion of what they see despite having a rich phenomenological experience of their immediate environment. Critically, one major way to nearly eliminate change blindness effects is to direct people's attention to objects^{126,140}, which ensures that objects are encoded into memory.

Other work suggests that encoding in visual working memory predicts the bandwidth of visual long-term memory. For instance, the effects of interference can be offset via attentional processes^{141–145} that induce repulsive biases at encoding and individuate similar memory representations in visual long-term memory^{58,146}. Other researchers have found that individual differences in performance on visual working memory tasks predicted performance on visual long-term memory tasks, but only under conditions in which visual working memory load is taxed¹⁴⁷. These results suggest that effects on visual long-term memory were due to failures of encoding rather than of maintenance capacity, and underscores how controlling for task demands aligns with a unidimensional model of visual memory across timescales. Other work also provides converging support for the view that manipulations of attention can help to upweight and improve long-term memory for objects¹⁴⁸. However, attentional processes might only be used to deprioritize irrelevant memory representations indirectly, under conditions of biased competition¹⁴⁹. More precisely, some evidence suggests that beneficial effects of attention occur only when multiple (two) stimuli were presented simultaneously and the irrelevant item was paired with an item than needed to be attended. This work directly connects to the finding that simultaneous as opposed to serial presentation in visual memory tasks can introduce attentional capacity limits and uncover structural bottlenecks¹⁵⁰.

Thus, attentional limits at encoding are common to both visual working and long-term memory despite being generally understudied in visual long-term memory. Attentional limits at encoding, as well as online activation, fit well with the view that memory representations are fundamentally noisy. These limits can therefore be viewed as one kind of resource that places limits on memory performance when encoding demands are high either in short-term or long-term visual memory tasks. More broadly, poor performance in most visual working memory tasks might reflect differences in encoding demands of visual working memory and visual long-term memory laboratory tasks in addition to or instead of differential memory maintenance abilities or differential memory representations over different timescales themselves. Thus, rather than viewing visual working and long-term memory as distinct systems that differ in capacity, a cognitive architecture in which attentional maintenance supports working activation of a subset of items that otherwise share the same underlying representation as long-term memory is a useful framework.

Influence of knowledge on visual memory

Visual memory is often assumed to be inherently perceptual in nature. However, existing knowledge and hierarchical knowledge structures help to scaffold memory of real-world visual stimuli in both working memory and long-term memory. This influence is taken for granted in visual long-term memory tasks, in which participants tend to be shown real-world scenes, objects or faces as stimuli^{143,151} and where conceptual knowledge has long been known to scaffold visual memories. For example, one of the most classic studies in constructive memory had people repeatedly draw an ambiguously face-like visual stimulus from memory and found that these drawings slowly morphed over repeated reproductions to look more like a genuine face²⁷. These results

demonstrate that people scaffold their memory for visual features using their knowledge of faces. Similarly, classic work has found that visual reproductions of ambiguous images are pulled in the direction of the labels associated with the images²⁸.

By contrast, the role of knowledge in visual working memory tasks has largely been underappreciated. Below, we describe how considering the role of meaning in shorter-term visual memory tasks has transformed views on how visual memory representations are maintained in memory tasks. We then disentangle the joint contribution of stimulus complexity and meaning in shaping visual memory representations. Finally, we integrate these points by discussing cognitive and neural evidence for the hierarchical nature of visual memory and discuss how hierarchical structures and priors scaffold memories for visual information in the real world.

Meaning in working memory

In traditional visual working memory tasks, people tend to be shown meaningless stimuli, such as circles, defined by simple features, such as colour^{50,54,66}. These simple, meaningless stimuli are often assumed to assess the core capacity of working memory best, because they have no semantic associations and therefore require ongoing active maintenance to remember them^{46,152}. Sensory recruitment models based on such tasks reinforce the idea that people store visual working memories within the sensory regions that process such stimuli (mainly low-level visual areas³⁸), which can lead to inferences because they are stored solely in an inherently perceptual format.

Using simple, single-feature stimuli probably does increase the need for attentional maintenance to protect stimuli from noise, because such stimuli provide very poor long-term memory traces¹⁵³.

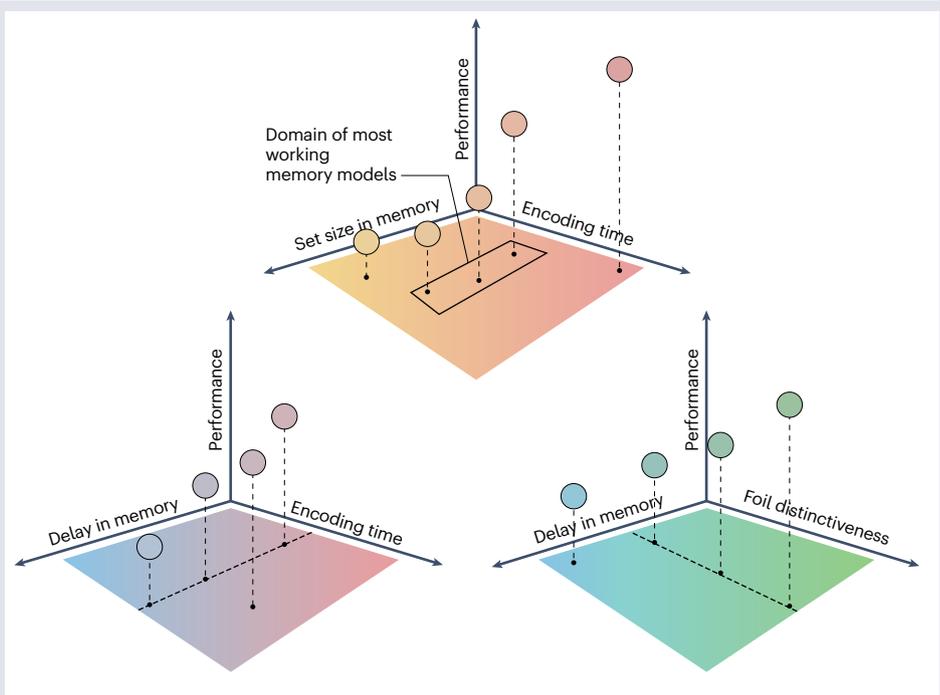
Box 3

Continuity in memory demands richer models

The finding that performance increases relatively continuously across different visual working and long-term memory manipulations and tasks raises the question of whether a single capacity limit is a useful construct, particularly in the domain of working memory where it is often used to motivate architectures and theoretical models and to fuel debates^{47,70}.

Models of visual working memory are typically developed to explain how performance varies as a function of how many items are simultaneously, briefly presented^{65,69,226}. As such, these models explain just a single slice through the space of possible levels of memory performance (top graph in figure). For instance, although modulating encoding time and set size are quite common, any mixture of these factors with (for example) delay in memory (left graph in figure), or with how distinctive test items are compared to those in memory, influences performance ('foil distinctiveness'; right graph in figure). However, these variables are not often manipulated in systematic ways because 'resource' models generally do not provide an integrative framework for understanding how these levers influence the availability of a resource or noise accrual.

Given these limitations, a fruitful direction for visual working memory research will be towards the study of long-term recognition and recall memory. In this domain, researchers often build integrative architectural models that describe processes that limit performance as a function of the particular encoding and retrieval manipulations for a given set of stimuli^{106,167,227}, as opposed to using vaguer constructs such as capacity or limited resources. For example, considerations of how the distinctiveness of stimuli relates



to the memory cues used will probably be useful in working memory research because the similarity of alternative items presented at test has an important role in shaping memory performance. Similarly, some stimuli spaces — such as real-world objects — have a higher upper bound on distinctiveness (right graph in figure) than those typically modelled (such as colours)¹⁰⁷. Taking into account the full set of manipulations that affect performance, as opposed to heavily prioritizing the number of items that are simultaneously present, will lead to a deeper understanding of how visual memory representations are affected by task demands and provide a common framework for thinking about visual memory at different timescales.

However, reducing the number of features that can be stored about an object, for example by using meaningless stimuli such as coloured shapes, also reduces how distinctive each object is in memory. This lack of distinctiveness leads to inter-item interference and confusability at retrieval that inevitably reduces performance in working memory tasks, leading to an underestimation of both visual working memory performance and the role of knowledge in scaffolding such performance. Although such simple stimuli help to ensure the use of attentional maintenance, their use does not necessarily uncover the true architecture or computations on which these visual working memory tasks rely^{154–157}.

Consider memory for Mooney faces, two-tone images that are sometimes perceived as a face and other times appear as a meaningless black and white blob (Fig. 4). In long-term memory, such stimuli have been used to show that recognition memory for exactly the same item can be improved when they are perceived as meaningful (as faces) versus not (as shape blobs)^{107,158}. Later work built on this idea in the study of working memory, demonstrating that perceiving a Mooney face as a face rather than a meaningless blob improves working memory performance and leads to an increase in the contralateral delay activity, a neural index of active maintenance engagement¹⁵⁹. There are further examples in which visually identical stimuli are better remembered in working memory tasks when they can be processed in a way that connects them to higher-level features or previous knowledge^{52,159–161}. These benefits extend to visual working memory performance whenever prior knowledge and expertise can be used to scaffold memory^{162,163}, such as for real-world objects¹⁶⁴, famous faces¹⁶⁵ and functional relationships between objects¹⁶⁶. Overall, this pattern suggests that visual working memory, like visual long-term memory, is not based solely on perceptual features. Instead, more meaningful stimuli are better remembered than less meaningful stimuli, even when perceptually identical.

Complexity and meaning

In contrast to the predictions of rudimentary information-theory models, people are far better at representing complex, meaningful stimuli than simple meaningless stimuli, an idea that dates to very early work

on ‘chunking’¹¹⁵ (Fig. 4a). Although this early work subscribed to very discrete views of chunking and capacity limits that are inconsistent with modern noise-based theories, its core message remains critical and underappreciated as an explanatory factor in visual working memory studies.

One explanation for the memory benefits of perceiving a meaningful and complex, multidimensional stimulus¹⁵⁹ is that doing so enhances the ability to extract features from that image. One specific hypothesis is that people are not completely flexible in their encoding, relying heavily on pre-learned features (for instance, using phonological features to store both binary and decimal digits, rather than adaptively switching to a code that takes advantage of the additional compressibility of binary digits¹²⁶). Thus, when remembering an unfamiliar, meaningless image one can only encode it in terms of its low-level shape, spatial frequency, and other low-level and mid-level features (Fig. 4b). However, when perceiving the same stimulus as meaningful, one unlocks higher level visual features, such as face-specific features including eye position, age and nose angle. This additional complexity probably improves performance because each stimulus is more distinct from the other stimuli (Fig. 4c). This role of featural distinctiveness is usually modelled in the long-term recognition memory literature¹⁶⁷ but rarely in working memory tasks. Furthermore, recognizing features at a higher level uncovers relations between features and makes it possible to efficiently encode and chunk multiple lower-level features jointly instead of separately, which also improves memory^{126,168}.

Importantly, adding arbitrary features to an object does not improve memory performance when those features are not part of one’s library of pre-learned features. For instance, although scrambled or inverted stimuli can be more complex and perceptually rich than simpler single-feature stimuli, they do not offer the opportunity to make use of existing higher-level features in the same way as realistic, meaningful stimuli presented in their pre-learned configuration. Indeed, memorizing more versus less visually complex meaningless stimuli hurts visual memory performance^{107,169,170}. Therefore, the benefits of extracting additional features from a visual object are only observed with stimuli for which there are existing ways to encode high-level

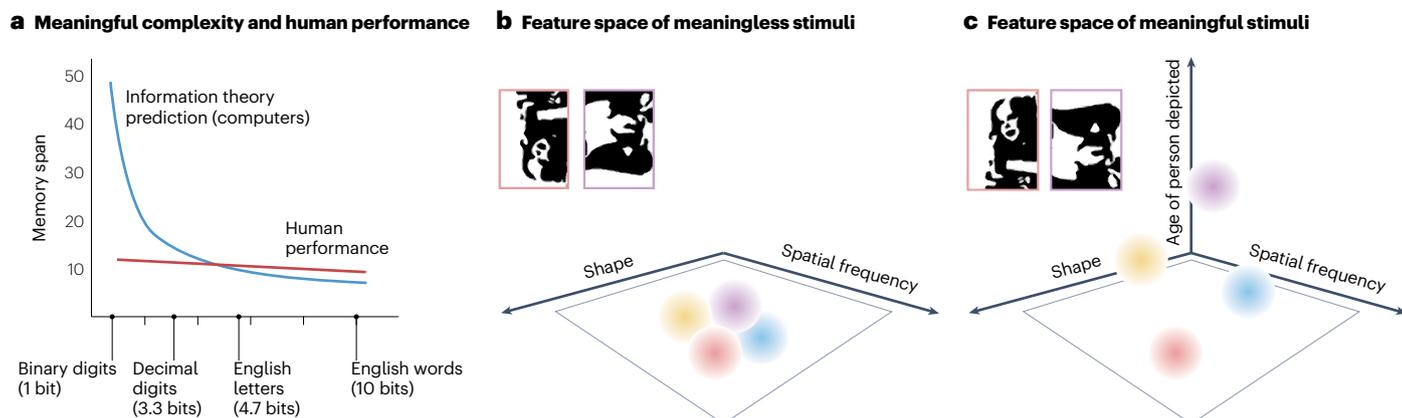


Fig. 4 | The effect of complexity and meaning. **a**, As individual items become more complex, a computer or other agent that encodes the stimulus relatively optimally and flexibly can store many fewer items (because they have a fixed capacity in bits). However, human performance remains relatively constant as complexity increases, as long as the information is meaningful, or capable of being scaffolded by previously acquired information. **b**, When inverted Mooney

faces (meaningless stimuli) are maintained in memory, human performance is relatively poor compared to when the same faces are upright (and recognized as faces). One hypothesis is that existing high-level features that apply to perceived faces but not to shapes add meaningful dimensions to the stimuli, therefore making them more separable and robust to interference and noise (such as age of the face, shown on the z axis).

visual and semantic information, but not for scrambled yet perceptually complex stimuli^{161,166}. This concept is sometimes formalized in working memory as a kind of ‘model mismatch’¹²⁷: stimuli that match an internal model of the world are easier to remember than those that do not.

As we elaborate below, the beneficial role of prior knowledge bridges many findings on how complexity and meaning can jointly improve visual working memory performance. These include benefits when multiple features are chunked into a single object as compared to separate ones⁸⁸, better performance at remembering recognized Mooney faces than unrecognized Mooney faces¹⁵⁹ and other aspects of model mismatch¹²⁷, as well as improvements in working memory performance after learning which features tend to go together¹²⁶. Conceptualizing chunking in a more continuous way, consistent with noisy representations, remains an important goal for future work on visual memory¹⁷¹. Conceptions of chunking construed in terms of more sophisticated formal models of information compression or rate-distortion theory^{126,128}, which have the potential to take into account differential effects of stimulus complexity, will probably be critical to future progress in understanding these phenomena. Through this more contemporary lens, visual memory chunking can be seen as a form of ‘lossy’ compression of continuous multidimensional variables: compression of memory representations, resulting in some information loss by forming more efficient, more abstract representations that are linked to prior knowledge structures. More generally, although item-based theories have often been used to explain chunking in memory^{172–174}, continuous theories of memory strength are equally compatible with the view that memory representations can be translated into more efficient formats, and that the common link between discrete-slot models and compressibility is historical rather than logical.

Elaborative encoding and hierarchical memories

The idea that recognizing and connecting an object to high-level visual and semantic features improves memory is closely related to the popular idea of ‘levels of processing’, which is more often studied for verbal than visual information^{175,176}. According to this view, processing of stimuli at a ‘deeper’ level leads to more elaborated memory traces¹⁷⁷. Such memory traces are then more distinctive^{178,179} and easier to retrieve because they are more connected to multiple distinct knowledge structures that can serve as retrieval cues^{175,180}. Elaborative encoding in visual memory can occur when encoding semantic features, but also mid-level (such as shape) and high-level (such as the distance between two eyes on a face) visual features. Thus, the hierarchical structure of visual objects and scenes, in terms of low-level to high-level visual features and semantic information, can enable improved memory performance.

Currently, how the hierarchy of visual features scaffolds memory performance and improves the ability to remember realistic stimuli has not been explained by most models of visual working memory performance. Instead, even models that focus on prior knowledge tend to focus on memory capacity rather than memory representations (such as rate distortion or information-theory models, as discussed above). Thus, these models tend to focus on how performance varies as a function of the number of simple visual features shown. Important recent work has attempted to model representations of realistic stimuli using generative deep learning neural nets to capture how stimuli are represented at multiple levels of the visual hierarchy¹⁸¹. This model attempts to mimic the hierarchical structure of the ventral visual pathway, capturing the encoding of low-level (such as orientation) and mid-level (such as shape) visual features. Critically, by using recurrent feedback mechanisms, the model can be used to reconstruct

important memory phenomena, such as more resource-efficient representation of familiar items, and efficient recognition of novel objects that share high-level features with familiar items. These findings and related work^{182,183} highlight that interactions across the visual hierarchy are bidirectional, meaning that signals from early areas of the visual hierarchy influence representations at higher levels, and vice versa. Studying these interactions is an important step toward expanding models of visual working memory to account for a feature hierarchy.

In contrast to the visual working memory literature, the role of elaborative encoding and the availability of high-level or semantic features has more frequently been studied in the visual long-term memory literature. For example, the role of categories in memory, such as the finding that memories might be biased towards often-encountered instances of a category, suggests that semantic features, in addition to perceptual features, play a substantial part in memory encoding^{184–186}. Other work has also shown that normal, unmodified faces activate higher-dimensional neural representations (involving high-level face-processing regions) and result in better memory compared to modified faces¹⁸⁷. The literature on memorability also taps into these concepts. Memorability is based on the empirical finding that some stimuli are more often remembered successfully, compared to other – sometimes very similar – stimuli. Although this concept is sometimes labelled intrinsic memorability, classic memory theories^{175,178} would suggest that some stimuli are more memorable because they connect more directly to higher-level visual and semantic features, enabling more elaborative encoding, and are therefore more distinctive (with a lowered summed similarity to the current memory set¹⁸⁸), which is not an intrinsic property of the stimulus. Teasing apart the relative contributions of elaborative encoding and distinctiveness within the set of items in memory to the reliability of responses across individuals remains an active area of work in the memorability literature^{4,189}.

In addition to the fact that features of a visual stimulus can be represented at multiple levels in the visual hierarchy (such as more perceptual or more conceptual), there is an even stronger sense in which visual memory can be considered hierarchical. Visual memory might be fundamentally structured, with distinct representations at different levels of abstraction. One reliable phenomenon that supports this view is that people’s memories are biased towards, but are not replaced with, their priors^{57,190,191}. For example, memories for the size of specific instances of natural objects (such as fruits and vegetables) are biased towards the average size of the object categories (size of all apples) and superordinate categories (size of all fruits¹⁹²; Box 1). This effect of priors on memory even holds in tasks in which priors are not informative for the stimulus selection process. For instance, people use knowledge of colour categories even in a task in which items’ colours are randomly chosen¹⁹³ – such that even for extremely simple stimuli like coloured dots, visual memories naturally drift toward prior knowledge of colour categories¹⁹⁴. More precisely, evidence suggests that people see certain hues of a colour as being more representative of the colour category and their memory will tend to drift towards these representative feature values.

In addition to memory depending on knowledge acquired before encoding, memories also seem to be encoded hierarchically within an episode. For example, when faced with sequences of items, in visual working memory or long-term memory tasks, people make use of similarity among multiple items to store abstractions and those abstractions influence subsequent memory for individual items^{195–197}. For instance, scenes help to provide a structure to the objects within them, as shown by evidence that if people are presented with an array

of objects that are embedded within a scene, people encode both the ensemble structure of the entire scene in addition to individual items^{102,198,199}. Similarly, when trying to remember real-world scenes, people encode the gist plus detailed information about some specific objects^{200,201}. Moreover, they use the gist to guide their choice of which specific objects to remember^{202,203} and are influenced by the gist when later trying to recall the details of the scene, tending to remember objects that are consistent with the scene gist even if they were not present^{204–206}. Even in randomly generated displays of simple colours, people are typically able to extract some spatial or feature-based structure from the display, which they use to encode the items with respect to each other rather than totally independently¹⁹⁹. Finally, in some situations, people seem to encode relatively separate texture or ensemble information in addition to individual items and use it to alter their responses²⁰⁷.

In long-term memory, it is also thought that episodic memories are hierarchical. Episodic long-term memories integrate across several levels²⁰⁸ of storing both gist and detailed item information and trading them off³⁴. These episodic memory processes flexibly integrate disparate components (including details and related information) of an autobiographical event into a coherent representation of the past or future^{208,209}.

To summarize, visual memory in the real-world operates on meaningful and familiar items, such as real-world objects and scenes. Memory for real-world objects is jointly but differently influenced by stimulus meaning and complexity. Meaning can bolster memory by distributing information across levels of the visual hierarchy, with memory for perceptual features influenced by and supported by memory for more abstract conceptual features. Added complexity can further scaffold memory by providing additional relational information that can be used to compress meaningful information, or serve as an additional retrieval cue. Likewise, such information is heavily influenced by prior knowledge, which can adaptively bias memory representations to optimize memory performance. Collectively, the study of memory for meaningful as opposed to meaningless, degraded visual information can help to ground theories and models of visual memory in the real world.

Summary and future directions

The visual system is usually conceived of in terms of a hierarchy of population codes, with representations that become more complex from the primary visual cortex up to more anterior visual regions. In this Review, we have suggested that a similar set of ideas can provide a useful mental model for understanding visual memory representations, with representations of objects or scenes conceived of as population-based, noisy (variable in strength) and stored in terms of hierarchical and distributed memory representations.

This conception of visual memory raises major challenges for measuring memory performance. In particular, modelling the decision-making process people use to integrate across an entire hierarchy of noisy representations will be critical to successfully measuring and understanding the memory representations themselves. In this context, a fruitful direction is to examine how model-based approximations of representations – such as those instantiated via convolutional neural nets^{210,211} or probabilistic models of knowledge structures²¹² – can yield insights into how continuous representations arise and are integrated across the visual hierarchy.

Greater synergy between the kinds of model and mechanism proposed in visual working memory and visual long-term memory tasks will be critical to progress across both domains, but particularly in the

domain of working memory. Existing working memory models are largely focused on manipulations of memory load for brief, simultaneous presentations of low-level visual features, and more cohesive models that use many of the same principles proposed in models of long-term memory are needed to examine and explain how encoding, delay, frequency of item presentation, context, testing conditions²¹³ and item influence visual working memory processes. Such modelling approaches could also help to illuminate the intersection between visual working and long-term memory and the role of limits on attentional protection of items from noise accumulation in real-world tasks. It is therefore also important to continue work that precisely defines – at the cognitive and neural level – how differences in ‘activation’ or attentional engagement can distinguish visual working memory and long-term memory.

Finally, a focus on precise quantification of memory representations in visual memory (as in continuous reproduction) and on research domains in which representations are well understood (such as perception and other domains where neuroscience informs the precise representation structure of items) will enable the development of detailed computational models that can help to elucidate performance. However, it is also critical to consider that impoverished stimuli like single visual features do not capture the full breadth of scenarios where memory is used in the real world, and so models – particularly of visual working memory performance – must carefully consider the role of elaborative encoding of more realistic stimuli.

Published online: 08 February 2024

References

1. Foer, J. *Moonwalking with Einstein: The Art and Science of Remembering Everything* (Penguin, 2012).
2. Squire, L. R., Knowlton, B. & Musen, G. The structure and organization of memory. *Annu. Rev. Psychol.* **44**, 453–495 (1993).
3. Eichenbaum, H., Yonelinas, A. R. & Ranganath, C. The medial temporal lobe and recognition memory. *Annu. Rev. Neurosci.* **30**, 123 (2007).
4. Brady, T. F. & Bainbridge, W. A. *Visual Memory* (Routledge, 2022).
5. Luck, S. J. & Hollingworth, A. *Visual Memory* (Oxford Univ. Press, 2008).
6. Hubel, D. H. & Wiesel, T. N. Ferrier lecture – Functional architecture of macaque monkey visual cortex. *Proc. R. Soc. Lond. B* **198**, 1–59 (1977).
7. Logothetis, N. K. & Sheinberg, D. L. Visual object recognition. *Annu. Rev. Neurosci.* **19**, 577–621 (1996).
8. Yamins, D. L. K. et al. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proc. Natl Acad. Sci. USA* **111**, 8619–8624 (2014).
9. Todd, J. J. & Marois, R. Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature* **428**, 751–754 (2004).
10. Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).
11. DiCarlo, J. J., Zoccolan, D. & Rust, N. C. How does the brain solve visual object recognition? *Neuron* **73**, 415–434 (2012).
12. Beschin, N., Cocchini, G., Della Sala, S. & Logie, R. H. What the eyes perceive, the brain ignores: a case of pure unilateral representational neglect. *Cortex* **33**, 3–26 (1997).
13. Guariglia, C., Padovani, A., Pantano, P. & Pizzamiglio, L. Unilateral neglect restricted to visual imagery. *Nature* **364**, 235–237 (1993).
14. Hebb, D. O. *The Organization of Behavior: a Psychological Theory* (Wiley, 1949).
15. Fuster, J. M. & Alexander, G. E. Neuron activity related to short-term memory. *Science* **173**, 652–654 (1971).
16. Goldman-Rakic, P. S. Working memory and the mind. *Sci. Am.* **267**, 110–117 (1992).
17. Berggren, N. & Eimer, M. Does contralateral delay activity reflect working memory storage or the current focus of spatial attention within visual working memory? *J. Cogn. Neurosci.* **28**, 2003–2020 (2016).
18. Ikkai, A., McCollough, A. W. & Vogel, E. K. Contralateral delay activity provides a neural measure of the number of representations in visual working memory. *J. Neurophysiol.* **103**, 1963–1968 (2010).
19. Kang, M.-S. & Woodman, G. F. The neurophysiological index of visual working memory maintenance is not due to load dependent eye movements. *Neuropsychologia* **56**, 63–72 (2014).
20. Klaver, P., Talsma, D., Wijers, A. A., Heinze, H.-J. & Mulder, G. An event-related brain potential correlate of visual short-term memory. *NeuroReport* **10**, 2001–2005 (1999).
21. Pomper, U., Ditye, T. & Ansorge, U. Contralateral delay activity during temporal order memory. *Neuropsychologia* **129**, 104–116 (2019).

22. Carlisle, N. B., Arita, J. T., Pardo, D. & Woodman, G. F. Attentional templates in visual working memory. *J. Neurosci.* **31**, 9315–9322 (2011).
23. Barbosa, J., Lozano-Soldevilla, D. & Compte, A. Pinging the brain with visual impulses reveals electrically active, not activity-silent, working memories. *PLoS Biol.* **19**, e3001436 (2021).
24. Harrison, S. A. & Tong, F. Decoding reveals the contents of visual working memory in early visual areas. *Nature* **458**, 632–635 (2009).
25. Serences, J. T., Ester, E. F., Vogel, E. K. & Awh, E. Stimulus-specific delay activity in human primary visual cortex. *Psychol. Sci.* **20**, 207–214 (2009).
26. Morey, C. C. The case against specialized visual-spatial short-term memory. *Psychol. Bull.* **144**, 849–883 (2018).
27. Bartlett, F. C. *Remembering: A study in Experimental and Social Psychology* (Cambridge Univ. Press, 1995).
28. Carmichael, L., Hogan, H. P. & Walter, A. A. An experimental study of the effect of language on the reproduction of visually perceived form. *J. Exp. Psychol.* **15**, 73–86 (1932).
29. Logie, R. H. *Visuo-spatial Working Memory* (Psychology Press, 2014).
30. Logie, R. H., Belletier, C. & Doherty, J. M. in *Working Memory: State of the Science* (eds Logie, R. H. et al.) 389–429 (Oxford Univ. Press, 2021).
31. Atkinson, R. C. & Shiffrin, R. M. in *Psychology of Learning and Motivation* vol. 2 (eds Spence, K. W. & Spence, J. T.) 89–195 (Elsevier, 1968).
32. Baddeley, A. D. The influence of acoustic and semantic similarity on long-term memory for word sequences. *Q. J. Exp. Psychol.* **18**, 302–309 (1966).
33. Cowan, N. What are the differences between long-term, short-term, and working memory? *Prog. Brain Res.* **169**, 323–338 (2008).
34. Brainerd, C. J. & Reyna, V. F. Fuzzy-trace theory and false memory. *Curr. Dir. Psychol. Sci.* **11**, 164–169 (2002).
35. Miner, A. E., Schurgin, M. W. & Brady, T. F. Is working memory inherently more “precise” than long-term memory? Extremely high fidelity visual long-term memories for frequently encountered objects. *J. Exp. Psychol. Hum. Percept. Perform.* **46**, 813–830 (2020).
36. Schurgin, M. W., Wixted, J. T. & Brady, T. F. Psychophysical scaling reveals a unified theory of visual memory strength. *Nat. Hum. Behav.* **4**, 1156–1172 (2020).
37. Serences, J. T. & Yantis, S. Selective visual attention and perceptual coherence. *Trends Cogn. Sci.* **10**, 38–45 (2006).
38. Sprague, T. C., Ester, E. F. & Serences, J. T. Restoring latent visual working memory representations in human cortex. *Neuron* **91**, 694–707 (2016).
39. Stokes, M. G. ‘Activity-silent’ working memory in prefrontal cortex: a dynamic coding framework. *Trends Cogn. Sci.* **19**, 394–405 (2015).
40. Cunningham, C. A. & Wolfe, J. M. The role of object categories in hybrid visual and memory search. *J. Exp. Psychol. Gen.* **143**, 1585 (2014).
41. Wolfe, J. M., Boettcher, S. E., Josephs, E. L., Cunningham, C. A. & Drew, T. You look familiar, but I don’t care: lure rejection in hybrid visual and memory search is not based on familiarity. *J. Exp. Psychol. Hum. Percept. Perform.* **41**, 1576 (2015).
42. Baddeley, A. Working memory: looking back and looking forward. *Nat. Rev. Neurosci.* **4**, 829–839 (2003).
43. Chun, M. M., Golomb, J. D. & Turk-Browne, N. B. A taxonomy of external and internal attention. *Annu. Rev. Psychol.* **62**, 73–101 (2011).
44. Postle, B. R., Druzgal, T. J. & D’Esposito, M. Seeking the neural substrates of visual working memory storage. *Cortex* **39**, 927–946 (2003).
45. Oberauer, K. Access to information in working memory: exploring the focus of attention. *J. Exp. Psychol. Learn. Mem. Cogn.* **28**, 411 (2002).
46. Cowan, N. The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav. Brain Sci.* **24**, 87–114 (2001).
47. Luck, S. J. & Vogel, E. K. Visual working memory capacity: from psychophysics and neurobiology to individual differences. *Trends Cogn. Sci.* **17**, 391–400 (2013).
48. Markov, Y. A., Utochkin, I. S. & Brady, T. F. Real-world objects are not stored in holistic representations in visual working memory. *J. Vis.* **21**, 18 (2021).
49. Rademaker, R. L., Park, Y. E., Sack, A. T. & Tong, F. Evidence of gradual loss of precision for simple features and complex objects in visual working memory. *J. Exp. Psychol. Hum. Percept. Perform.* **44**, 925–940 (2018).
50. Wilken, P. & Ma, W. J. A detection theory account of change detection. *J. Vis.* **4**, 1120–1135 (2004).
51. Williams, J. R., Robinson, M. M., Schurgin, M. W., Wixted, J. T. & Brady, T. F. You cannot “count” how many items people remember in visual working memory: the importance of signal detection-based measures for understanding change detection performance. *J. Exp. Psychol. Hum. Percept. Perform.* **48**, 1390–1409 (2022).
52. Alvarez, G. A. & Cavanagh, P. The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychol. Sci.* **15**, 106–111 (2004).
53. Awh, E., Barton, B. & Vogel, E. K. Visual working memory represents a fixed number of items regardless of complexity. *Psychol. Sci.* **18**, 622–628 (2007).
54. Luck, S. J. & Vogel, E. K. The capacity of visual working memory for features and conjunctions. *Nature* **390**, 279–281 (1997).
55. Brady, T. F., Konkle, T., Alvarez, G. A. & Oliva, A. Visual long-term memory has a massive storage capacity for object details. *Proc. Natl Acad. Sci. USA* **105**, 14325–14329 (2008).
56. Bae, G. Y. & Luck, S. J. Interactions between visual working memory representations. *Atten. Percept. Psychophys.* **79**, 2376–2395 (2017).
57. Brady, T. F. & Alvarez, G. A. Hierarchical encoding in visual working memory: ensemble statistics bias memory for individual items. *Psychol. Sci.* **22**, 384–392 (2011).
58. Chunharas, C., Rademaker, R. L., Brady, T. F. & Serences, J. T. An adaptive perspective on visual working memory distortions. *J. Exp. Psychol. Gen.* **151**, 2300–2323 (2022).
59. Golomb, J. D. Divided spatial attention and feature-mixing errors. *Atten. Percept. Psychophys.* **77**, 2562–2569 (2015).
60. Lively, Z., Robinson, M. M. & Benjamin, A. S. Memory fidelity reveals qualitative changes in interactions between items in visual working memory. *Psychol. Sci.* **32**, 1426–1441 (2021).
61. Scotti, P. S., Hong, Y., Golomb, J. D. & Leber, A. B. Statistical learning as a reference point for memory distortions: swap and shift errors. *Atten. Percept. Psychophys.* **83**, 1652–1672 (2021).
62. Pertzov, Y., Bays, P. M., Joseph, S. & Husain, M. Rapid forgetting prevented by retrospective attention cues. *J. Exp. Psychol. Hum. Percept. Perform.* **39**, 1224 (2013).
63. Makovski, T. & Jiang, Y. V. Distributing versus focusing attention in visual short-term memory. *Psychon. Bull. Rev.* **14**, 1072–1078 (2007).
64. Marini, F., Scott, J., Aron, A. R. & Ester, E. F. Task-irrelevant distractors in the delay period interfere selectively with visual short-term memory for spatial locations. *Atten. Percept. Psychophys.* **79**, 1384–1392 (2017).
65. van den Berg, R., Shin, H., Chou, W.-C., George, R. & Ma, W. J. Variability in encoding precision accounts for visual short-term memory limitations. *Proc. Natl Acad. Sci. USA* **109**, 8780–8785 (2012).
66. Bays, P. M. Noise in neural populations accounts for errors in working memory. *J. Neurosci.* **34**, 3632–3645 (2014).
67. van den Berg, R., Awh, E. & Ma, W. J. Factorial comparison of working memory models. *Psychol. Rev.* **121**, 124–149 (2014).
68. Adam, K. C. S., Vogel, E. K. & Awh, E. Clear evidence for item limits in visual working memory. *Cogn. Psychol.* **97**, 79–97 (2017).
69. Zhang, W. & Luck, S. J. Discrete fixed-resolution representations in visual working memory. *Nature* **453**, 233–235 (2008).
70. Schneegans, S., Taylor, R. & Bays, P. M. Stochastic sampling provides a unifying account of visual working memory limits. *Proc. Natl Acad. Sci. USA* **117**, 20959–20968 (2020).
71. Pratte, M. S. Set size effects on working memory precision are not due to an averaging of slots. *Atten. Percept. Psychophys.* **82**, 2937–2949 (2020).
72. Bays, P., Schneegans, S., Ma, W. J. & Brady, T. F. Representation and computation in working memory. Preprint at [PsyArXiv https://doi.org/10.31234/osf.io/kubr9](https://doi.org/10.31234/osf.io/kubr9) (2022).
73. Swets, J. A. Form of empirical ROCs in discrimination and diagnostic tasks: implications for theory and measurement of performance. *Psychol. Bull.* **99**, 181–198 (1986).
74. Wixted, J. T. Dual-process theory and signal-detection theory of recognition memory. *Psychol. Rev.* **114**, 152–176 (2007).
75. Wixted, J. T. & Mickes, L. A continuous dual-process model of remember/know judgments. *Psychol. Rev.* **117**, 1025–1054 (2010).
76. Taylor, R. & Bays, P. M. Theory of neural coding predicts an upper bound on estimates of memory variability. *Psychol. Rev.* **127**, 700 (2020).
77. Delay, C. G. & Wixted, J. T. Discrete-state versus continuous models of the confidence-accuracy relationship in recognition memory. *Psychon. Bull. Rev.* **28**, 556–564 (2021).
78. Brady, T. F., Robinson, M. M., Williams, J. R. & Wixted, J. T. Measuring memory is harder than you think: how to avoid problematic measurement practices in memory research. *Psychon. Bull. Rev.* **30**, 421–449 (2023).
79. Honig, M., Ma, W. J. & Fougny, D. Humans incorporate trial-to-trial working memory uncertainty into rewarded decisions. *Proc. Natl Acad. Sci. USA* **117**, 8391–8397 (2020).
80. Jabar, S. B. et al. Using a betting game to reveal the rich nature of visual working memories. Preprint at [bioRxiv https://doi.org/10.1101/2020.10.28.357442](https://doi.org/10.1101/2020.10.28.357442) (2020).
81. Yoo, A. H., Acerbi, L. & Ma, W. J. Uncertainty is maintained and used in working memory. *J. Vis.* **21**, 13 (2021).
82. Treisman, A. Features and objects in visual processing. *Sci. Am.* **13**, 114–125 (1986).
83. Schneegans, S. & Bays, P. M. Neural architecture for feature binding in visual working memory. *J. Neurosci.* **37**, 3913–3925 (2017).
84. Bays, P. M., Wu, E. Y. & Husain, M. Storage and binding of object features in visual working memory. *Neuropsychologia* **49**, 1622–1631 (2011).
85. Fougny, D. & Alvarez, G. A. Object features fail independently in visual working memory: evidence for a probabilistic feature-store model. *J. Vis.* **11**, 3 (2011).
86. Shin, H. & Ma, W. J. Visual short-term memory for oriented, colored objects. *J. Vis.* **17**, 12 (2017).
87. Utochkin, I. S. & Brady, T. F. Independent storage of different features of real-world objects in long-term memory. *J. Exp. Psychol. Gen.* **149**, 530–549 (2020).
88. Fougny, D., Cormiea, S. M. & Alvarez, G. A. Object-based benefits without object-based representations. *J. Exp. Psychol. Gen.* **142**, 621–626 (2013).
89. Wang, B., Cao, X., Theeuwes, J., Olivers, C. N. & Wang, Z. Location-based effects underlie feature conjunction benefits in visual working memory. *J. Vis.* **16**, 12 (2016).
90. Egly, R., Driver, J. & Rafal, R. D. Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *J. Exp. Psychol. Gen.* **123**, 161–177 (1994).
91. Francis, G. & Thunell, E. Excess success in articles on object-based attention. *Atten. Percept. Psychophys.* **84**, 700–714 (2022).
92. Chater, N., Tenenbaum, J. B. & Yuille, A. Probabilistic models of cognition: conceptual foundations. *Trends Cogn. Sci.* **10**, 287–291 (2006).
93. Ma, W. J. Organizing probabilistic models of perception. *Trends Cogn. Sci.* **16**, 511–518 (2012).
94. Kellen, D., Winiger, S., Dunn, J. C. & Singmann, H. Testing the foundations of signal detection theory in recognition memory. *Psychol. Rev.* **128**, 1022–1050 (2021).

95. Rotello, C. M., Heit, E. & Dubé, C. When more data steer us wrong: replications with the wrong dependent measure perpetuate erroneous conclusions. *Psychon. Bull. Rev.* **22**, 944–954 (2015).
96. Starns, J. J. et al. Assessing theoretical conclusions with blinded inference to investigate a potential inference crisis. *Adv. Meth. Pract. Psychol. Sci.* **2**, 335–349 (2019).
97. Wixted, J. T. & Mickes, L. Theoretical vs. empirical discriminability: the application of ROC methods to eyewitness identification. *Cogn. Res. Princ. Implic.* **3**, 9 (2018).
98. Shen, K. J., Colloff, M. F., Vul, E., Wilson, B. M. & Wixted, J. T. Modeling face similarity in police lineups. *Psychol. Rev.* **130**, 432–461 (2023).
99. Adler, W. T. & Ma, W. J. Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLoS Comput. Biol.* **14**, e1006572 (2018).
100. Ma, W. J. Bayesian decision models: a primer. *Neuron* **104**, 164–175 (2019).
101. Pouget, A., Beck, J. M., Ma, W. J. & Latham, P. E. Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* **16**, 1170–1178 (2013).
102. Brady, T. F. & Tenenbaum, J. B. A probabilistic model of visual working memory: incorporating higher order regularities into working memory capacity estimates. *Psychol. Rev.* **120**, 85–109 (2013).
103. Smith, P. L., Saber, S., Corbett, E. A. & Lilburn, S. D. Modeling continuous outcome color decisions with the circular diffusion model: metric and categorical properties. *Psychol. Rev.* **127**, 562 (2020).
104. Ma, W. J., Husain, M. & Bays, P. M. Changing concepts of working memory. *Nat. Neurosci.* **17**, 347–356 (2014).
105. Schurgin, M. W. & Flombaum, J. I. Visual working memory is more tolerant than visual long-term memory. *J. Exp. Psychol. Hum. Percept. Perform.* **44**, 1216–1227 (2018).
106. Howard, M. W. & Kahana, M. J. A distributed representation of temporal context. *J. Math. Psychol.* **46**, 269–299 (2002).
107. Brady, T. F. & Störmer, V. S. The role of meaning in visual working memory: real-world objects, but not simple features, benefit from deeper processing. *J. Exp. Psychol. Learn. Mem. Cogn.* **48**, 942–958 (2022).
108. Endress, A. D. & Potter, M. C. Large capacity temporary visual memory. *J. Exp. Psychol. Gen.* **143**, 548–565 (2014).
109. Tsubomi, H., Fukuda, K., Watanabe, K. & Vogel, E. K. Neural limits to representing objects still within view. *J. Neurosci.* **33**, 8257–8263 (2013).
110. Vogel, E. K., Woodman, G. F. & Luck, S. J. The time course of consolidation in visual working memory. *J. Exp. Psychol. Hum. Percept. Perform.* **32**, 1436–1451 (2006).
111. Ricker, T. J. & Hardman, K. O. The nature of short-term consolidation in visual working memory. *J. Exp. Psychol. Gen.* **146**, 1551–1573 (2017).
112. Li, X., Xiong, Z., Theeuwes, J. & Wang, B. Visual memory benefits from prolonged encoding time regardless of stimulus type. *J. Exp. Psychol. Learn. Mem. Cogn.* **46**, 1998–2005 (2020).
113. Draschkow, D., Kallmayer, M. & Nobre, A. C. When natural behavior engages working memory. *Curr. Biol.* **31**, 869–874 (2021).
114. Williams, J. & Störmer, V. S. Working memory: how much is it used in natural behavior? *Curr. Biol.* **31**, R205–R206 (2021).
115. Miller, G. A. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychol. Rev.* **63**, 81–97 (1956).
116. Balaban, H. & Luria, R. The number of objects determines visual working memory capacity allocation for complex items. *NeuroImage* **119**, 54–62 (2015).
117. Barton, B., Ester, E. F. & Awh, E. Discrete resource allocation in visual working memory. *J. Exp. Psychol. Hum. Percept. Perform.* **35**, 1359–1367 (2009).
118. Rouder, J. N. et al. An assessment of fixed-capacity models of visual working memory. *Proc. Natl Acad. Sci. USA* **105**, 5975–5979 (2008).
119. Navon, D. Resources — a theoretical soup stone? *Psychol. Rev.* **99**, 216–234 (1984).
120. Franconeri, S. L., Alvarez, G. A. & Cavanagh, P. Flexible cognitive resources: competitive content maps for attention and memory. *Trends Cogn. Sci.* **17**, 134–141 (2013).
121. Bays, P. M. Spikes not slots: noise in neural populations limits working memory. *Trends Cogn. Sci.* **19**, 431–438 (2015).
122. Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 (2012).
123. Smith, P. L., Lilburn, S. D., Corbett, E. A., Sewell, D. K. & Kyllingsbæk, S. The attention-weighted sample-size model of visual short-term memory: attention capture predicts resource allocation and memory load. *Cogn. Psychol.* **89**, 71–105 (2016).
124. Green, D. M. & Swets, J. A. *Signal Detection Theory and Psychophysics* (John Wiley, 1966).
125. Palmer, J. Attentional limits on the perception and memory of visual information. *J. Exp. Psychol. Hum. Percept. Perform.* **16**, 332 (1990).
126. Brady, T. F., Konkle, T. & Alvarez, G. A. Compression in visual working memory: using statistical regularities to form more efficient memory representations. *J. Exp. Psychol. Gen.* **138**, 487–502 (2009).
127. Orhan, A. E., Sims, C. R., Jacobs, R. A. & Knill, D. C. The adaptive nature of visual working memory. *Curr. Dir. Psychol. Sci.* **23**, 164–170 (2014).
128. Bates, C. J. & Jacobs, R. A. Efficient data compression in perception and perceptual memory. *Psychol. Rev.* **127**, 891–917 (2020).
129. Van den Berg, R. & Ma, W. J. A resource-rational theory of set size effects in human visual working memory. *eLife* **7**, e34963 (2018).
130. Griffiths, T. L., Lieder, F. & Goodman, N. D. Rational use of cognitive resources: levels of analysis between the computational and the algorithmic. *Top. Cogn. Sci.* **7**, 217–229 (2015).
131. Kristjánsson, Á. & Draschkow, D. Keeping it real: looking beyond capacity limits in visual cognition. *Attent. Percept. Psychophys.* **83**, 1375–1390 (2021).
132. Hayhoe, M. & Ballard, D. Eye movements in natural behavior. *Trends Cogn. Sci.* **9**, 188–194 (2005).
133. Ballard, D. H., Hayhoe, M. M. & Pelz, J. B. Memory representations in natural tasks. *J. Cogn. Neurosci.* **7**, 66–80 (1995).
134. Simons, D. J. & Rensink, R. A. Change blindness: past, present, and future. *Trends Cogn. Sci.* **9**, 16–20 (2005).
135. Simons, D. J. & Ambinder, M. S. Change blindness: theory and consequences. *Curr. Dir. Psychol. Sci.* **14**, 44–48 (2005).
136. Dretske, F. What change blindness teaches about consciousness. *Phil. Persp.* **21**, 215–230 (2007).
137. O’Regan, J. K. & Noë, A. A sensorimotor account of vision and visual consciousness. *Behav. Brain Sci.* **24**, 939–973 (2001).
138. O’Regan, J. K. Solving the ‘real’ mysteries of visual perception: the world as an outside memory. *Can. J. Psychol.* **46**, 461–488 (1992).
139. Rensink, R. A. Change blindness. In *Neurobiology of Attention* 76–81 (Elsevier, 2005).
140. Hollingworth, A. Scene and position specificity in visual memory for objects. *J. Exp. Psychol. Learn. Mem. Cogn.* **32**, 58–69 (2006).
141. Jost, K. et al. Controlling conflict from interfering long-term memory representations. *J. Cogn. Neurosci.* **24**, 1173–1190 (2012).
142. Konkle, T., Brady, T. F., Alvarez, G. A. & Oliva, A. Scene memory is more detailed than you think: the role of categories in visual long-term memory. *Psychol. Sci.* **21**, 1551–1556 (2010).
143. Konkle, T., Brady, T. F., Alvarez, G. A. & Oliva, A. Conceptual distinctiveness supports detailed visual long-term memory for real-world objects. *J. Exp. Psychol. Gen.* **139**, 558 (2010).
144. Robertson, E. M. New insights in human memory interference and consolidation. *Curr. Biol.* **22**, R66–R71 (2012).
145. Wais, P. E. & Gazzaley, A. Distractibility during retrieval of long-term memory: domain-general interference, neural networks and increased susceptibility in normal aging. *Front. Psychol.* **5**, 280 (2014).
146. Drascher, M. L. & Kuhl, B. A. Long-term memory interference is resolved via repulsion and precision along diagnostic memory dimensions. *Psychon. Bull. Rev.* **29**, 1898–1912 (2022).
147. Fukuda, K. & Vogel, E. K. Visual short-term memory capacity predicts the ‘bandwidth’ of visual long-term memory encoding. *Mem. Cogn.* **47**, 1481–1497 (2019).
148. Sundby, C. S., Woodman, G. F. & Fukuda, K. Electrophysiological and behavioral evidence for attentional up-regulation, but not down-regulation, when encoding pictures into long-term memory. *Mem. Cogn.* **47**, 351–364 (2019).
149. Desimone, R. & Duncan, J. Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* **18**, 193–222 (2019).
150. Beck, D. M. & Kastner, S. Top-down and bottom-up mechanisms in biasing competition in the human brain. *Vis. Res.* **49**, 1154–1165 (2009).
151. Rugo, K. F., Tamler, K. N., Woodman, G. F. & Maxcey, A. M. Recognition-induced forgetting of faces in visual long-term memory. *Attent. Percept. Psychophys.* **79**, 1878–1885 (2017).
152. Lin, P.-H. & Luck, S. J. Proactive interference does not meaningfully distort visual working memory capacity estimates in the canonical change detection task. *Front. Psychol.* **3**, 42 (2012).
153. Olson, I. R. & Jiang, Y. Visual short-term memory is not improved by training. *Mem. Cogn.* **32**, 1326–1332 (2004).
154. Lu, X., Dai, A., Guo, Y., Shen, M. & Gao, Z. Is the social chunking of agent actions in working memory resource-demanding? *Cognition* **229**, 105249 (2022).
155. Snow, J. C. & Culham, J. C. The treachery of images: how realism influences brain and behavior. *Trends Cogn. Sci.* **25**, 506–519 (2021).
156. Vestner, T., Over, H., Gray, K. L. & Cook, R. Objects that direct visuospatial attention produce the search advantage for facing dyads. *J. Exp. Psychol. Gen.* **151**, 161–171 (2022).
157. Brady, T. F. et al. Scaling up visual attention and visual working memory to the real world. In *Psychology of Learning and Motivation* Vol. 70, 29–69 (Elsevier, 2019).
158. Wiseman, S. & Neisser, U. Perceptual organization as a determinant of visual recognition memory. *Am. J. Psychol.* **87**, 675–681 (1974).
159. Asp, I. E., Störmer, V. S. & Brady, T. F. Greater visual working memory capacity for visually matched stimuli when they are perceived as meaningful. *J. Cogn. Neurosci.* **33**, 902–918 (2021).
160. Ngiam, W. X. Q., Brissenden, J. A. & Awh, E. “Memory compression” effects in visual working memory are contingent on explicit long-term memory. *J. Exp. Psychol. Gen.* **148**, 1373–1385 (2019).
161. Starr, A., Srinivasan, M. & Bunge, S. A. Semantic knowledge influences visual working memory in adults and children. *PLoS ONE* **15**, e0241110 (2020).
162. Moore, C. D., Cohen, M. X. & Ranganath, C. Neural mechanisms of expert skills in visual working memory. *J. Neurosci.* **26**, 11187–11196 (2006).
163. Xie, W. & Zhang, W. Familiarity increases the number of remembered Pokémon in visual short-term memory. *Mem. Cogn.* **45**, 677–689 (2017).
164. Brady, T. F., Störmer, V. S. & Alvarez, G. A. Working memory is not fixed-capacity: more active storage capacity for real-world objects than for simple stimuli. *Proc. Natl Acad. Sci. USA* **113**, 7459–7464 (2016).
165. Jackson, M. C. & Raymond, J. E. Familiarity enhances visual working memory for faces. *J. Exp. Psychol. Hum. Percept. Perform.* **34**, 556–568 (2008).

166. O'Donnell, R. E., Clement, A. & Brockmole, J. R. Semantic and functional relationships among objects increase the capacity of visual working memory. *J. Exp. Psychol. Learn. Mem. Cogn.* **44**, 1151–1158 (2018).
167. Nosofsky, R. M. Tests of an exemplar model for relating perceptual classification and recognition memory. *J. Exp. Psychol. Hum. Percept. Perform.* **17**, 3–27 (1991).
168. Wyble, B., Swan, G. & Callahan-Flintoft, C. Measuring visual memory in its native format. *Trends Cogn. Sci.* **20**, 790–791 (2016).
169. Sahar, T., Sidi, Y. & Makovski, T. A metacognitive perspective of visual working memory with rich complex objects. *Front. Psychol.* **11**, 179 (2020).
170. Stojanoski, B., Emrich, S. M. & Cusack, R. Representation of semantic information in ventral areas during encoding is associated with improved visual short-term memory. Preprint at *bioRxiv* <https://doi.org/10.1101/2019.12.13.875542> (2020).
171. Allen, M. G., Destefano, I. & Brady, T. F. Chunks are not 'content-free': hierarchical representations preserve perceptual detail within chunks. In *Proc. Ann. Meet. Cogn. Sci. Soc.* **43**, 721–727 (2021).
172. Mathy, F. & Feldman, J. What's magic about magic numbers? Chunking and data compression in short-term memory. *Cognition* **122**, 346–362 (2012).
173. Wood, J. N. Visual working memory for observed actions. *J. Exp. Psychol. Gen.* **136**, 639–652 (2007).
174. Shen, M., Gao, Z., Ding, X., Zhou, B. & Huang, X. Holding biological motion information in working memory. *J. Exp. Psychol. Hum. Percept. Perform.* **40**, 1332–1345 (2014).
175. Craik, F. I. Levels of processing: past, present... and future? *Memory* **10**, 305–318 (2002).
176. Craik, F. I. & Lockhart, R. S. Levels of processing: a framework for memory research. *J. Verbal Learn. Verbal Behav.* **11**, 671–684 (1972).
177. Bradshaw, G. L. & Anderson, J. R. Elaborative encoding as an explanation of levels of processing. *J. Verbal Learn. Verbal Behav.* **21**, 165–174 (1982).
178. Nairne, J. S. Remembering over the short-term. *Annu. Rev. Psychol.* **53**, 53–81 (2002).
179. Nelson, D. L., Cermak, L. & Craik, F. Remembering pictures and words: appearance, significance and name. In *Levels of Processing in Human Memory* 45–76 (Taylor & Francis, 1979).
180. Tulving, E. & Thomson, D. M. Encoding specificity and retrieval processes in episodic memory. *Psychol. Rev.* **80**, 352–373 (1973).
181. Hedayati, S., O'Donnell, R. E. & Wyble, B. A model of working memory for latent representations. *Nat. Hum. Behav.* **6**, 709–719 (2022).
182. van Kerkoerle, T., Self, M. W. & Roelfsema, P. R. Layer-specificity in the effects of attention and working memory on activity in primary visual cortex. *Nat. Commun.* **8**, 13804 (2017).
183. Self, M. W., van Kerkoerle, T., Goebel, R. & Roelfsema, P. R. Benchmarking laminar fMRI: neuronal spiking and synaptic activity during top-down and bottom-up processing in the different layers of cortex. *NeuroImage* **197**, 806–817 (2019).
184. Maxcey, A. M. & Woodman, G. F. Forgetting induced by recognition of visual images. *Vis. Cogn.* **22**, 789–808 (2014).
185. Konkle, T. & Alvarez, G. Deepnets do not need category supervision to predict visual system responses to objects. *J. Vis.* **20**, 498–498 (2020).
186. Naspi, L., Hoffman, P., Devereux, B. & Morcom, A. M. Perceptual and semantic representations at encoding contribute to true and false recognition of objects. *J. Neurosci.* **41**, 8375–8389 (2021).
187. Sheng, J. et al. Higher-dimensional neural representations predict better episodic memory. *Sci. Adv.* **8**, eabm3829 (2022).
188. Kahana, M. J. *Foundations of Human Memory* (Oxford Univ. Press, 2012).
189. Rust, N. C. & Mehrpour, V. Understanding image memorability. *Trends Cogn. Sci.* **24**, 557–568 (2020).
190. Hemmer, P. & Steyvers, M. A Bayesian account of reconstructive memory. *Top. Cogn. Sci.* **1**, 189–202 (2009).
191. Huttenlocher, J., Hedges, L. V. & Vevea, J. L. Why do categories affect stimulus judgment? *J. Exp. Psychol. Gen.* **129**, 220–241 (2000).
192. Hemmer, P. & Steyvers, M. Integrating episodic memories and prior knowledge at multiple levels of abstraction. *Psychon. Bull. Rev.* **16**, 80–87 (2009).
193. Bae, G.-Y., Olkkonen, M., Allred, S. R. & Flombaum, J. I. Why some colors appear more memorable than others: a model combining categories and particulars in color working memory. *J. Exp. Psychol. Gen.* **144**, 744–763 (2015).
194. Destefano, I., Brady, T. & Vul, E. Predicting memory errors with a Bayesian model of concept generalization. In *Proc. Ann. Meeting Cogn. Sci. Soc.* **43**, 1760–1766 (2021).
195. Brady, T. F., Schacter, D. L. & Alvarez, G. A. The adaptive nature of false memories is revealed by gist-based distortion of true memories. *J. Vis.* **15**, 948 (2015).
196. Bruning, A. L. & Lewis-Peacock, J. A. Long-term memory guides resource allocation in working memory. *Sci. Rep.* **10**, 22161 (2020).
197. Dubé, C. Central tendency representation and exemplar matching in visual short-term memory. *Mem. Cogn.* **47**, 589–602 (2019).
198. Orhan, A. E. & Jacobs, R. A. A probabilistic clustering theory of the organization of visual short-term memory. *Psychol. Rev.* **120**, 297 (2013).
199. Brady, T. F. & Alvarez, G. A. No evidence for a fixed object limit in working memory: spatial ensemble representations inflate estimates of working memory capacity for complex objects. *J. Exp. Psychol. Learn. Mem. Cogn.* **41**, 921–929 (2015).
200. Hollingworth, A. & Henderson, J. M. Testing a conceptual locus. *Mem. Cogn.* **31**, 930–940 (2003).
201. Oliva, A. In *Neurobiology of Attention* (eds Itti, L. et al.) 251–256 (Elsevier, 2005).
202. Friedman, A. Framing pictures: the role of knowledge in automatized encoding and memory for gist. *J. Exp. Psychol. Gen.* **108**, 316–355 (1979).
203. Hollingworth, A. & Henderson, J. M. Semantic informativeness mediates the detection of changes in natural scenes. *Vis. Cogn.* **7**, 213–235 (2000).
204. Miller, M. B. & Gazzaniga, M. S. Creating false memories for visual scenes. *Neuropsychologia* **36**, 513–520 (1998).
205. Brewer, W. F. & Treyns, J. C. Role of schemata in memory for places. *Cogn. Psychol.* **13**, 207–230 (1981).
206. Lampinen, J. M., Copeland, S. M. & Neuschatz, J. S. Recollections of things schematic: room schemas revisited. *J. Exp. Psychol. Learn. Mem. Cogn.* **27**, 1211–1222 (2001).
207. Schurgin, M. W. & Brady, T. F. When "capacity" changes with set size: ensemble representations support the detection of across-category changes in visual working memory. *J. Vis.* **19**, 3 (2019).
208. Moscovitch, M., Cabeza, R., Winocur, G. & Nadel, L. Episodic memory and beyond: the hippocampus and neocortex in transformation. *Annu. Rev. Psychol.* **67**, 105–134 (2016).
209. Moscovitch, M., Nadel, L., Winocur, G., Gilboa, A. & Rosenbaum, R. S. The cognitive neuroscience of remote episodic, semantic and spatial memory. *Curr. Opin. Neurobiol.* **16**, 179–190 (2006).
210. Jaiswal, S., Fernando, B. & Tan, C. TDAM: top-down attention module for contextually guided feature selection in CNNs. In *Proc. Computer Vision—ECCV 2022: 17th Eur. Conf. XXV* 259–276 (Springer, 2022).
211. Bates, C. J., Alvarez, G. & Gershman, S. J. Scaling models of visual working memory to natural images. Preprint at *bioRxiv* <https://doi.org/10.1101/2023.03.17.533050> (2023).
212. Lake, B. M., Salakhutdinov, R. & Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science* **350**, 1332–1338 (2015).
213. Cleary, A. M. Dependent measures in memory research: from free recall to recognition. In *Handbook of Research Methods in Human Memory* 19–35 (Routledge, 2018).
214. Harlow, I. M. & Yonelinas, A. P. Distinguishing between the success and precision of recollection. *Memory* **24**, 114–127 (2016).
215. Nilakantan, A. S., Bridge, D. J., VanHaerents, S. & Voss, J. L. Distinguishing the precision of spatial recollection from its success: evidence from healthy aging and unilateral mesial temporal lobe resection. *Neuropsychologia* **119**, 101–106 (2018).
216. Sutterer, D. W. & Awh, E. Retrieval practice enhances the accessibility but not the quality of memory. *Psychon. Bull. Rev.* **23**, 831–841 (2016).
217. Pratte, M. S. Swap errors in spatial working memory are guesses. *Psychon. Bull. Rev.* **26**, 958–966 (2018).
218. Konkle, T. & Oliva, A. A familiar-size Stroop effect: real-world size is an automatic property of object representation. *J. Exp. Psychol. Hum. Percept. Perform.* **38**, 561–569 (2012).
219. Pouget, A., Dayan, P. & Zemel, R. Information processing with population codes. *Nat. Rev. Neurosci.* **1**, 125–132 (2000).
220. Rust, N. C. & Cohen, M. R. Priority coding in the visual system. *Nat. Rev. Neurosci.* **23**, 376–388 (2022).
221. Yonelinas, A. P. Receiver-operating characteristics in recognition memory: evidence for a dual-process model. *J. Exp. Psychol. Learn. Mem. Cogn.* **20**, 1341–1354 (1994).
222. Marr, D. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (MIT Press, 2010).
223. van Bergen, R. S. & Jehee, J. F. Probabilistic representation in human visual cortex reflects uncertainty in serial decisions. *J. Neurosci.* **39**, 8164–8176 (2019).
224. Yeon, J. & Rahnev, D. The suboptimality of perceptual decision making with multiple alternatives. *Nat. Commun.* **11**, 3857 (2020).
225. Rahnev, D., Block, N., Denison, R. N. & Jehee, J. Is perception probabilistic? Clarifying the definitions. Preprint at *PsyArXiv* <https://doi.org/10.31234/osf.io/f8v5r> (2021).
226. Bays, P. M., Catalao, R. F. G. & Husain, M. The precision of visual working memory is set by allocation of a shared resource. *J. Vis.* **9**, 7 (2009).
227. Shiffrin, R. M. & Steyvers, M. A model for recognition memory: REM — retrieving effectively from memory. *Psychon. Bull. Rev.* **4**, 145–166 (1997).

Acknowledgements

The authors thank J. Wixted and V. Störmer for helpful discussions.

Author contributions

All authors contributed equally to aspects of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Peer review information *Nature Reviews Psychology* thanks Zaifeng Gao, Robert Logie, and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© Springer Nature America, Inc. 2024