Priority coding in the visual system

Nicole C. Rust[1]✉ *and Marlene R. Cohen*[2,3]

Abstract | Although we are continuously bombarded with visual input, only a fraction of incoming visual events is perceived, remembered or acted on. The neural underpinnings of various forms of visual priority coding, including perceptual expertise, goal-directed attention, visual salience, image memorability and preferential looking, have been studied. Here, we synthesize information from these different examples to review recent developments in our understanding of visual priority coding and its neural correlates, with a focus on the role of behaviour to evaluate candidate correlates. We propose that the brain combines different types of priority into a unified priority signal while also retaining the ability to differentiate between them, and that this happens by leveraging partially overlapping low-dimensional neural subspaces for each type of priority that are shared with the downstream neural populations involved in decision-making. Finally, we describe the gulfs in understanding that have resulted from different research approaches, and we point towards future directions that will lead to fundamental insights about neural coding and how prioritization influences visually guided behaviours.

Only a subset of the visual information that falls on our retinae is perceived or remembered, or affects behaviour. Many visual and cognitive factors combine to determine what visual events are prioritized. To what degree are they reflected via common neural correlates? How are they combined and how are they distinguished?

In this Review, we highlight similarities and differences in how distinct forms of visual priority coding affect several metrics of neural population activity at different stages of visual processing. By definition, all forms of priority coding have similar behavioural effects, to the extent that they prioritize some stimuli over others. However, they are in many ways diverse, operating on timescales that range from tens of milliseconds to months, depending on factors that range from properties of images to those that are cognitive in origin, and operating on a range of visual features from low-level features, such as orientation, to high-level concepts such as object category.

In principle, each form of prioritization could be associated with a unique neural correlate. There are some advantages to such a scheme, including ease in distinguishing whether a change in neural representation came from learning or salience, for example. Indeed, there are notable differences in the neural correlates of different forms of priority coding. However, perhaps more remarkable are their similarities: the vast number of studies searching for the neural correlates of different forms of priority have uncovered only a modest number of neural correlates. This leads to the tantalizing hypothesis that a small number of mechanisms might underlie a broad range of behaviours guided by prioritized stimuli.

The goal of this Review is to determine whether the similarity between the behavioural effects of different forms of prioritization reflects generalized principles for priority coding that are also reflected as commonalities across neural coding schemes, neural circuits and/or computational mechanisms. To answer these questions, we highlight insights from new data analysis methods that focus on subspaces of neural population activity and that have the potential to uncover general relationships between priority coding, holistic measures of neuronal population activity and behaviour. We propose that the visual system uses partially overlapping subspaces of neural activity to combine different types of priority into a unified priority signal while also retaining the ability to differentiate between them. We describe how this mechanism may enable different priority signals to support flexible behaviour through flexible, functional communication with downstream neurons involved in decision-making. Finally, we conclude by describing what we regard as the next important steps towards understanding how the brain prioritizes incoming visual information, including the experiments, data analysis techniques and models through which the behavioural implications of different forms of priority coding on populations of neurons can be compared.

Forms of priority coding

Here, we use the term 'behavioural priority signatures' to refer to the behavioural outcomes of tasks that involve prioritizing some stimuli over others, and we refer to that act as 'prioritization'. Prioritization can be reflected as an increase in behavioural performance, a decrease in

[1]Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA.

[2]Department of Neuroscience, University of Pittsburgh, Pittsburgh, PA, USA.

[3]Center for the Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, PA, USA.

✉e-mail: nrust@psych.upenn.edu

https://doi.org/10.1038/s41583-022-00582-9

reaction time and/or an increase in the time spent engaging with a stimulus. Prioritization is crucial in all species and systems, ranging from honeybees, which selectively respond to visual stimuli that are both salient and associated with reward[1], to elephants, which learn to imitate the voices of the elephants or humans most important to them[2]. The ubiquity of these behaviours has led to much discussion about whether different forms of prioritization should be thought of as part of the same category or differentiated. Rather than focusing on definitions of terms or categories, we focus here on whether different behavioural priority signatures are mediated by overlapping neural correlates. In terms of neurophysiological investigation coupled with behaviour, the best-studied behavioural priority signatures concern the visual system of non-human primates. Therefore, we highlight five example behavioural priority signatures in the primate visual system selected on the basis of their depth of explanation and their range of investigation across brain areas.

The visual cortex of primates contains a hierarchy of brain areas. In earlier visual areas, beginning with the primary visual cortex (V1), receptive fields are small and neurons respond primarily to low-level features such as orientation, colour or depth[3]. From there, response selectivity in the ventral stream progresses to mid-level features (such as texture and curvature in V4)[4] and to higher-level properties that enable object recognition in the inferotemporal cortex (IT)[5].

The first three behavioural priority signatures that we consider — perceptual learning, goal-directed visual attention and visual salience — are often studied at earlier stages of the visual system, such as in V1, V2 and V4, because their spatial specificity suggests that they are mediated by neurons with small receptive fields such as those in these areas. However, all three have been associated with modulations in every visual area in which they have been studied. Perceptual learning improves observers' ability to discriminate specific, well-practised stimuli over periods of weeks or months, leading to a prioritization of practised stimuli at practised locations in visual space[6–21] (FIG. 1a). Goal-directed visual attention affects the activity of neurons throughout the visual pathway and enables observers to flexibly prioritize different locations, features, objects or other parts of a visual scene and, critically, to filter out task-irrelevant information (for reviews, see REFS[22–26]) (FIG. 1b). Visual salience enables certain regions of complex images to be prioritized because properties of these regions are intrinsically attention-grabbing. Psychophysically, salience is commonly studied in three ways. The first is by comparing the precepts of stimuli that do and do not 'pop out', depending on whether they differ from the background in visual features[7–29] or in onset time[30,31] and can therefore be prioritized through bottom-up mechanisms, or whether they fit a search template[29,32,33] (FIG. 1c). The second way to study visual salience processing is using visual search tasks in which observers search for targets that vary from distractors by different amounts[29,34,35]. Third, free-viewing paradigms can be used in which the salience of an image is quantified on the basis of the regularity with which patterns of gaze fixations are consistent across participants[36].

The next two behavioural priority signatures that we examine — image memorability and visual novelty — are most often linked to processing at higher stages of the visual system (such as IT) because they are associated with higher-level, object-centric properties and rapid, long-lasting plasticity. Image memorability enables some stimuli to be prioritized because properties of those images are intrinsically more easily remembered. For example, in a visual recognition memory task, participants are asked to view novel and repeated images and report whether they have seen them before (FIG. 1d). In these studies, some images are found to be more memorable than others[37–44]. Investigations of the properties that affect image memorability have implicated higher-level attributes more strongly than low-level features, including object type (for example, faces are typically more memorable than nature scenes) and object atypicality (an unusual chair, shaped like a hand, tends to be memorable)[37,40]. Images are also prioritized when they are novel, leading to a preference by human infants, monkeys and other animals to look at images that are unlike those encountered before, called preferential looking[45–47] (FIG. 1e). In addition to this implicit report of visual novelty, novelty can be measured by explicit reports of whether an image is repeated or novel in a visual recognition memory task[48,49]. Whereas image memorability is associated with the prioritization of more memorable images across a set of (all novel) images, novelty measured by preferential looking refers to the prioritization of novel over repeated stimuli.

Neural correlates of priority coding

Despite their differences, the behaviours depicted in FIG. 1 all involve prioritizing a subset of visual information. Several non-mutually exclusive neural correlates have been put forward for how priority coding might manifest in the brain (FIG. 2). Here, we define a neural correlate as a manifestation of neural population activity that, when decoded, could result in stimulus prioritization.

Proposals for priority coding. The simplest proposed neural correlate underlying priority coding is magnitude variation, whereby a particular image evokes more vigorous population responses when it is prioritized[31,50–53] (FIG. 2a). The second proposal is reflected as a change in the magnitude or structure of correlated trial variability (also termed noise correlations, spike count correlations or r_{SC}), which reflects shared variability in responses to repeated presentations of the same stimulus, and could change the fidelity with which information about a prioritized object can be decoded by downstream neurons[54] (FIG. 2b). The third proposal is mathematically similar to the second but arises from a different source. This proposal is restricted to tasks that require grouping multiple stimuli together to extract a parameter of interest while disregarding other differences between them, for example, object identification, which can be thought of as grouping together all the images containing the same object while ignoring variation in the details with which those objects appear such as their spatial position, size and background context. In tasks like these, 'nuisance' variability refers to the

Receptive fields
The restricted region of visual space within which changes in the visual stimulus lead to changes in a neuron's firing rate response.

Ventral stream
Also called the 'form processing' or 'what are you looking at?' pathway owing to its association with object identification. Includes primate visual brain areas of visual cortex V1, V2, and V4 and the inferotemporal cortex.

Search template
Defines the combined set of features that are sought in a visual search task.

Population responses
Snapshots of the spiking activity of a collection of individual neurons in response to a single trial in one experimental condition.

Trial variability
Variability in the responses of an individual neuron across repeated instances of the same experimental conditions and visual stimulus.

Noise correlations
The degree to which trial variability is correlated between different units in response to repeated presentations of the same visual stimuli and other experimental conditions.

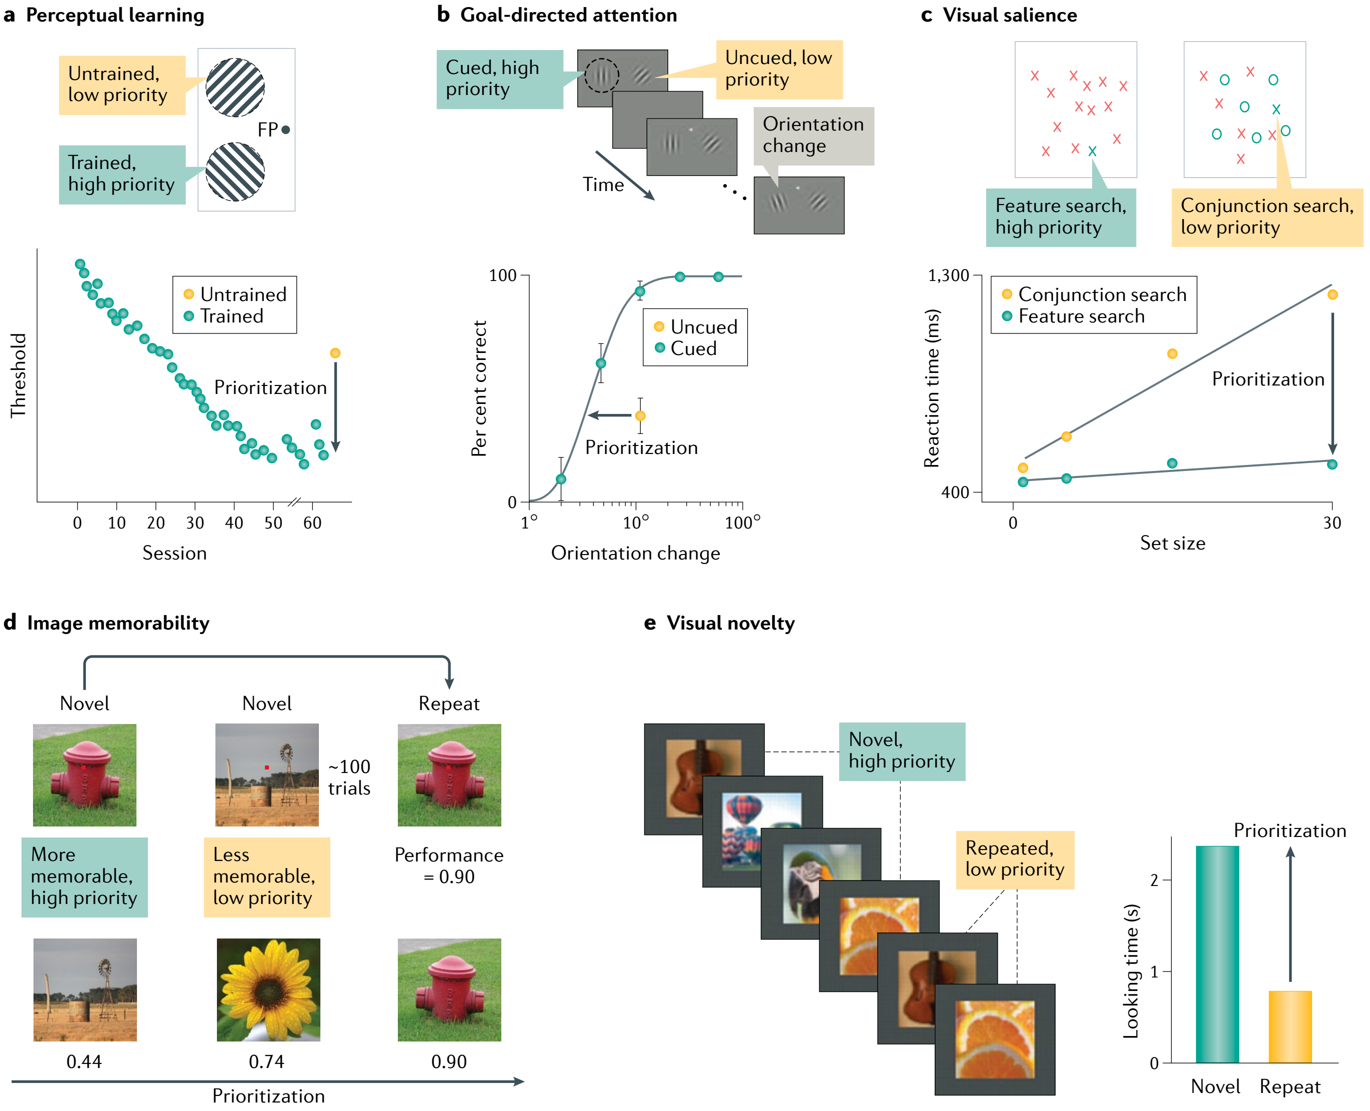

Fig. 1 | **Forms of priority coding.** Each panel depicts one behavioural prioritization signature illustrated with one representative psychophysical example of prioritization. **a** | Perceptual learning involves extensive training on a task such as discriminating subtly different stimulus orientations at a specific location (high priority). Performance improvements do not generalize to untrained locations (low priority). The plot shows how the psychophysical thresholds at the trained location (green circles) decrease as a function of training session number and are largely retained following breaks in training (hash marks). By contrast, discrimination thresholds at an untrained location (yellow circle) remain high. **b** | Goal-directed attention is studied using a Posner cuing paradigm, which includes two repeated, flashing stimuli and a cue that indicates the location at which to expect a stimulus change (high priority). Changes also occur at the uncued location (low priority) on a small fraction of trials. Shown is a psychometric curve, which is shifted to the left for stimuli at the cued location, indicating better performance at detecting changes there than for the uncued location. **c** | Visual salience can be studied by determining how long it takes for participants to find a stimulus (a green cross) embedded in a background of distractors. Stimuli with high salience differ by only one feature and pop out (high priority), and search time remains independent of set size. Stimuli with low salience differ by one of multiple features (low priority), and search time increases as a function of the number of distractors. **d** | Image memorability can be studied in a recognition memory paradigm in which participants view one image per trial and report whether it is novel or repeated. Image memorability is scored 0–1 and corresponds to the fraction of participants that remember seeing an image when it is repeated. Images with high memorability (high priority) are better remembered than images with low memorability[37,142] (low priority). **e** | Visual novelty can be studied as the amount of time that a monkey spends free-viewing novel images (high priority) as compared to repeated images (low priority). FP, fixed point. Panel **a** adapted from REF.[70], Springer Nature Limited. Panel **b** adapted from REF.[90], Springer Nature Limited. Panel **c** adapted with permission from REF.[27], Elsevier. Panel **d** adapted from REF.[142], CC BY 4.0 (https://creativecommons.org/licenses/by/4.0/). Panel **e** adapted with permission from REF.[45], National Academy of Sciences.

dispersion of neural responses across different stimuli in a group. The third proposal is reflected as changes in correlated nuisance variability, which reflects correlations between neurons in their responses to different stimuli that constitute noise in a grouping task, or equivalently, stimulus-to-stimulus response variability[55] (FIG. 2c).

A fourth proposal differs from the others in that it operates by changing the way that a fixed neural representation is read out or decoded by a downstream population, whereas the others operate by changing neural population activity in the brain area that represents the relevant visual feature or object[56] (FIG. 2d). A fifth proposal suggests that

Nuisance variability
The spiking variability induced by parameters not relevant to a task such as object position or size in a task that requires extraction of, for example, object identity.

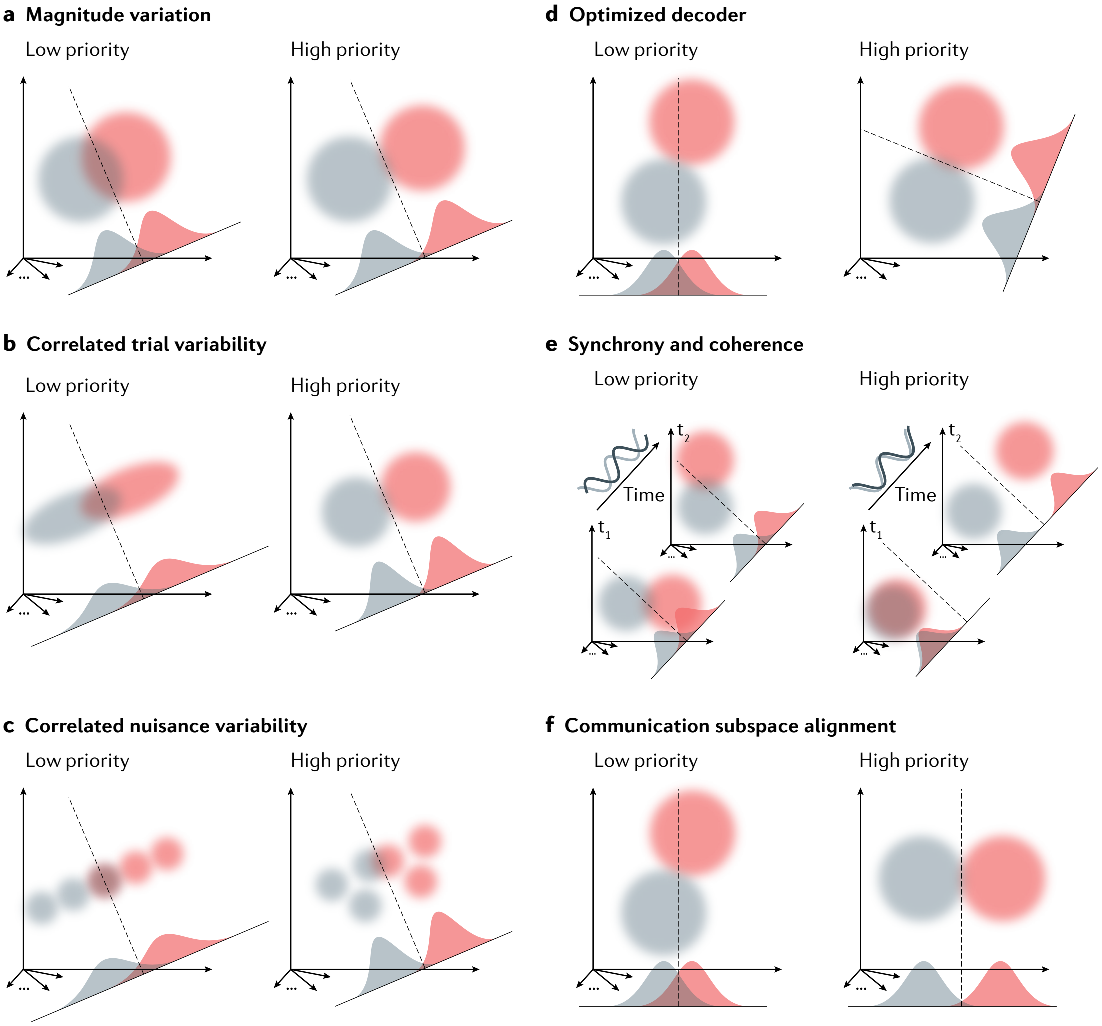

Fig. 2 | **Proposals for priority coding.** Each panel depicts one proposed neural correlate of priority coding, shown in the context of a task that requires a two-way classification (such as 'is this image an A or a B?', where A and B can have different priority statuses, such as high-priority attended and low-priority unattended stimuli). Each plot depicts a subset of two dimensions that exist within a high-dimensional space where each axis is the response of one neuron. Shown are hypothetical 2D probability distributions, with each colour representing one class of responses. Also shown are projections along a linear-decoding axis (indicated by the dotted line) perpendicular to the decoding axis; the decoding axis also reflects the threshold for the classification. All six proposals operate by increasing the ability to discriminate between stimulus A and B by increasing the separation between the grey and red distributions in the high-priority versus low-priority condition. What differs between the proposals is how this end point is achieved. These proposals are not mutually exclusive and multiple types may contribute to any one priority coding behaviour. **a** | Magnitude coding: prioritization follows from a more vigorous population response to that stimulus. **b** | Correlated trial variability: prioritization follows from a change in the magnitude or structure of the variability shared across neurons in response to repeated presentations of the same stimulus in a manner that affects how information is decoded downstream. **c** | Correlated nuisance variability: similar to correlated trial variability, but prioritization results from a change in correlational structure between neurons in their responses to different stimuli. **d** | Optimized decoder: prioritization follows from increasing the efficacy of a downstream decoder as opposed to changing the population response itself. **e** | Synchrony and coherence: prioritization follows changing synchronous activity on very fast timescales, possibly through sharing of oscillatory activity. **f** | Communication subspace alignment: prioritization follows from a reorganization in the population response in a manner that maintains the same total information about a stimulus but changes how that information aligns with the communication subspace.

Synchrony
Simultaneous activation and/or inactivation of different neurons on fast timescales (less than 10 ms). Often measured as the coherence between activity in two areas in a particular temporal frequency band.

prioritization is modulated through changes in synchrony (sometimes inferred by measuring coherence) between different subgroups of neurons in a population on very fast time scales, potentially mediated by oscillatory activity that is shared between populations[57,58] (FIG. 2e).

Most recently, a sixth overlapping but conceptually distinct framework for thinking about priority coding has emerged around the subspaces of neural population space, which is defined as the space in which the response of each simultaneously recorded neuron represents one dimension[59–61]. As a simple illustration, one can envision a population of two neurons that are exact copies of one another and, as a consequence of their overlap, the responses of these two neurons occupy a

Weights

Used to determine the output of a linear decoder, computed as a weighted sum of the population response on a single trial (for example, output = weight 1 × neuron 1 response + weight 2 × neuron 2 response…).

dimensionality of one as opposed to two; for example, no information would be lost if one of those neurons were eliminated from the population. In this example, the perfectly correlated activity of the two neurons reflects the encoding subspace (weights of (1,1)), and the 'null space' not occupied by the population is defined by the perpendicular direction, in which the two neurons are anticorrelated (weights of (1,–1)). As the numbers of neurons in the population grows, the principle remains similar, and the subspace dimensionality is defined by the different types of activity that are elicited by the population.

These same ideas are used not only to discuss the visual information encoded by a population of neurons but also the information communicated from one brain area to the next. Indeed, simultaneous recordings from multiple neurons in multiple brain areas suggest that the information conveyed from one area to the next is a subset of the total information reflected in that brain area[59,60,62,63]. The communicated information is described as existing in a 'communication subspace', whereas information that is not communicated is described as existing in the remaining null or private space. The communication subspace is influenced not only by anatomy but also by the weighting of inputs from one brain area to the next. This idea can be extended to account for priority coding, leading to a sixth proposed neural correlate: that priority coding is achieved by affecting the extent to which prioritized visual information aligns with the communication subspace, while maintaining the same total information about a stimulus (FIG. 2f). This proposal can be distinguished from the alternative that is implied by many, although not all, of the other proposals, in which priority coding operates by modulating the total information about a stimulus in the visual cortex.

When considering these different proposals, it is important to understand that the same behaviour may manifest in different ways in different brain areas. For example, magnitude coding in a given brain area is consistent with but does not necessarily imply the decoded transformation of any of the other proposals reflected in the input to that brain area. Therefore, it is important to measure neural activity at multiple levels of the visual hierarchy when investigating the neural correlates of priority coding.

Challenges in comparing types of priority coding. The most obvious challenge in drawing comparisons between the different forms of priority coding in FIG. 1 is that, because of varying traditions and concerns in distinct subfields, they have typically been investigated very differently. For example, if an experiment does not measure correlated variability in a particular brain area (FIG. 2b), it will not identify it as a neural correlate even if it exists. Very few studies[64–66] have investigated correlated trial variability at higher stages of the ventral pathway (such as in the IT). Similarly, very few studies[67,68] of earlier stages of the ventral visual pathway have investigated nuisance variation (FIG. 2c) in the context of grouping tasks that require a participant to abstract a variable of interest (such as object identity) across different images. Given the differences in the way various forms of priority coding have been studied, the similarities in the neural correlates of the various forms of priority coding reported thus far are typically more informative than the differences.

Behaviour as a constraint for evaluating candidate priority-coding schemes. A remarkably high number of neural coding schemes have been implicated in priority coding (FIG. 2). How do we move beyond the myriad of observations that prioritization during behaviour X corresponds to a change in candidate neural correlate Y? That is, how do we distinguish putative neural correlates that are possible candidates for having causal links to behaviour from those that are epiphenomena? Building on the classic proposals of Parker and Newsome[69], modern insight has been gained through focusing on behaviour. Although all forms of priority coding are, by definition, related to behaviour in at least a pairwise manner (for example, more of both behaviour X and neural correlate Y occur in condition A than in condition B), a more powerful demonstration would involve evaluating the degree to which systematic variation in behaviour correlates with systematic variation in candidate neural correlates. Focusing on behaviour has been important for progress in understanding several of the forms of priority coding that are highlighted in FIG. 1.

In the case of perceptual learning, a compelling link between changes in the properties of individual neurons and the prioritization that happens with learning was difficult to come by for quite some time. Studies that focused on single neurons in the early visual cortex reported diverse effects that sometimes included increases in overall mean firing rate in response to the prioritized stimulus and various other neural correlates in many areas of the visual cortex. However, these studies did not identify a systematic relationship between these changes and the behavioural improvements with learning[56,70–79] (although see REF.[80]). One study, focused on motion perception, made an illuminating step forward by reporting that, although there were no detectable changes in the activity of individual neurons in the visual cortex (the middle temporal area) with learning, there were changes in the response of neurons in a downstream projection area implicated in decision-making (the lateral intraparietal cortex) that systematically reflected the amount of prioritization during learning[56]. These results suggest that prioritization happens by optimizing the ability of lateral intraparietal cortex neurons to preferentially weigh and thus decode inputs from the most sensitive middle temporal area neurons (FIG. 2d). More recent investigations recording populations of neurons have found correlates of priority coding within the visual cortex itself (typically mid-level areas such as V4) reflected as changes in the magnitude of correlated trial variability between neurons (FIG. 2b). The magnitude of correlated trial variability in the medial superior temporal area or area V4 was reduced systematically with learning[74,81], and the magnitude of correlated variability was strongly related to behavioural performance on a day-to-day basis[74]. An exciting possibility is that changes in correlated variability might lead to changes in the way inputs from the visual cortex are read out by association areas[82,83].

Goal-directed visual attention is perhaps the most extensively studied form of priority coding, and it has been associated with changes in a staggering variety of measures of neural activity in almost every visual and visuomotor brain area, including trial-averaged firing rates[50,52,53,84–88] (FIG. 2a), shared variability between pairs of neurons (FIG. 2b) in the same brain area[74,89–103] and in different brain areas[98,104–106], and synchrony, which is sometimes measured as coherence[57,58,106–124] (FIG. 2d). The number of studies that have directly linked those changes with behaviour is much smaller. Given that visually guided behaviours are thought to be mediated by large numbers of neurons in multiple brain areas, it is perhaps unsurprising that the measures that are most reliably associated with behaviour are those that encompass the population (such as the magnitude of correlated variability or synchrony[74,125,126]). Below, we return to the topic of how these results relate to new efforts to understand how priority coding changes the subspaces of neural activity in which visual stimuli are encoded and communicated (FIG. 2f).

Magnitude coding (FIG. 2a) has been implicated in the other priority coding behaviours depicted in FIG. 1. Visually salient, pop-out stimuli modulate the magnitude of responses in the early visual cortex[127,128] (but see REF.[129]) as does visual search, which involves a confluence of bottom-up processes ('what are you looking at?') and top-down processes ('what are you looking for?')[107,109,130–136]. A recent study[137] used a causal manipulation to demonstrate that reversible inactivation of the parietal cortex reduces salience coding in prefrontal neurons and reduces the ability to predict eye movements in a free-viewing paradigm from models that predict the salience of images[138–141].

Magnitude variation has also been linked to image memorability: more memorable images are reported to evoke IT population responses that are about 20% larger than those evoked by less memorable images, and continuous changes in magnitude coding strongly correlate with continuous changes in image memorability[142]. Similarly, image novelty is also reflected in the IT through a magnitude coding scheme as novel stimuli evoke a more vigorous IT response than do familiar stimuli[48,51,143–145]. Although this phenomenon is often referred to as 'repetition suppression' to emphasize its relationship with familiarity memory, it could also be titled 'novelty enhancement' to emphasize its relationship with priority coding (although whether the underlying neural mechanisms that shape this phenomenon are enhancing or suppressive remains unclear). Indeed, the extent of this phenomenon in the hippocampus predicts continuous changes in novelty behaviour assessed by preferential looking[45] and, over timescales of minutes, its magnitude aligns with behavioural reports of novelty versus familiarity[48]. Consequently, it too passes the benchmark of accounting for continuous behavioural variation.

Synchrony on the timescale of a small number of milliseconds (sometimes measured as coherence) has been implicated in various behavioural priority signatures. It has been most broadly studied in the context of attention: increases in attention (and concomitant changes in performance and reaction time) are associated with enhanced synchrony or coherence[57,58,106–124,126,146]. Enhanced synchrony has also been associated with salience in the context of visual pop-out stimuli[147] and with visual memory[148–153].

In summary, although many neural correlates of priority coding have been identified (FIG. 2), to date, only a few have met the bar of illustrating a systematic relationship with behaviour. A crucial focus of future work should be determining whether other correlates, such as the others highlighted in FIG. 2, can meet this bar and whether multiple behavioural priority signatures are supported by the same underlying mechanism.

A unified yet distinguishable code

Recently, a framework for thinking about neural coding has emerged[62,83,99,154–170] around examining the subspaces of neural population space (defined by the response of each neuron as one dimension) in which visual information is encoded or communicated between brain areas. These ideas have delivered insight on many aspects of neural coding and given rise to fresh ideas related to priority coding.

The first insight is that, in many tasks, most of the relevant population activity exists in a much lower dimensional subspace than the theoretical upper bound. Although the information represented in any given brain area has the potential to have a dimensionality equal to the number of neurons it contains (for example, many millions[5] in primate V1, V2, V4 or IT), emerging evidence suggests that, in various tasks and in response to various stimuli, the brain often uses neuronal population activity that resides in much lower-dimensional subspaces[62,83,99,154–170]. Priority coding seems to affect particularly low-dimensional subspaces. For example, attention seems to modulate trial-to-trial variability along a very small number of dimensions (approximately one[74,171,172]). Similarly, in the IT, the mapping of neural activity to novelty behaviour is linearly decodable (one-dimensional)[48,173].

The second insight is that the subspace in which priority is associated with response modulations seems to be configured such that priority does not dramatically affect representations of visual identity. Representations of identity are thought to be largely reflected as population spike patterns or, equivalently, as population vector direction[5] (FIG. 3a). By contrast, magnitude coding implies a modulation of population vector magnitude and, when priority modulations multiplicatively scale firing rates, they change magnitude without changing vector direction (FIG. 3a). Magnitude modulations are reported to be approximately multiplicative for attention[50,53,174–177] and novelty[178,179] (although see REFS[180,181]).

This insight provides a potential explanation for why overall modulations of the population response can be relatively large yet have minimal impact on choices in perceptual tasks[48,182]. The subspace hypothesis is not unique in hypothesizing a behavioural effect that is not a direct consequence of the magnitude of response

Coherence
A measure of the similarity of oscillatory activity between two brain regions.

Population vector direction
The position of a population response vector in an *N*-dimensional space (where *N* equals the number of neurons) after normalizing for population vector length (or magnitude).

Multiplicatively
Modulations that impact a neuron's response by multiplying it by a factor.

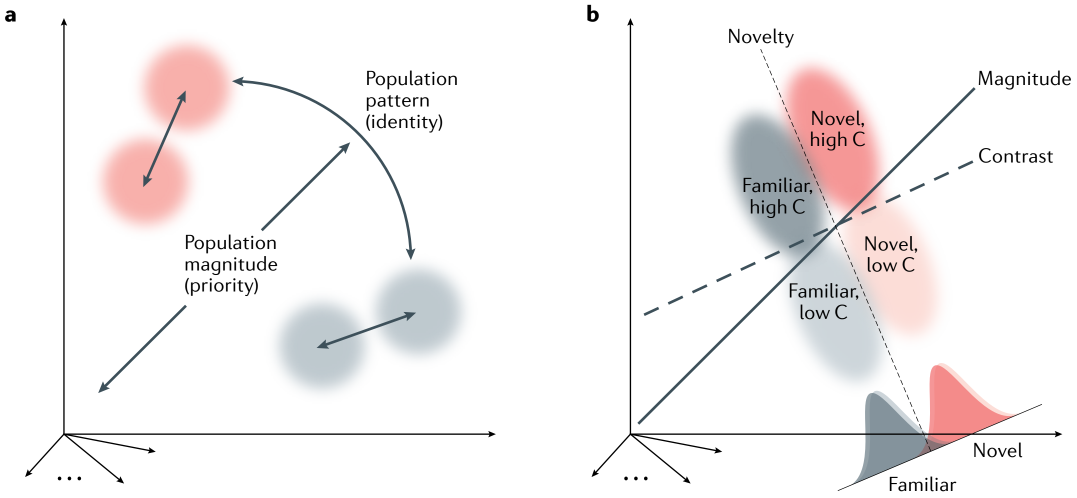

Fig. 3 | **Subspaces for priority coding. a** | Illustration of the complementary coding schemes reflected as changes in population response patterns (to reflect, for example, changes in image or object identity) versus changes in population response magnitude (to reflect, for example, priority). **b** | Schematic of the 2D linear subspace of responses in the inferotemporal cortex recorded while monkeys reported image novelty at the same time as ignoring randomized changes in contrast[173]. Although information about novelty (novel versus familiar) and contrast (C; high versus low) overlaps with the magnitude decoder (that is, the vector with weights of [1,1,1,1...]), a novelty decoder orthogonalized to contrast could account for the mapping of inferotemporal cortex neural responses to the monkeys' behavioural patterns.

modulation[57,58] (although some of the simplest population codes, such as magnitude coding, do not). However, viewing population activity in terms of subspaces that correspond to particular visual or cognitive variables of interest provides a unifying intuition about what types of modulation — those aligned with behaviourally relevant variables — should affect behaviour.

A third insight suggests that different forms of prioritization affect partially overlapping subspaces of neural activity. This insight begins with the observation that many forms of priority coding are associated with changes in the grand mean firing rate[45,48,50–53,84–88,124,127,128,130–136,143–145], which implies that their population representations are at least partially overlapping when using a 'magnitude' linear decoder whose weights are equal for all neurons [1,1,1, ...] (FIG. 3b). This means that the combination of different types of priority can be accessed in a simple way: by monitoring population response magnitude. However, this, in turn, presents a puzzle: the brain must have ways of differentiating diverse forms of priority coding, for example, disambiguating whether an incoming stimulus is prioritized as a consequence of novelty or goal-directed attention — how does it manage to do this?

The answer is that the overlap between the population responses associated with different forms of priority is only partial: optimally weighted linear decoders do not align precisely with the magnitude axis (that is, by affecting all neurons equally) but they are also not orthogonal to it. This is supported by evidence in V4. There, variables such as contrast, adaptation, spatial attention and feature attention all affect magnitude coding; however, at the same time, the neural representations for each of these variables lie in a low-dimensional subspace, and the multiplicative scaling of a given neuron's response by one of those variables only weakly correlates with the extent to which it is modulated by the others[172]. In other words, there is great neuron-to-neuron heterogeneity in the extent to which response magnitude is affected by any particular form of priority coding. As a result, although the subspaces that are engaged by the population responses that are modulated by these different variables are distinct, they are partially overlapping[172].

The suggestion of partial overlap of the different subspaces that are engaged in diverse forms of priority is supported by evidence from a study that compared neural responses in the IT with the ability of monkeys to distinguish changes in image novelty from changes in image contrast[173]. The study revealed that, although changes in both contrast and novelty affected the overall magnitude of the population response, a novelty decoder that classified responses along a dimension that fell orthogonal to the dimension along which population responses changed according to contrast could account for the mapping of IT neural responses to behaviour[173] (FIG. 3b). In principle, all that is required to differentiate diverse priority variables is for their effects on population responses to be at least partially non-overlapping with one another. This would be expected when individual units have heterogeneous sensitivities for the multiple factors that modulate the population response[172,183,184].

As a final insight, functional communication between brain areas might be altered without changing the subspace of activity that is shared between two populations (for example, by changing the nature or amount of information that is aligned with a fixed subspace). One study supported this idea by demonstrating that shifting attention between opposite hemifields has negligible effects on both the amount of information encoded in the visual cortex and the subspace of population activity that related activity in the visual cortex to premotor neurons or the animals' behaviour[99]. Instead, attention acts

Decoder
A single (typically linear) axis in a high-dimensional space, most often created to extract a particular type of information (such as 'is this image an A or a B?') from a neuronal population.

Adaptation
Changes in the response of an individual (behavioural) or neuron with repeated or prolonged exposure to a stimulus.

Units
Individual neurons or groups of a few neurons whose spiking activity is recorded typically via extracellular techniques. Measures of unit activity may or may not reflect the responses of a single neuron.

Linear subspaces
Given a population of *N* neurons that define a population dimensionality with an upper bound of *N*, a linear subspace is a subset of the full space with dimensions $M < N$.

Covariance matrix
Describes the covariation between different neurons across visual stimuli and repeated trials.

Divisive normalization
A model that describes the responses of an individual neuron or population as a combination of the image within its 'classic' receptive field, adjusted by the combined response of other neurons.

primarily to change how neural representations of visual information are aligned with relatively fixed decoding strategies, and it seems to change the information shared between visual and premotor areas without changing the dimensionality of the communication subspace[61,185]. Understanding how changes in the alignment of a population response with the communication subspace manifest at the level of the population and are reflected by the other mechanisms depicted in FIG. 2 is an important avenue for future work.

Putting it all together and allowing for a bit of speculation, the insights described above suggest that the visual system combines different types of priority into a unified priority signal that is strongly related to the population response magnitude (FIG. 3a). At the same time, these different types of priority remain distinguishable from one another as a consequence of their reflection in low-dimensional linear subspaces that are only partially overlapping (FIG. 3b) owing to heterogeneity across individual neurons in their sensitivities to different types of priority modulation. Although priority modulations of population responses seem to be sufficient in size to affect the types of downstream processing that are sensitive to population response magnitude (such as memory), they also seem to be configured to minimize interference with types of representations that are reflected primarily as patterns of spikes (such as object identity) by way of acting approximately (albeit imperfectly) multiplicatively (FIG. 3a). Finally, priority coding associated with cognitive processes that enable flexible behaviour, such as goal-directed attention, may operate by changing the information that is aligned with fixed subspaces (FIG. 2f).

A corollary to the idea that the magnitude axis is preferentially read out to guide at least some forms of visual behaviour is that changes in mean correlated variability or in the shape of the distribution of correlated nuisance variability can have big changes on the accuracy of the information that is communicated to decision-making areas. Many studies have observed that spike count correlations between neurons that encode prioritized information are low[54], even when the correlation changes associated with priority coding have negligible impact on the information that could be gleaned using a decoder optimized for the stimulus and task[54,186–188]. Perhaps the negligible change in the amount of information represented in the population as a whole is adaptive: changing the mean correlated variability (as opposed to more complex, specific or higher-dimensional changes to the structure of the covariance matrix) is a way to change the signal-to-noise ratio of the visual information projected along the magnitude coding axis while preserving currently irrelevant information for future action or memory.

Magnitude coding is clearly not the whole story. Deviations from magnitude coding are plentiful and clearly important. It is those deviations that enable sensory signals to be distinguished from cognitive signals or different forms of priority coding to be distinguished from each other (FIG. 3b). However, the low-dimensional structure of neuronal population activity may impose limits on the number of those features or cognitive processes that can be distinguished. An exciting direction for future work will be to understand the relationship between these limits on neural coding and the perceptual and cognitive limitations of humans and other animals.

Perhaps the greatest contribution of the subspace framework that we have outlined here is its ability to unify many observations and provide intuitions that can guide future experiments. It rarely comes in direct conflict with other theories but, nonetheless, provides distinct insights. For example, the principle of divisive normalization provides an explanation for the scaling of responses associated with several forms of normalization, including visual attention[175,189–191]. The normalization model allows for the response magnitudes of different neurons to be modulated to varying extents by attention. Thus, it implies that attention can change the direction of the population response vector. The complementary contribution of the subspace framework is that it provides powerful geometric intuitions about the quantities of interest (the dimensions along which the modulation occurs) that can be visualized, which in turn facilitates the conceptualization of how the neural correlates of priority coding are connected across different experiments. The insights of the subspace framework are consistent with traditional ways of thinking about population coding, but the subspace framework has already begun to dramatically affect experiment design and new theories about neural coding.

Merging network models

There is a long history of modelling prioritization from diverse perspectives that leverage data from behavioural, physiological and human imaging studies[50,53,175,192–194]. Here, we highlight two largely non-overlapping types of models because we see their merging as a crucial path towards understanding the mechanisms underlying priority coding. The first category uses as its core a long history of biophysically realistic circuit models[82,195–204]. The key constraints on these models are measurements of variability, including the variability of single neurons, changes in shared variability mediated by the balance between excitation and inhibition, noise correlations within and between brain areas, and the dimensionality of noise within each area. These models are also constrained by network properties that arise from heterogeneous connectivity — for example, the strong relationship between the extent to which a neuron is modulated by spatial attention and the extent to which it exhibits divisive normalization when multiple stimuli are placed within its receptive field[100,106,189,205]. Recent observations indicating that only a low-dimensional subspace of information is shared between areas[59,83] will provide key constraints on future models of this type. The strength of this body of work is its connection to circuit mechanisms and its ability to account for observations about how different measures of variability are changed by processes such as attention and perceptual learning. The limitation is that these model networks include only limited representations of visual information and cannot perform tasks for which prioritization is necessary.

The second category of models uses deep artificial neural networks (DANNs) to model transformations across different stages of the ventral stream. This effort builds on foundational work demonstrating that, when DANNs are trained to perform object categorization, the functional organization of their different layers bears considerable resemblance to the representations of images and object identity in different brain areas as measured by single-unit recordings in the primate brain[206,207] (reviewed elsewere[208]). Extending these results, one study determined that one form of priority coding emerges from these same networks: image memorability variation[142]. Like the brain, the IT-analogous layers of DANNs trained to categorize objects respond more vigorously to some images than to others, and the magnitudes of the population responses in these layers are correlated with how memorable humans find those images. The strength of this type of modelling approach is that these models can be used to predict the responses of individual neurons or a population to complex, natural scenes (a model property often referred to as 'image computability'). Consequently, such models have proved effective in helping our understanding of how the brain deals with one type of noise — nuisance variation — in the context of tasks that require extracting information about one type of variable, such as object identity, while ignoring stimulus-induced changes in other types of irrelevant information such as changes in object position, size or background context. For example, these models do a reasonable (albeit imperfect) job of recapitulating the hierarchical transformation of object identity representations from a highly inaccessible and nonlinear format at earlier stages (such as in the V1) to a more accessible and linear format at later stages (such as in the IT), as well as predicting how individual neurons at different stages will respond to arbitrary natural images[206,207]. The limitation of these models is that they are 'noiseless' insofar as they lack trial variability — once they are trained and their weights fixed, they respond in the exact same way on different trials with the same stimulus, unlike the brain. As a result, they cannot recapitulate the changes in noise correlations (FIG. 2b) associated with attention and perceptual learning.

Although considerable modelling work remains to be done, our current understanding sheds some insight into the mechanisms that might support a unified yet distinguishable priority code. Namely, the observation that most of the variance in the activity of populations of visual cortical neurons occupies a number of dimensions far smaller than the number of neurons in the population constrains plausible mechanisms for priority coding. In a low-dimensional space where the range of firing rates is limited (from 0 to about 100–200 spikes per second per neuron), modulating the magnitude of a population response might be a particularly straightforward way to achieve prioritization. The mechanism is simple: small changes in the ratio of excitation to inhibition can cause big changes in mean rates. Additionally, the impact of these changes is big: more spikes leads to more (and probably more nuanced) activity in downstream populations, meaning that changes in response magnitude might have preferentially large effects on behaviour. In addition, as long as these changes are approximately multiplicative, this approach results in stimulus prioritization without affecting representations of visual identity (FIG. 3a). Whether this or other mechanisms support the neural correlates of priority coding is an important topic of future research.

Towards a unified account

After so many years of studying priority coding, what is needed to test speculative ideas about a unified account of priority coding? The field is well positioned to move in several directions that we believe are essential to achieve that goal. Here, we highlight four.

First, we need data that systematically compare different forms of priority. The first step towards a unified account of priority coding is to determine whether priority coding is in fact unified, in terms of whether different behavioural priority signatures are reflected via one or more shared neural correlates as well as underlying neural mechanisms. The traditional tasks, stimuli, brain areas and neural measurements typical of different subfields have made it difficult to directly compare different forms of priority coding. In recent years, the field has strongly encouraged collaborations between experimentalists and theorists through grant programmes such as the Collaborative Research in Computational Neuroscience (CRCNS) programme. However, this key obstacle to understanding priority coding will be best overcome by collaborations between teams of experimentalists who commit to making directly comparable measurements during different forms of priority coding and recording from multiple brain areas. These experiments will need controlled and robust behaviours as an anchor with which to evaluate modulation by different forms of priority as well as results of causal manipulations[137]. These experiments will help shed insight into the sources of input that lead to prioritization and the degree to which prioritization arises from feedforward as compared to feedback inputs. These experiments will also need to build on the behavioural repertoire in FIG. 2 to highlight other behavioural priority signatures, such as emotional salience[209–211].

Second, we need unified models that capture multiple types of noise. To date, two distinct classes of models have focused on different forms of noise. Circuit models focus on trial-to-trial variability at the expense of an account of how networks respond to naturalistic stimuli. By contrast, DANNs focus on stimulus variability, but their deterministic structure cannot capture priority-driven changes in trial-to-trial variability. Although some efforts have been made to reconcile these two different approaches, such as by inserting noise into early layers of DANNs and determining how connectivity affects the way that noise is changed throughout the pathway[212], these differences highlight the importance of creating models that can account for and make predictions about different aspects of neural population activity[213].

Third, we need experimental and theoretical efforts to understand the roles of subspaces of population activity in priority coding. Uncovering the relationship between the subspaces of neuronal population activity that encode different stimuli or features, forms of priority coding, and interactions between areas and behavioural prioritization is a critical direction for future work. New experiments are needed to measure the extent to which activity inside and outside those subspaces influences behaviour and how that activity is modulated by different forms of prioritization. Current instantiations of circuit models and DANNs do not typically account for current observations about subspaces, and doing so may place critical constraints on their architecture.

Fourth, we need data and models that capture how different forms of priority are integrated as well as the dynamic variation in priority behaviours. In natural behaviour, prioritization involves the integration of multiple forms of priority as a dynamic and continuous process to navigate a continually changing environment. By contrast, in laboratory studies, prioritization is typically studied by isolating different forms of priority and, for each type, switching between a small number of unambiguous task or stimulus conditions. Behavioural evidence supports notion that understanding the interactions between different types of prioritization is important for understanding priority coding; for example, perceptual learning can increase the saliency of familiar targets among distractors and reduce dependency on top-down attention[81,214]. Similarly, dynamic and continuous tasks are likely to change our understanding of the neural mechanisms by which cognition and perception interact[89,215–218] and help us more effectively identify the mapping between neurons and behaviour than binary tasks[168].

Prioritization is a critical part of natural vision, and the failure to appropriately prioritize visual information is a particularly debilitating symptom of neuropsychiatric disorders ranging from Alzheimer disease to autism, and is particularly relevant during recovery from blindness. Efforts to understand the neural basis of prioritization have resulted in a tantalizing smorgasbord of observations, and the time is ripe for merging different experimental and theoretical approaches to understand how they fit together. For future studies, doing so should be a priority.

Citation diversity statement

Recent work in neuroscience and related fields has identified citation biases whereby work from women and minorities are under-cited relative to other papers in the field[219–221]. Inclusion of citation diversity statements has been proposed as a way of increasing transparency surrounding citation practice[219,222]. Similar to REF.[219], the gender balance of citations was quantified according to the first names of the first/last authors using open source code[223] followed by manual adjustment to resolve ambiguities. Excluding self-citations, this article contains 64.9% man/man, 11.0% man/woman, 13.6% woman/man and 10.4% woman/woman citations. For comparison, proportions estimated from articles in five top neuroscience journals (as reported in REF.[219]) are 58.4% man/man, 9.4% man/woman, 25.5% woman/man and 6.7% woman/woman. Note that the used databases may not always be indicative of gender identity and do not account for intersex, non-binary or transgender individuals.

We also obtained predicted racial/ethnic category of the first and last author of each reference by databases that store the probability of a first and last name being carried by an author of colour[224,225] using the approaches described in REF.[226]. By this measure (and excluding self-citations), our references contain 11.4% authors of colour (first)/author of colour (last), 11.9% white author/author of colour, 21.3% author of colour/white author and 55.5% white author/white author. For comparison, proportions estimated from articles in five top neuroscience journals (as reported in REF.[226]) are 14.1% author of colour/author of colour, 15.6% white author/author of colour, 23.5% author of colour/white author and 46.7% white author/white author. Note that these databases are automated and may not always be indicative of racial or ethnic identity.

1. Srinivasan, M. V. Honey bees as a model for vision, perception, and cognition. *Annu. Rev. Entomol.* **55**, 267–284 (2010).
2. Stoeger, A. S. & Manger, P. Vocal learning in elephants: neural bases and adaptive context. *Curr. Opin. Neurobiol.* **28**, 101–107 (2014).
3. Ungerleider, L. G., Mishkin, M. in *Analysis of Visual Behavior* 549–586 (MIT Press, 1982).
4. Pasupathy, A., Popovkina, D. V. & Kim, T. Visual functions of primate area V4. *Annu. Rev. Vis. Sci.* **6**, 363–385 (2020).
5. DiCarlo, J. J., Zoccolan, D. & Rust, N. C. How does the brain solve visual object recognition? *Neuron* **73**, 415–34 (2012).
6. Gilbert, C. D., Sigman, M. & Crist, R. E. The neural basis of perceptual learning. *Neuron* **31**, 681–697 (2001).
7. Fahle, M. Perceptual learning: specificity versus generalization. *Curr. Opin. Neurobiol.* **15**, 154–160 (2005).
8. Seitz, A. R. & Dinse, H. R. A common framework for perceptual learning. *Curr. Opin. Neurobiol.* **17**, 148–153 (2007).
9. Sasaki, Y., Nanez, J. E. & Watanabe, T. Advances in visual perceptual learning and plasticity. *Nat. Rev. Neurosci.* **11**, 53–60 (2010).
10. Lu, Z.-L., Hua, T., Huang, C.-B., Zhou, Y. & Dosher, B. A. Visual perceptual learning. *Neurobiol. Learn. Mem.* **95**, 145–151 (2011).
11. Sagi, D. Perceptual learning in vision research. *Vis. Res.* **51**, 1552–1566 (2011).
12. Watanabe, T. & Sasaki, Y. Perceptual learning: toward a comprehensive theory. *Annu. Rev. Psychol.* **66**, 197–221 (2015).
13. Tamaki, M. et al. Reward does not facilitate visual perceptual learning until sleep occurs. *Proc. Natl Acad. Sci. USA* **117**, 959–968 (2020).
14. Tamaki, M., Wang, Z., Watanabe, T. & Sasaki, Y. Trained-feature-specific offline learning by sleep in an orientation detection task. *J. Vis.* **19**, 12 (2019).
15. Szpiro, S. F. A. & Carrasco, M. Exogenous attention enables perceptual learning. *Psychol. Sci.* **26**, 1854–1862 (2015).
16. Kim, R., Seitz, A., Feenstra, H. & Shams, L. Testing assumptions of statistical learning: is it long-term and implicit? *Neurosci. Lett.* **461**, 145–149 (2009).
17. Fiorentini, A. & Berardi, N. Perceptual learning specific for orientation and spatial frequency. *Nature* **287**, 43–44 (1980).
18. Pourtois, G., Rauss, K. S., Vuilleumier, P. & Schwartz, S. Effects of perceptual learning on primary visual cortex activity in humans. *Vis. Res.* **48**, 55–62 (2008).
19. Mollon, J. D. & Danilova, M. V. Three remarks on perceptual learning. *Spat. Vis.* **10**, 51–58 (1996).
20. Little, D. F., Zhang, Y.-X. & Wright, B. A. Disruption of perceptual learning by a brief practice break. *Curr. Biol.* **27**, 3699–3705.e3 (2017).
21. Kang, D.-W. et al. Structural and functional connectivity changes beyond visual cortex in a later phase of visual perceptual learning. *Sci. Rep.* **8**, 5186 (2018).
22. Kastner, S. & Ungerleider, L. G. Mechanisms of visual attention in the human cortex. *Annu. Rev. Neurosci.* **23**, 315–341 (2000).
23. Anton-Erxleben, K. & Carrasco, M. Attentional enhancement of spatial resolution: linking behavioural and neurophysiological evidence. *Nat. Rev. Neurosci.* **14**, 188–200 (2013).
24. Krauzlis, R. J., Lovejoy, L. P. & Zénon, A. Superior colliculus and visual spatial attention. *Annu. Rev. Neurosci.* **36**, 165–182 (2013).
25. Amso, D. & Scerif, G. The attentive brain: insights from developmental cognitive neuroscience. *Nat. Rev. Neurosci.* **16**, 606–619 (2015).
26. Clark, K., Squire, R. F., Merrikhi, Y. & Noudoost, B. Visual attention: linking prefrontal sources to neuronal and behavioral correlates. *Prog. Neurobiol.* **132**, 59–80 (2015).
27. Treisman, A. M. & Gelade, G. A feature-integration theory of attention. *Cogn. Psychol.* **12**, 97–136 (1980).
28. Treisman, A. & Sato, S. Conjunction search revisited. *J. Exp. Psychol. Hum. Percept. Perform.* **16**, 459–478 (1990).
29. Wolfe, J. M. Visual search in continuous, naturalistic stimuli. *Vis. Res.* **34**, 1187–1195 (1994).
30. Posner, M. I. Orienting of attention. *Q. J. Exp. Psychol.* **32**, 3–25 (1980).
31. Carrasco, M. Visual attention: the past 25 years. *Vis. Res.* **51**, 1484–1525 (2011).
32. Bacon, W. F. & Egeth, H. E. Overriding stimulus-driven attentional capture. *Percept. Psychophys.* **55**, 485–496 (1994).
33. Wolfe, J. M., Cave, K. R. & Franzel, S. L. Guided search: an alternative to the feature integration model for visual search. *J. Exp. Psychol. Hum. Percept. Perform.* **15**, 419–433 (1989).
34. Nakayama, K. & Martini, P. Situating visual search. *Vis. Res.* **51**, 1526–1537 (2011).

35. Eckstein, M. P. Visual search: a retrospective. *J. Vis.* **11**, 14 (2011).
36. Judd, T., Ehinger, K., Durand, F. & Torralba, A. Learning to predict where humans look. *IEEE 12th International Conference on Computer Vision* 2106–2113 (IEEE, 2009).
37. Isola, P., Xiao, J., Parikh, D., Torralba, A. & Oliva, A. What makes a photograph memorable? *IEEE Trans. Pattern Anal. Mach. Intell.* **36**, 1469–1482 (2014). **This paper describes image memorability-associated behavioural variation across a large and arbitrarily selected set of photographs.**
38. Bainbridge, W. A., Isola, P. & Oliva, A. The intrinsic memorability of face photographs. *J. Exp. Psychol. Gen.* **142**, 1323–1334 (2013).
39. Goetschalckx, L. & Wagemans, J. MemCat: a new category-based image set quantified on memorability. *PeerJ* **7**, e8169 (2019).
40. Khosla, A., Raju, A. S., Torralba, A. & Oliva, A. Understanding and predicting image memorability at a large scale. *International Conference on Computer Vision* 2390–2398 (IEEE, 2015).
41. Rust, N. C. & Mehrpour, V. Understanding image memorability. *Trends Cogn. Sci.* **24**, 557–568 (2020).
42. Bainbridge, W. A. The memorability of people: intrinsic memorability across transformations of a person's face. *J. Exp. Psychol. Learn. Mem. Cogn.* **43**, 706–716 (2017).
43. Bainbridge, W. A. in *Psychology of Learning and Motivation* (eds Federmeier, K. D. & Beck, D. M.) Ch. 1 vol. 70 1–27 (Academic, 2019).
44. Bylinskii, Z., Isola, P., Bainbridge, C., Torralba, A. & Oliva, A. Intrinsic and extrinsic effects on image memorability. *Vis. Res.* **116**, 165–178 (2015).
45. Jutras, M. J. & Buffalo, E. A. Recognition memory signals in the macaque hippocampus. *Proc. Natl Acad. Sci. USA* **107**, 401–406 (2010).
46. Hall, B. A., Melfi, V., Burns, A., McGill, D. M. & Doyle, R. E. Curious creatures: a multi-taxa investigation of responses to novelty in a zoo environment. *PeerJ* **6**, e4454 (2018).
47. Teller, D. Y. The forced-choice preferential looking procedure: a psychophysical technique for use with human infants. *Infant. Behav. Dev.* **2**, 135–153 (1979).
48. Meyer, T. & Rust, N. C. Single-exposure visual memory judgments are reflected in inferotemporal cortex. *eLife* **7**, e32259 (2018).
49. Rich, J. B. in *Encyclopedia of Clinical Neuropsychology* (eds Kreutzer, J. S., DeLuca, J. & Caplan, B.) 2124–2128 (Springer, 2011).
50. Desimone, R. & Duncan, J. Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* **18**, 193–222 (1995).
51. Desimone, R. Neural mechanisms for visual memory and their role in attention. *Proc. Natl Acad. Sci. USA* **93**, 13494–13499 (1996).
52. Reynolds, J. H. & Chelazzi, L. Attentional modulation of visual processing. *Annu. Rev. Neurosci.* **27**, 611–647 (2004).
53. Maunsell, J. H. R. Neuronal mechanisms of visual attention. *Annu. Rev. Vis. Sci.* **1**, 373–391 (2015).
54. Cohen, M. R. & Kohn, A. Measuring and interpreting neuronal correlations. *Nat. Neurosci.* **14**, 811–819 (2011).
55. Roth, N. & Rust, N. C. Rethinking assumptions about how trial and nuisance variability impact neural task performance in a fast-processing regime. *J. Neurophysiol.* **121**, 115–130 (2019).
56. Law, C.-T. & Gold, J. I. Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat. Neurosci.* **11**, 505–513 (2008).
57. Fries, P. Rhythms for cognition: communication through coherence. *Neuron* **88**, 220–235 (2015).
58. Womelsdorf, T. & Fries, P. The role of neuronal synchronization in selective attention. *Curr. Opin. Neurobiol.* **17**, 154–160 (2007).
59. Kohn, A. et al. Principles of corticocortical communication: proposed schemes and design considerations. *Trends Neurosci.* **43**, 725–737 (2020).
60. Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation through neural population dynamics. *Annu. Rev. Neurosci.* **43**, 249–275 (2020).
61. Srinath, R., Ruff, D. A. & Cohen, M. R. Attention improves information flow between neuronal populations without changing the communication subspace. *Curr. Biol.* **31**, 5299–5313 (2021). **This study demonstrated that attention improves the ability to predict activity in decision neurons from the activity of neurons in the visual cortex without changing the dimensionality of the communication subspace.**
62. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* **17**, 440–448 (2014). **This study demonstrated that only a subspace of activity in a given population affects behaviour, while the remaining subset (null space) does not, laying a foundation for the subspace framework.**
63. Saxena, S. & Cunningham, J. P. Towards the neural population doctrine. *Curr. Opin. Neurobiol.* **55**, 103–111 (2019).
64. Gawne, T. J. & Richmond, B. J. How independent are the messages carried by adjacent inferior temporal cortical neurons? *J. Neurosci.* **13**, 2758–2771 (1993).
65. Chen, Y.-P., Lin, C.-P., Hsu, Y.-C. & Hung, C. P. Network anisotropy trumps noise for efficient object coding in macaque inferior temporal cortex. *J. Neurosci.* **35**, 9889–9899 (2015).
66. Hung, C. P., Cui, D., Chen, Y.-P., Lin, C.-P. & Levine, M. R. Correlated activity supports efficient cortical processing. *Front. Comput. Neurosci.* **8**, 171 (2014).
67. Ziemba, C. M., Freeman, J., Movshon, J. A. & Simoncelli, E. P. Selectivity and tolerance for visual texture in macaque V2. *Proc. Natl Acad. Sci. USA* **113**, E3140–3149 (2016).
68. Dapello, J. et al. Simulating a primary visual cortex at the front of CNNs improves robustness to image perturbations. *Advances in Neural Information Processing Systems 33* https://proceedings.neurips.cc/paper/2020/hash/98b17f068d5d9b7668e19fb8ae470841-Abstract.html (2020).
69. Parker, A. J. & Newsome, W. T. Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* **21**, 227–277 (1998).
70. Schoups, A., Vogels, R., Qian, N. & Orban, G. Practising orientation identification improves orientation coding in V1 neurons. *Nature* **412**, 549–553 (2001).
71. Ghose, G. M., Yang, T. & Maunsell, J. H. R. Physiological correlates of perceptual learning in monkey V1 and V2. *J. Neurophysiol.* **87**, 1867–1888 (2002).
72. Raiguel, S., Vogels, R., Mysore, S. G. & Orban, G. A. Learning to see the difference specifically alters the most informative V4 neurons. *J. Neurosci.* **26**, 6589–6602 (2006).
73. Uka, T., Sasaki, R. & Kumano, H. Change in choice-related response modulation in area MT during learning of a depth-discrimination task is consistent with task learning. *J. Neurosci.* **32**, 13689–13700 (2012).
74. Ni, A. M., Ruff, D. A., Alberts, J. J., Symmonds, J. & Cohen, M. R. Learning and attention reveal a general relationship between population activity and behavior. *Science* **359**, 463–465 (2018). **This paper demonstrated that perceptual learning and attention were reflected in V4 primarily by an overlapping neural correlate: changes in correlated trial variability, which is strongly linked to performance.**
75. Yotsumoto, Y. et al. Location-specific cortical activation changes during sleep after training for perceptual learning. *Curr. Biol.* **19**, 1278–1282 (2009).
76. Poort, J. et al. Learning enhances sensory and multiple non-sensory representations in primary visual cortex. *Neuron* **86**, 1478–1490 (2015).
77. Karmarkar, U. R. & Dan, Y. Experience-dependent plasticity in adult visual cortex. *Neuron* **52**, 577–585 (2006).
78. Yan, Y., Zhaoping, L. & Li, W. Bottom-up saliency and top-down learning in the primary visual cortex of monkeys. *Proc. Natl Acad. Sci. USA* **115**, 10499–10504 (2018).
79. Jia, K. et al. Recurrent processing drives perceptual plasticity. *Curr. Biol.* **30**, 4177–4187.e4 (2020).
80. Yang, T. & Maunsell, J. H. R. The effect of perceptual learning on neuronal responses in monkey visual area V4. *J. Neurosci.* **24**, 1617–1626 (2004).
81. Gu, Y. et al. Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* **71**, 750–761 (2011).
82. Huang, C. et al. Circuit models of low-dimensional shared variability in cortical networks. *Neuron* **101**, 337–348.e4 (2019). **This study presents a circuit model in which realistic variability that is low-dimensional and modulated by attention arises from a biophysically realistic network.**
83. Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M. & Kohn, A. Cortical areas interact through a communication subspace. *Neuron* **102**, 249–259.e4 (2019). **This study demonstrated that populations of neurons in visual areas V1 and V2, which are highly interconnected, share activity via a subspace of activity that is lower-dimensional than the activity in either area.**
84. Treue, S. & Martínez Trujillo, J. C. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* **399**, 575–579 (1999).
85. McAlonan, K., Cavanaugh, J. & Wurtz, R. H. Guarding the gateway to cortex with attention in visual thalamus. *Nature* **456**, 391–394 (2008).
86. Moran, J. & Desimone, R. Selective attention gates visual processing in the extrastriate cortex. *Science* **229**, 782–784 (1985).
87. Buffalo, E. A., Fries, P., Landman, R., Liang, H. & Desimone, R. A backward progression of attentional effects in the ventral stream. *Proc. Natl Acad. Sci. USA* **107**, 361–365 (2010).
88. Saenz, M., Buracas, G. T. & Boynton, G. M. Global effects of feature-based attention in human visual cortex. *Nat. Neurosci.* **5**, 631–632 (2002).
89. Cohen, M. R. & Maunsell, J. H. R. When attention wanders: how uncontrolled fluctuations in attention affect performance. *J. Neurosci.* **31**, 15802–15806 (2011).
90. Cohen, M. R. & Maunsell, J. H. R. Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* **12**, 1594–1600 (2009).
91. Gregoriou, G. G., Rossi, A. F., Ungerleider, L. G. & Desimone, R. Lesions of prefrontal cortex reduce attentional modulation of neuronal responses and synchrony in V4. *Nat. Neurosci.* **17**, 1003–1011 (2014).
92. Herrero, J. L., Gieselmann, M. A., Sanayei, M. & Thiele, A. Attention-induced variance and noise correlation reduction in macaque V1 is mediated by NMDA receptors. *Neuron* **78**, 729–739 (2013).
93. Luo, T. Z. & Maunsell, J. H. R. Neuronal modulations in visual cortex are associated with only one of multiple components of attention. *Neuron* **86**, 1182–1188 (2015).
94. Mayo, J. P. & Maunsell, J. H. R. Graded neuronal modulations related to visual spatial attention. *J. Neurosci.* **36**, 5353–5361 (2016).
95. Mitchell, J. F., Sundberg, K. A. & Reynolds, J. H. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* **63**, 879–888 (2009).
96. Nandy, A. S., Nassi, J. J. & Reynolds, J. H. Laminar organization of attentional modulation in macaque visual area V4. *Neuron* **93**, 235–246 (2017).
97. Ruff, D. A. & Cohen, M. R. Attention can either increase or decrease spike count correlations in visual cortex. *Nat. Neurosci.* **17**, 1591–1597 (2014).
98. Ruff, D. A. & Cohen, M. R. Stimulus dependence of correlated variability across cortical areas. *J. Neurosci.* **36**, 7546–7556 (2016).
99. Ruff, D. A. & Cohen, M. R. Simultaneous multi-area recordings suggest that attention improves performance by reshaping stimulus representations. *Nat. Neurosci.* **22**, 1669–1676 (2019).
100. Verhoef, B.-E. & Maunsell, J. H. R. Attention-related changes in correlated neuronal activity arise from normalization mechanisms. *Nat. Neurosci.* **20**, 969–977 (2017).
101. Yan, Y. et al. Perceptual training continuously refines neuronal population codes in primary visual cortex. *Nat. Neurosci.* **17**, 1380–1387 (2014).
102. Zénon, A. & Krauzlis, R. J. Attention deficits without cortical neuronal deficits. *Nature* **489**, 434–437 (2012).
103. Ruff, D. A. & Cohen, M. R. Global cognitive factors modulate correlated response variability between V4 neurons. *J. Neurosci.* **34**, 16408–16416 (2014).
104. Oemisch, M., Westendorff, S., Everling, S. & Womelsdorf, T. Interareal spike-train correlations of anterior cingulate and dorsal prefrontal cortex during attention shifts. *J. Neurosci.* **35**, 13076–13089 (2015).
105. Pooresmaeili, A. & Roelfsema, P. R. A growth-cone model for the spread of object-based attention during contour grouping. *Curr. Biol.* **24**, 2869–2877 (2014).
106. Ruff, D. A., Alberts, J. J. & Cohen, M. R. Relating normalization to neuronal populations across cortical areas. *J. Neurophysiol.* **116**, 1375–1386 (2016).
107. Bichot, N. P., Rossi, A. F. & Desimone, R. Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* **308**, 529–534 (2005).

108. Bosman, C. A. et al. Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron* **75**, 875–888 (2012).
109. Buschman, T. J. & Miller, E. K. Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* **315**, 1860–1862 (2007).
110. Fries, P., Reynolds, J. H., Rorie, A. E. & Desimone, R. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* **291**, 1560–1563 (2001).
111. Gregoriou, G. G., Gotts, S. J., Zhou, H. & Desimone, R. Long-range neural coupling through synchronization with attention. *Prog. Brain Res.* **176**, 35–45 (2009).
112. Lakatos, P., Karmos, G., Mehta, A. D., Ulbert, I. & Schroeder, C. E. Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* **320**, 110–113 (2008).
113. Miller, E. K. & Buschman, T. J. Cortical circuits for the control of attention. *Curr. Opin. Neurobiol.* **23**, 216–222 (2013).
114. Saalmann, Y. B., Pigarev, I. N. & Vidyasagar, T. R. Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. *Science* **316**, 1612–1615 (2007).
115. Saproo, S. & Serences, J. T. Attention improves transfer of motion information between V1 and MT. *J. Neurosci.* **34**, 3586–3596 (2014).
116. Womelsdorf, T. & Fries, P. Neuronal coherence during selective attentional processing and sensory-motor integration. *J. Physiol. Paris* **100**, 182–193 (2006).
117. Briggs, F., Mangun, G. R. & Usrey, W. M. Attention enhances synaptic efficacy and the signal-to-noise ratio in neural circuits. *Nature* **499**, 476–480 (2013).
118. Dagnino, B., Gariel-Mathis, M.-A. & Roelfsema, P. R. Microstimulation of area V4 has little effect on spatial attention and on perception of phosphenes evoked in area V1. *J. Neurophysiol.* **113**, 730–739 (2015).
119. Klink, P. C., Jeurissen, D., Theeuwes, J., Denys, D. & Roelfsema, P. R. Working memory accuracy for multiple targets is driven by reward expectation and stimulus contrast with different time-courses. *Sci. Rep.* **7**, 9082 (2017).
120. Moore, T. & Armstrong, K. M. Selective gating of visual signals by microstimulation of frontal cortex. *Nature* **421**, 370–373 (2003).
121. Ruff, D. A. & Cohen, M. R. A normalization model suggests that attention changes the weighting of inputs between visual areas. *Proc. Natl Acad. Sci. USA* **114**, E4085–E4094 (2017).
122. Salinas, E. & Sejnowski, T. J. Correlated neuronal activity and the flow of neural information. *Nat. Rev. Neurosci.* **2**, 539–550 (2001).
123. Azouz, R. & Gray, C. M. Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. *Neuron* **37**, 513–523 (2003).
124. Buffalo, E. A., Fries, P., Landman, R., Buschman, T. J. & Desimone, R. Laminar differences in gamma and alpha coherence in the ventral stream. *Proc. Natl Acad. Sci. USA* **108**, 11262–11267 (2011).
125. Womelsdorf, T., Fries, P., Mitra, P. P. & Desimone, R. Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* **439**, 733–736 (2006).
126. Bosman, C. A., Womelsdorf, T., Desimone, R. & Fries, P. A microsaccadic rhythm modulates gamma-band synchronization and behavior. *J. Neurosci.* **29**, 9471–9480 (2009).
127. Knierim, J. J. & van Essen, D. C. Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J. Neurophysiol.* **67**, 961–980 (1992).
128. Burrows, B. E. & Moore, T. Influence and limitations of popout in the selection of salient visual stimuli by area V4 neurons. *J. Neurosci.* **29**, 15169–15177 (2009).
129. Hegdé, J. & Felleman, D. J. How selective are V1 cells for pop-out stimuli? *J. Neurosci.* **23**, 9968–9980 (2003).
130. Chelazzi, L., Duncan, J., Miller, E. K. & Desimone, R. Responses of neurons in inferior temporal cortex during memory-guided visual search. *J. Neurophysiol.* **80**, 2918–2940 (1998).
131. Thompson, K. G. & Bichot, N. P. A visual salience map in the primate frontal eye field. *Prog. Brain Res.* **147**, 251–262 (2005).
132. Ogawa, T. & Komatsu, H. Neuronal dynamics of bottom-up and top-down processes in area V4 of macaque monkeys performing a visual search. *Exp. Brain Res.* **173**, 1–13 (2006).
133. Bichot, N. P., Heard, M. T., DeGennaro, E. M. & Desimone, R. A source for feature-based attention in the prefrontal cortex. *Neuron* **88**, 832–844 (2015).
134. Pagan, M., Urban, L. S., Wohl, M. P. & Rust, N. C. Signals in inferotemporal cortex and perirhinal cortex suggest an untangling of visual target information. *Nat. Neurosci.* **16**, 1132–1139 (2013).
135. Roth, N. & Rust, N. C. Inferotemporal cortex multiplexes behaviorally-relevant target match signals and visual representations in a manner that minimizes their interference. *PLoS One* **13**, e0200528 (2018).
136. Roth, N. & Rust, N. C. The integration of visual and target signals in V4 and IT during visual object search. *J. Neurophysiol.* **122**, 2522–2540 (2019).
137. Chen, X. et al. Parietal cortex regulates visual salience and salience-driven behavior. *Neuron* **106**, 177–187. e4 (2020).
138. Koch, C. & Ullman, S. Shifts in selective visual attention: towards the underlying neural circuitry. *Hum. Neurobiol.* **4**, 219–227 (1985).
139. Borji, A., Sihite, D. N. & Itti, L. What stands out in a scene? A study of human explicit saliency judgment. *Vis. Res.* **91**, 62–77 (2013).
140. Itti, L., Koch, C. & Niebur, E. A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**, 1254–1259 (1998).
141. Wang, J., Borji, A., Jay Kuo, C.-C. & Itti, L. Learning a combined model of visual saliency for fixation prediction. *IEEE Trans. Image Process.* **25**, 1566–1579 (2016).
142. Jaegle, A. et al. Population response magnitude variation in inferotemporal cortex predicts image memorability. *eLife* **8**, e47596 (2019).
This paper describes the neural correlate of one behavioural priority signature, image memorability, and demonstrates that it emerges naturally from deep neural networks trained to identify objects.
143. Fahy, F. L., Riches, I. P. & Brown, M. W. Neuronal activity related to visual recognition memory: long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex. *Exp. Brain Res.* **96**, 457–472 (1993).
144. Li, L., Miller, E. K. & Desimone, R. The representation of stimulus familiarity in anterior inferior temporal cortex. *J. Neurophysiol.* **69**, 1918–1929 (1993).
145. Xiang, J. Z. & Brown, M. W. Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology* **37**, 657–676 (1998).
146. Womelsdorf, T. Translating expectation into visual selection through a beta-synchronous fronto-parietal neural subnetwork. *Neuron* **109**, 8–10 (2021).
147. Ossandón, T. et al. Efficient 'pop-out' visual search elicits sustained broadband γ activity in the dorsal attention network. *J. Neurosci.* **32**, 3414–3421 (2012).
148. Jutras, M. J., Fries, P. & Buffalo, E. A. Oscillatory activity in the monkey hippocampus during visual exploration and memory formation. *Proc. Natl Acad. Sci. USA* **110**, 13144–13149 (2013).
149. Jutras, M. J., Fries, P. & Buffalo, E. A. Gamma-band synchronization in the macaque hippocampus and memory formation. *J. Neurosci.* **29**, 12521–12531 (2009).
150. Sederberg, P. B. et al. Gamma oscillations distinguish true from false memories. *Psychol. Sci.* **18**, 927–932 (2007).
151. Sederberg, P. B. et al. Hippocampal and neocortical gamma oscillations predict memory formation in humans. *Cereb. Cortex* **17**, 1190–1196 (2007).
152. Fell, J. et al. Human memory formation is accompanied by rhinal-hippocampal coupling and decoupling. *Nat. Neurosci.* **4**, 1259–1264 (2001).
153. Naya, Y. & Suzuki, W. A. Integrating what and when across the primate medial temporal lobe. *Science* **333**, 773–776 (2011).
154. Kiani, R., Esteky, H., Mirpour, K. & Tanaka, K. Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *J. Neurophysiol.* **97**, 4296–4309 (2007).
155. Yu, B. M. et al. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* **102**, 614–635 (2009).
156. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17**, 1500–1509 (2014).
157. Sadtler, P. T. et al. Neural constraints on learning. *Nature* **512**, 423–426 (2014).
158. Kiani, R. et al. Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. *Neuron* **85**, 1359–1373 (2015).
159. Cowley, B. R., Smith, M. A., Kohn, A. & Yu, B. M. Stimulus-driven population activity patterns in macaque primary visual cortex. *PLoS Comput. Biol.* **12**, e1005185 (2016).
160. Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nat. Commun.* **7**, 13239 (2016).
161. Golub, M. D., Chase, S. M., Batista, A. P. & Yu, B. M. Brain-computer interfaces for dissecting cognitive processes underlying sensorimotor control. *Curr. Opin. Neurobiol.* **37**, 53–58 (2016).
162. Morcos, A. S. & Harvey, C. D. History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat. Neurosci.* **19**, 1672–1681 (2016).
163. Elsayed, G. F. & Cunningham, J. P. Structure in neural population recordings: an expected byproduct of simpler phenomena? *Nat. Neurosci.* **20**, 1310–1318 (2017).
164. Jazayeri, M. & Afraz, A. Navigating the neural space in search of the neural code. *Neuron* **93**, 1003–1014 (2017).
165. Miri, A. et al. Behaviorally selective engagement of short-latency effector pathways by motor cortex. *Neuron* **95**, 683–696.e11 (2017).
166. Pitkow, X. & Angelaki, D. E. Inference in the brain: statistics flowing in redundant population codes. *Neuron* **94**, 943–953 (2017).
167. Pandarinath, C. et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat. Methods* **15**, 805–815 (2018).
168. Ruff, D. A., Ni, A. M. & Cohen, M. R. Cognition as a window into neuronal population space. *Annu. Rev. Neurosci.* **41**, 77–97 (2018).
169. Stringer, C., Pachitariu, M., Steinmetz, N., Carandini, M. & Harris, K. D. High-dimensional geometry of population responses in visual cortex. *Nature* **571**, 361–365 (2019).
170. Rajan, K. & Bialek, W. Maximally informative "stimulus energies" in the analysis of neural responses to natural signals. *PLoS One* **8**, e71959 (2013).
171. Ecker, A. S. et al. State dependence of noise correlations in macaque primary visual cortex. *Neuron* **82**, 235–48 (2014).
172. Ruff, D. A., Xue, C., Kramer, L. E., Baqai, F. & Cohen, M. R. Low rank mechanisms underlying flexible visual representations. *Proc. Natl Acad. Sci. USA* **117**, 29321–29329 (2020).
173. Mehrpour, V., Meyer, T., Simoncelli, E. P. & Rust, N. C. Pinpointing the neural signatures of single-exposure recognition memory. *Proc. Natl Acad. Sci. USA* **118**, e2021660118 (2021).
In the context of a visual memory task, this study demonstrated that the population subspaces encoding novelty and a task-irrelevant variable, contrast, were partially overlapping and could be disambiguated.
174. McAdams, C. J. & Maunsell, J. H. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* **19**, 431–441 (1999).
175. Reynolds, J. H. & Heeger, D. J. The normalization model of attention. *Neuron* **61**, 168–185 (2009).
176. Herrmann, K., Heeger, D. J. & Carrasco, M. Feature-based attention enhances performance by increasing response gain. *Vis. Res.* **74**, 10–20 (2012).
177. Eldar, E., Cohen, J. D. & Niv, Y. The effects of neural gain on attention and learning. *Nat. Neurosci.* **16**, 1146–1153 (2013).
178. McMahon, D. B. & Olson, C. R. Repetition suppression in monkey inferotemporal cortex: relation to behavioral priming. *J. Neurophysiol.* **97**, 3532–43 (2007).
179. Grill-Spector, K., Henson, R. & Martin, A. Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn. Sci.* **10**, 14–23 (2006).
180. Woloszyn, L. & Sheinberg, D. L. Effects of long-term visual experience on responses of distinct classes of single units in inferior temporal cortex. *Neuron* **74**, 193–205 (2012).
181. Arandia-Romero, I., Tanabe, S., Drugowitsch, J., Kohn, A. & Moreno-Bote, R. Multiplicative and additive modulation of neuronal tuning with population activity affects encoded information. *Neuron* **89**, 1305–1316 (2016).
182. Heeger, D. J. Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* **9**, 181–197 (1992).
183. Rigotti, M. et al. The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).

184. Fusi, S., Miller, E. K. & Rigotti, M. Why neurons mix: high dimensionality for higher cognition. *Curr. Opin. Neurobiol.* **37**, 66–74 (2016).
185. Panichello, M. F. & Buschman, T. J. Shared mechanisms underlie the control of working memory and attention. *Nature* **592**, 601–605 (2021).
This study directly compared the neural correlates of working memory and attention, which is a critical step toward determining whether different forms of prioritization are mediated by similar mechanisms.
186. Moreno-Bote, R. et al. Information-limiting correlations. *Nat. Neurosci.* **17**, 1410–1417 (2014).
187. Kohn, A., Coen-Cagli, R., Kanitscheider, I. & Pouget, A. Correlations and neuronal population information. *Annu. Rev. Neurosci.* **39**, 237–256 (2016).
188. Deneve, S., Latham, P. E. & Pouget, A. Efficient computation and cue integration with noisy population codes. *Nat. Neurosci.* **4**, 826–831 (2001).
189. Lee, J. & Maunsell, J. H. R. A normalization model of attentional modulation of single unit responses. *PLoS One* **4**, e4651 (2009).
190. Boynton, G. M. A framework for describing the effects of attention on visual responses. *Vis. Res.* **49**, 1129–1143 (2009).
191. Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 (2011).
192. Zhaoping, L. A new framework for understanding vision from the perspective of the primary visual cortex. *Curr. Opin. Neurobiol.* **58**, 1–10 (2019).
193. Bogacz, R. & Brown, M. W. Comparison of computational models of familiarity discrimination in the perirhinal cortex. *Hippocampus* **13**, 494–524 (2003).
194. Itti, L. & Koch, C. Computational modelling of visual attention. *Nat. Rev. Neurosci.* **2**, 194–203 (2001).
195. Shadlen, M. N. & Newsome, W. T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
196. van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).
197. Amit, D. J. & Brunel, N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* **7**, 237–252 (1997).
198. Doiron, B., Litwin-Kumar, A., Rosenbaum, R., Ocker, G. K. & Josić, K. The mechanics of state-dependent neural correlations. *Nat. Neurosci.* **19**, 383–393 (2016).
199. Deco, G. et al. Resting-state functional connectivity emerges from structurally and dynamically shaped slow linear fluctuations. *J. Neurosci.* **33**, 11239–11252 (2013).
200. Wimmer, K. et al. Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. *Nat. Commun.* **6**, 6177 (2015).
201. Kanashiro, T., Ocker, G. K., Cohen, M. R. & Doiron, B. Attentional modulation of neuronal variability in circuit models of cortex. *eLife* **6**, e23978 (2017).
202. Hennequin, G., Ahmadian, Y., Rubin, D. B., Lengyel, M. & Miller, K. D. The dynamical regime of sensory cortex: stable dynamics around a single stimulus-tuned attractor account for patterns of noise variability. *Neuron* **98**, 846–860.e5 (2018).
203. Sadeh, S. & Clopath, C. Theory of neuronal perturbome in cortical networks. *Proc. Natl Acad. Sci. USA* **117**, 26966–26976 (2020).
204. Sweeney, Y. & Clopath, C. Population coupling predicts the plasticity of stimulus responses in cortical circuits. *eLife* **9**, e56053 (2020).
205. Ni, A. M., Ray, S. & Maunsell, J. H. R. Tuned normalization explains the size of attention modulations. *Neuron* **73**, 803–813 (2012).
206. Khaligh-Razavi, S.-M. & Kriegeskorte, N. Deep supervised, but not unsupervised, models may explain IT cortical representation. *PLoS Comput. Biol.* **10**, e1003915 (2014).
207. 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).
208. Yamins, D. L. K. & DiCarlo, J. J. Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* **19**, 356–365 (2016).
209. Tamietto, M. & de Gelder, B. Neural bases of the non-conscious perception of emotional signals. *Nat. Rev. Neurosci.* **11**, 697–709 (2010).
210. Pessoa, L., Kastner, S. & Ungerleider, L. G. Neuroimaging studies of attention: from modulation of sensory processing to top-down control. *J. Neurosci.* **23**, 3990–3998 (2003).
211. Markovic, J., Anderson, A. K. & Todd, R. M. Tuning to the significant: neural and genetic processes underlying affective enhancement of visual perception and memory. *Behav. Brain Res.* **259**, 229–241 (2014).
212. Wenliang, L. K. & Seitz, A. R. Deep neural networks for modeling visual perceptual learning. *J. Neurosci.* **38**, 6028–6044 (2018).
One of the only studies to investigate the consequences of adding trial variability that is modulated by priority coding to an otherwise deterministic deep network.
213. Zenke, F. et al. Visualizing a joint future of neuroscience and neuromorphic engineering. *Neuron* **109**, 571–575 (2021).
214. Sigman, M. & Gilbert, C. D. Learning to find a shape. *Nat. Neurosci.* **3**, 264–269 (2000).
215. Sarafyazd, M. & Jazayeri, M. Hierarchical reasoning by neural circuits in the frontal cortex. *Science* **364**, eaav8911 (2019).
216. Cohen, M. R. & Maunsell, J. H. R. Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron* **70**, 1192–1204 (2011).
217. Purcell, B. A. & Kiani, R. Neural mechanisms of post-error adjustments of decision policy in parietal cortex. *Neuron* **89**, 658–671 (2016).
218. Ebitz, R. B., Tu, J. C. & Hayden, B. Y. Rules warp feature encoding in decision-making circuits. *PLoS Biol.* **18**, e3000951 (2020).
219. Dworkin, J. D. et al. The extent and drivers of gender imbalance in neuroscience reference lists. *Nat. Neurosci.* **23**, 918–926 (2020).
220. Maliniak, D., Powers, R. & Walter, B. F. The gender citation gap in international relations. *Int. Organ.* **67**, 889–922 (2013).
221. Caplar, N., Tacchella, S. & Birrer, S. Quantitative evaluation of gender bias in astronomical publications from citation counts. *Nat. Astron.* **1**, 0141 (2017).
222. Zurn, P., Bassett, D. S. & Rust, N. C. The citation diversity statement: a practice of transparency, a way of life. *Trends Cogn. Sci.* **24**, 669–672 (2020).
223. Bertolero, M. *Balanced Citer* https://github.com/mb3152/balanced_citer (2021).
224. Ambekar, A., Ward, C., Mohammed, J., Male, S. & Skiena, S. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 49–58 (ACM, 2009).
225. Sood, G. & Laohaprapanon, S. Predicting race and ethnicity from the sequence of characters in a name. *arXiv* https://arxiv.org/abs/1805.02109 (2018).
226. Bertolero, M. A. et al. Racial and ethnic imbalance in neuroscience reference lists and intersections with gender. *bioRxiv* https://doi.org/10.1101/2020.10.12.336230 (2020).

Acknowledgements

This work was supported by the Simons Foundation (Simons Collaboration on the Global Brain award 543033 to N.C.R. and 542961SPI to M.R.C.), the US National Eye Institute of the National Institutes of Health (awards R01EY020851 and R01EY032878 to N.C.R. and awards R01EY022930 and R01NS121913 to M.R.C.) and the US National Science Foundation (award 2043255 to N.C.R.).

Author contributions

The authors contributed equally to all aspects of the article.

Competing interests

The authors declare no competing interests.

Peer review information

Nature Reviews Neuroscience thanks the anonymous reviewers for their contribution to the peer review of this work.

Publisher's note