



The structures and functions of correlations in neural population codes

Stefano Panzeri^{1,2}, Monica Moroni², Houman Safaai³ and Christopher D. Harvey³

Abstract | The collective activity of a population of neurons, beyond the properties of individual cells, is crucial for many brain functions. A fundamental question is how activity correlations between neurons affect how neural populations process information. Over the past 30 years, major progress has been made on how the levels and structures of correlations shape the encoding of information in population codes. Correlations influence population coding through the organization of pairwise-activity correlations with respect to the similarity of tuning of individual neurons, by their stimulus modulation and by the presence of higher-order correlations. Recent work has shown that correlations also profoundly shape other important functions performed by neural populations, including generating codes across multiple timescales and facilitating information transmission to, and readout by, downstream brain areas to guide behaviour. Here, we review this recent work and discuss how the structures of correlations can have opposite effects on the different functions of neural populations, thus creating trade-offs and constraints for the structure–function relationships of population codes. Further, we present ideas on how to combine large-scale simultaneous recordings of neural populations, computational models, analyses of behaviour, optogenetics and anatomy to unravel how the structures of correlations might be optimized to serve multiple functions.

Population code

The features and patterns of activity of neural populations that are used to perform key information-processing computations, such as encoding information and/or transmitting information.

¹Department of Excellence for Neural Information Processing, Center for Molecular Neurobiology (ZMNH), University Medical Center Hamburg-Eppendorf (UKE), Hamburg, Germany.

²Istituto Italiano di Tecnologia, Rovereto, Italy.

³Department of Neurobiology, Harvard Medical School, Boston, MA, USA.

✉e-mail: s.panzeri@uke.de; harvey@hms.harvard.edu
<https://doi.org/10.1038/s41583-022-00606-4>

The brain performs many functions that rely on the distributed and coordinated activity of populations of neurons, including perception, decision-making, memory, planning and navigation. These functions arise from the ability of populations of neurons to encode information about the sensory environment and to transmit this information to other brain structures to inform appropriate behavioural outputs.

A fundamental question is whether there are emergent functions of populations of neurons that are crucial for information processing beyond what can be performed by single neurons. Emergent functions of neural populations rely on the formation of a population code, which can be defined in two ways. First, some studies consider a population code by joining together the coding or tuning properties computed separately for many individual neurons and measuring the information available in the joint representation. One example of this definition of a population code is place cells in the hippocampus: individual place cells are tuned to specific locations in an environment, with different place cells tuned to different locations. By assembling place cells into

a population, it is possible to form a cognitive map and decode an animal's location accurately — neither of which can be done from individual neurons¹. Recently, other properties of population codes defined in this way have been identified, including the importance of the dimensionality of the space of neural population activity and the geometry of the encoding of abstract cognitive variables^{2,3}. These neural population properties depend only on the tuning or coding features of individual neurons and neglect the functional interactions between neurons. We thus refer to this type of population code as an independent population code. This type of code can be understood even by recording neurons one at a time and then placing them together into a pseudo-population (FIG. 1a).

Alternatively, a population code can be defined by the features of population activity that arise from the functional interactions within a group of neurons and that cannot be identified by examining individual neurons or by assembling a pseudo-population (FIG. 1b). In this case, the population code is based on the emergence of macroscopic properties from the microscopic interactions

Space of neural population activity

A space where each dimension represents the activity of a neuron and each point is a population vector.

Functional interactions

The statistical relationships between the activity of different neurons, often quantified as correlations between the activity of different neurons.

between individual elements, as is fundamental in many physical and biological systems. Functional interactions between neurons are commonly measured as correlations between the activity of different neurons. We thus refer to this type of code as a correlated population code. Although aspects of both independent and correlated population codes have been studied extensively, we here focus on correlated population codes because they express genuine emergent properties that are specific to real populations of simultaneously recorded neurons.

A second fundamental question regards whether the structure, or organization, of the correlations between neurons shapes information processing in correlated population codes. In theory, two populations of neurons could have similar distributions of correlation values but

differ in how these correlations are arranged between pairs of neurons within the network. As an analogy, in physical systems, the same set of carbon atoms can form a diamond or graphite depending on how their interactions are spatially organized. For neurons, the key question is whether there is an underlying structure that organizes the functional interactions between neurons in a population code and what consequences may arise from this structure for neural computations⁴. The interaction structures that are potentially important for information processing include the arrangements of activity correlations between neurons with similar or dissimilar tuning (FIG. 1c). For example, rather than the correlations being randomly arranged in a population, the correlations between neurons with similar tuning to

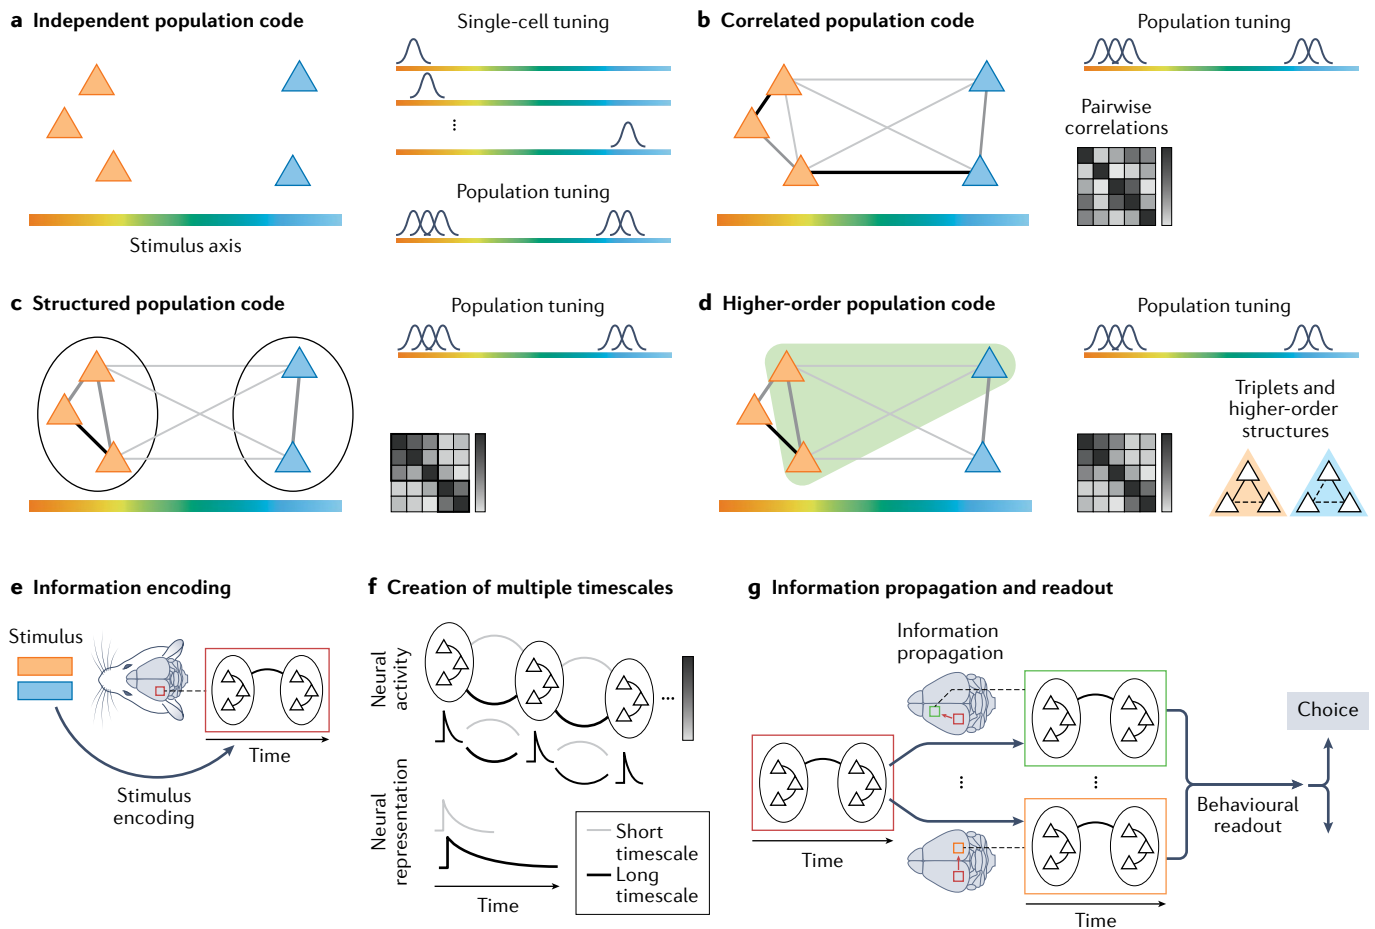


Fig. 1 | Different types of population codes and the functions they can perform. a–d | Schematic of a population of neurons (triangles) encoding information about a stimulus (whose identity is colour coded) using different types of population codes. An independent population code results from computing tuning curves and coding properties separately for each neuron, and then combining information from individual cells, neglecting single-trial interactions and correlations between neurons (part a). A correlated population code is critically shaped by the functional interactions (noise correlations) within the population. Pairwise correlations are indicated as links between pairs of neurons, and correlation strength is plotted in the pairwise-interaction matrix (part b). In a population code with structured correlations, functional interactions and correlations are structured in a nonrandom way. In this example, neurons with similar tuning also have stronger interactions (as highlighted by the outlined diagonal blocks in the pairwise-correlations matrix) (part c). Higher-order population

codes have higher-order correlations (represented graphically as shaded areas and often termed hyperedges in network theory³ because they represent structures that cannot be described as sets of pairwise edges) that contribute to population coding. Triplets of neurons in the bottom right part are coloured to indicate that the strength of their high-order interaction may be stimulus-modulated to carry stimulus information (part d). **e–g** | We hypothesize that population codes might perform multiple information-processing functions. They may encode information about sensory stimuli in a brain region (red box) (part e). They may contain correlations that contribute to the creation and maintenance of neural representations across multiple timescales, providing the flexibility necessary to perform different tasks (part f). They may contain correlations that influence the propagation of information across multiple brain areas (represented as the arrows from the red to the green and orange boxes) and its behavioural readout (part g).

Pseudo-population

Collections of activity of non-simultaneously measured neurons, either because they were recorded at a different time or from different experiments, or because they were created by trial shuffling.

Efficient coding theories

Theories that postulate that the properties of neurons in sensory areas are designed to maximize the information that these neurons carry about sensory stimuli with naturalistic features.

Signal correlations

The correlations of the trial-averaged neural responses across different stimuli.

Noise correlations

The correlated trial-to-trial variability of the activity of different neurons or of different neural populations over repeated presentations of the same stimulus.

Signal–noise angle

The angle between the noise axis and the signal axis.

Signal axis

The axis in neural population activity space of the largest stimulus-related variations, which in linear cases is measured as the axis that connects the trial-averaged population responses to the different stimuli.

Noise axis

The axis of largest variation in neural population activity for a fixed stimulus.

a sensory stimulus might be stronger than the correlations between neurons tuned to dissimilar stimuli. We define structured correlated population codes as those codes whose functions are shaped by the structure of correlations. Other potentially important structures (beyond pairwise correlations) could include the existence of higher-order interactions⁵ that critically shape information processing, leading to higher-order population codes (FIG. 1d).

Recent studies have made substantial progress in understanding how the magnitude and structure of correlations among neurons in a population relate to the functions that the neural population code can perform. These studies have identified at least three main potential functions for correlations in population codes: shaping the encoding of information; generating codes across multiple timescales; and facilitating information transmission to, and readout by, downstream brain areas to guide behaviour (FIG. 1e–g). Because each neural population can contribute to multiple functions, the structure and optimization of population codes is subject to multiple constraints. Thus, beyond optimizing the encoding of information, which was originally proposed as the goal of neural populations by efficient coding theories^{6–9}, correlations in a neural population must be organized to satisfy multiple constraints, including how this information is read out to guide behaviour. Here, we review these studies and we reason on how combining large-scale simultaneous recordings of populations of neurons, computational methods, neural perturbations and anatomy can further help to establish how population codes serve cognitive functions.

To keep this Review focused, we highlight principles of population coding using mainly examples from the encoding of sensory information for perception and decision-making. Also, we consider the recently reviewed¹⁰ field of population dynamics only as it relates to the timescales of population coding.

Encoding and correlation structure

The most studied aspect of information processing in neural populations is the encoding of information about sensory stimuli. Information encoding is usually used to refer to the amount of information present in population activity under the assumption that the population information is read out by an optimal decoder, without factoring in how it is transmitted downstream or used to inform behaviour.

Theoretical foundations. Are correlations a source of information or a nuisance for representing sensory stimuli? Theoretical work has established that the information in a population of neurons is influenced by the correlations between neurons^{11–14}.

Conventionally, a distinction is made between two types of correlations¹⁵. Signal correlations indicate the similarity of stimulus tuning of different neurons, with high signal correlations for neurons tuned to the same stimuli. Signal correlations can be defined even for independent population codes and can be measured in both simultaneously recorded and pseudo-population responses. By contrast, noise correlations measure activity

correlations beyond the stimulus tuning shared by the neurons and are often quantified as the correlation in individual trial responses between neurons for a given stimulus. For two neurons with positive noise correlation, one neuron tends to respond more strongly than usual on a trial with a given stimulus when the other neuron also responds more strongly than usual, and tends to respond less strongly than usual on trials with the same stimulus in which the other neuron's responses are also less strong. Noise correlations can be measured only in population codes of simultaneously recorded neurons and thus are the feature that distinguishes correlated population codes from independent population codes. Noise correlations can originate in various ways¹⁶ but, as discussed below, their effects on population codes are the same regardless of their origin. Here, we focus on their functional implications rather than on how they arise.

A key question about population coding is how noise correlations affect the amount of information encoded by the population. A major factor to consider when approaching this question is the relationship between signal and noise correlations. If signal and noise correlations have the same sign, the signal and noise will have a similar shape in population-activity space and will be more difficult to separate. In this case, noise correlations impair stimulus discrimination by creating an overlap between the distributions of responses to different stimuli (FIG. 2a, left) that is larger than in the absence of noise correlations (FIG. 2b). An example is when two neurons spike strongly to the same stimuli and thus have a positive signal correlation, while also having positively correlated variations in spiking from trial to trial, thus having a positive noise correlation (FIG. 2a, left). By contrast, if signal and noise correlations have different signs, such as a positive noise correlation for a pair of neurons that respond to different stimuli (negative signal correlation), then the noise correlations decrease the overlap between the response distributions to different stimuli and increase the amount of information encoded compared with the amount encoded in the absence of noise correlations (FIG. 2a, right). These effects can be geometrically described in terms of the signal–noise angle^{17,18} between the signal axis and the noise axis (FIG. 2d) in the space of neural population activity, a concept that also extends to higher-dimensional spaces with large population sizes. Information-enhancing correlations with opposite signs of signal and noise correlations correspond to large signal–noise angles, and information-limiting correlations with signal and noise correlations of the same sign correspond to small signal–noise angles (FIG. 2a–d).

The effects of the relationship between signal and noise correlations are more pronounced as the population size grows^{11,19,20}. When signal and noise correlations have the same sign, the information encoded by the whole population can be saturated^{14,19}. In such a case, noise correlations limit the benefit of averaging noisy information across neurons to the extent that they fundamentally limit the amount of information that can be encoded by a population regardless of the population's size.

A second major factor to consider relating to the question of how noise correlations affect information encoding is whether noise correlations change with the stimulus. If noise correlations are stimulus dependent, they can increase the information encoded in population activity by acting as a coding mechanism that is complementary to the firing rates of individual neurons^{11,12,21,22}. Stimulus-dependent noise correlations can act as the only information channel if the responses of neurons are similar across stimuli and only noise correlations vary. In this case, stimuli can be decoded based on jointly observed population activity even if the responses of individual neurons are similar across stimuli (FIG. 2c, left). This stimulus-dependent increase of the encoded information is separate from the signal–noise alignment discussed above and can offset the information-limiting effects of signal–noise alignment^{11,12,23} (FIG. 2c, right).

Empirical results. With the advent of technologies enabling simultaneous recordings from multiple neurons, experimental evidence has emerged that supports the theoretically possible effects of noise correlations on information encoding described above.

Initial results sparked excitement by reporting that noise correlations are stimulus dependent and add information not available in individual cells. For example, stimulus-dependent noise correlations in cat visual cortex were found to carry information about whether individual neurons respond to the same or separate objects²⁴. Other examples of stimulus-dependent noise correlations serving to enhance information coding have also been reported^{25–29}. In other cases, noise correlations were reported to be stimulus independent yet nevertheless to increase the information encoded in neural activity because the signs of the signal and noise correlations

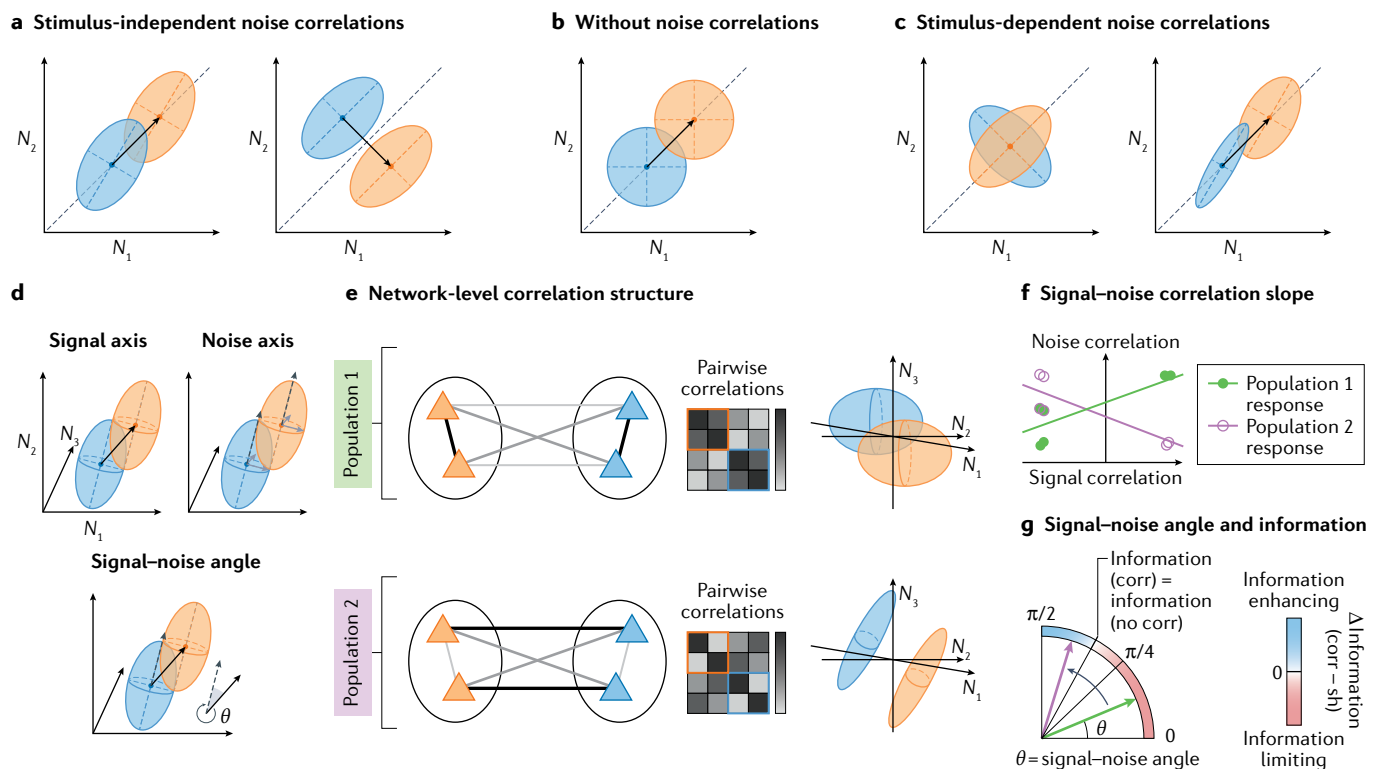


Fig. 2 | Noise correlations and information encoding. Illustration of how noise correlations shape stimulus–information encoding. We use cartoons of response distributions (ellipses) of populations of neurons (N_1 , N_2) in the space of neural population activity (using $n = 2$ or 3 neurons) in response to two stimuli (orange and blue). **a** | Stimulus-independent noise correlations can decrease (left) or increase (right) the amount of encoded information about stimulus (overlapping distributions being harder to discriminate) with respect to population activity without noise correlations (as in part **b**). **b** | Here, the properties of the individual neurons are the same but noise correlations are absent. **c** | Stimulus-dependent noise correlations, that vary in sign and/or strength across stimuli, might provide a separate channel for stimulus information encoding (left) or reverse the information-limiting effect of stimulus-independent noise correlations (right). **d** | In a high-dimensional space, the interplay between signal and noise can be described using the signal axis, which connects the average responses to different stimuli (solid arrow); the noise axis, which represents the direction of maximum variability of a response to a fixed stimulus (dashed arrow); and the signal–noise angle, which is the high-dimensional angle between the signal axis and noise axis. **e** | Example of two neural populations with

the same signal and noise correlation distribution but with a different noise correlation structure. Population 1 (population 2) has higher noise correlations for neurons with similar (dissimilar) tuning, resulting in a larger (smaller) overlap in the distributions of neural responses in high-dimensional population space (ellipsoids on the right). **f** | The network-level structure of noise correlations plotted in part **e** results in a positive (negative) signal–noise slope when noise correlations are higher for similarly (dissimilarly) tuned neurons. Green (purple) dots and lines represent the pairwise signal correlation and noise correlation values and slopes of the population in population 1 (population 2) of part **e**. Purple-outlined green dots are plotted where green and purple dots coincide. **g** | A network-level structure with higher noise correlations for similarly (dissimilarly) tuned neurons results in smaller (larger) signal–noise angles, represented by the green (purple) arrow, and in an information-limiting (information-enhancing) effect of noise correlations, illustrated by a colour bar plotting the difference of information encoded by the population with correlations intact ('corr') or after shuffling ('sh'). Green (purple) colours denote the results for population 1 (population 2) in part **e**. Results in parts **f** and **g** are schematic, intended to illustrate the sign of the effects, and are not quantitative.

Trial shuffling

An analytical procedure to remove the effect of noise correlations by combining responses of neurons taken from different trials to a given stimulus.

Redundant neuron pairs

Pairs of neurons that together carry less information than the sum of the information carried by the two neurons in each pair, owing to the information-limiting effect of noise and signal correlations.

Synergistic neuron pairs

Pairs of neurons that together carry more information than the sum of the information carried by the two neurons in each pair, owing to the information-enhancing effect of noise correlations.

were opposite^{30,31}. Other studies have shown that noise correlations increase information without explicitly identifying whether the information increase was due to stimulus dependence, signal–noise misalignment or both^{32,33}.

However, it has been commonly found that neurons with similar stimulus tuning have positive and strong noise correlations^{4,29,34–36}. As a consequence, in many cases, noise correlations limit the information encoded in a neural population^{4,37–44}. This conclusion has been reached by comparing the information contained in a simultaneously recorded neural population, with noise correlations intact, to a pseudo-population code with noise correlations removed by trial shuffling or other analytical means (BOX 1).

As predicted by theoretical models, noise correlations affect how the information encoded in a population increases with population size^{38,42}. For small populations, the effects of noise correlations have been found to be small^{13,17,31,39}, leading to the suggestion that noise correlations only weakly modulate the information carried by the tuning properties of individual cells. However, more recent studies of large populations found large information-limiting effects of noise correlations that lead to the saturation of information levels^{37,38,42}. Thus, noise correlations have effects on information encoding that can be as large as the effects of the tuning properties of the individual cells, and can impose fundamental limits on the amount of information that can be encoded in a population³⁸.

From the many empirical studies suggesting an information-limiting effect of noise correlations and from the accompanying theories^{19,45}, an influential view has emerged that noise correlations often limit the amount of information encoded in a population and thus impair perceptual accuracy^{19,46}.

Network-level organization of pairwise noise correlations. Much of the founding theoretical work on the role of noise correlations in information encoding has focused on the average value of correlations in a typical pair of neurons or has assumed that the properties of single cells and/or correlations are homogenous across the population. However, pairwise correlations span a wide range of values in a population, including positive and negative values^{37,42,43,47,48}. This raises questions about whether the distribution of correlations and their organization between pairs affect information encoding. Indeed, recent experiments recording from hundreds to thousands of neurons, together with theoretical developments, have revealed that the organization of correlations, and not only their average value, matters for encoding, providing evidence for population codes with structured correlations.

Mathematical analysis of network models has shown that heterogeneity of both the correlation values and the tuning properties across neurons in a population can reduce the information-limiting effect of noise correlations^{20,49–51}. Consider two populations with similar distributions of pairwise noise correlations (for example, those schematized in FIG. 2e): in the population with higher heterogeneity, correlations have less of an

information-limiting effect, because the heterogeneities reduce the probability that the noise variance aligns with the signal variance. Thus, separating signal from noise and decoding stimulus information from population activity may be easier to do in heterogeneous populations than in a homogeneous population²⁰.

The extent to which correlations limit information encoding depends on how tightly signal and noise correlations are aligned on a pair-by-pair basis. This association can be quantified as the slope of the graph plotting signal correlations versus noise correlations (FIG. 2f). A higher positive value for the slope is present when neuron pairs with high positive signal correlations have high positive noise correlations. For a given average value of noise correlations in a population, a higher value for the slope is associated with larger information-limiting effects (FIG. 2g).

The importance of the signal–noise slope is demonstrated by studies that show that improvements in behavioural performance with learning, attention, adaptation or task engagement correlate with changes in population activity. In some cases, changes in the properties of noise correlations account for improvements in behavioural performance more than changes in the selectivity of individual neurons do⁵². For small population sizes (handfuls of neurons), the change in noise correlations that best accounts for improvement in behavioural performance is a decrease in the values of the signal–noise slope^{53–56}. This change in signal–noise relationship results in a larger increase in encoded information than do behaviour-related changes in the average amplitude of noise correlations⁵⁷. A decrease in the value of the signal–noise slope and in the information-limiting effect of noise correlations is also observed when cholinergic modulation in the cortex is enhanced through optogenetic stimulation of basal forebrain cholinergic neurons, suggesting that such modulation is a potential mechanism for the changes in population encoding that are associated with attention and learning⁴⁹. Thus, the signal–noise slope, in addition to noise-correlation magnitudes, is one of the crucial factors that control information encoding in populations.

The effects on information encoding of the signal–noise slope and of the heterogeneity of correlations can be recapitulated geometrically in terms of shifts of the signal–noise angle. The information-limiting effect of noise correlations depends on how much of the noise variance ends up on the signal axis⁴⁵: the proportion of noise variance that ends up on the signal axis is highest when the signal–noise angle is small. Both high neural heterogeneities and correlation structures with lower-value signal–noise slopes shift a large part of the variance of noise correlations away from the signal axis, creating a large signal–noise angle and small information-limiting effects. Indeed, recent measurements of variance in large neural populations show that most correlated variance lies outside the signal axis, suggesting that the structure of noise correlations limits the amount of noise variance along the signal axis^{37,42}.

Finally, besides the redundant neuron pairs often reported in previous studies, a recent study found a prevalence of synergistic neuron pairs, which encode

Box 1 | Analytical tools to remove the effect of noise correlations from population codes

Quantifying the effects of noise correlations on population coding requires comparing the information encoded in the original data with that encoded in surrogate distributions of neural activity obtained after removing noise correlations while preserving other aspects of population activity¹³. Similar approaches can be used to estimate the effect of noise correlations on task performance⁶⁰.

A widely used way to generate pseudo-population activity vectors with noise correlations removed is to shuffle data. Shuffling population activity vectors across trials of the same stimulus, independently at each time point, destroys across-time noise correlations. Shuffling the trial labels, independently for each neuron, across trials of the same stimulus disrupts across-neuron noise correlations (see the figure, part a). Trial shuffling is simple and non-parametric, but it can be applied only when trials can be meaningfully grouped by stimulus.

Other approaches include fitting data to parametric models of population activity distributions. These models can generate probabilities of spike trains with arbitrary values of noise correlations, or remove specific aspects of noise correlations. These approaches can also be applied when stimuli cannot be categorized, but they make assumptions about the data. Such parametric models are described below (and see the figure, part b).

A caveat of removing rather than disrupting noise correlations is that some effects of the functional interactions between neurons probably remain in each neuron's recorded activity. Thus, disruption analyses probably underestimate how much functional interactions contribute to population coding⁸⁸.

Maximum entropy (ME) models

The ME model describes neural population activity with a probability distribution that has maximal entropy among the set of distributions that satisfy a set of constraints on the statistical properties of data¹⁶⁵. Imposing data constraints that contain interactions between neurons up to order k will give an ME model that captures interactions up to order k .

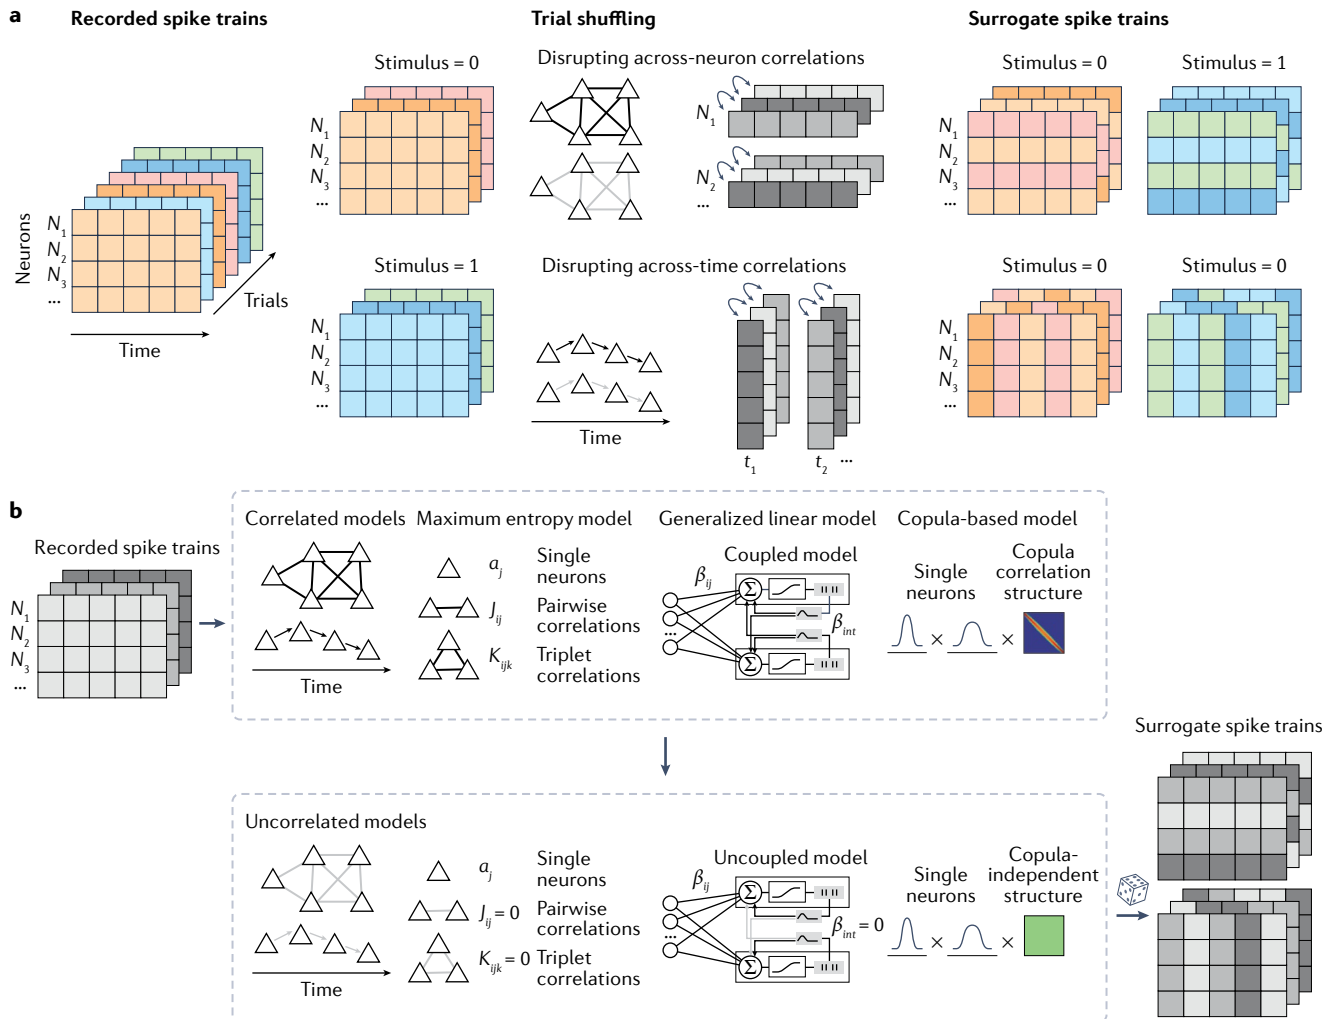
Coupled generalized linear models (c-GLMs)

c-GLMs model the conditional probability of populations of neurons expressing the input to each neuron as the sum (Σ) of linear filters, and then apply neural-like non-linearities to generate firing rates. Input filters include task-related filters that express contributions from sensory and behavioural variables and coupling filters that describe the effect of the firing of other neurons at previous times (across-neuron links with backward arrow in the figure)^{32,88}.

Copula-based models

These^{166,167} model the population activity probability as a product of the single-neuron marginal distributions and their interaction structure ('copula'). The copula can be fitted to data parametrically¹⁶⁶ or non-parametrically¹⁶⁸. Changing the copula structure generates data with different levels and structure of correlation.

a_j , ME coefficients of single-neuron activities; β_{ij} , GLM weights of different predictors; β_{int} , GLM weights of pairwise interactions; J_{ij} , ME coefficients of pairwise interactions; K_{ijk} , ME coefficients of triplet-wise interactions.



Redundant hubs

Neurons with high probability of having redundant interactions with other neurons.

Synergistic hubs

Neurons with high probability of having synergistic interactions with other neurons.

Population-wise correlations

Correlated variability of an entire population of neurons, usually measured applying dimensionality-reduction techniques to the population covariance matrix.

Across-neuron noise correlations

The noise correlation between the time-averaged activity of two different neurons or two different neural populations, quantifying the similarity of the time-averaged neural or population responses across trials with the same stimulus.

Across-time noise correlations

The noise correlation between the population activity vector of the same population at different times, quantifying the similarity of the population responses at different times across trials with the same stimulus.

Persistent activity

The activity of individual cells whose firing rate remains sustained over an entire task period, for example, during working memory or decision-making tasks.

Ramping activity

The activity of individual cells whose firing rate decreases or increases constantly over time during a task, for example, to reflect the accumulation of evidence to make a decision.

Attractor states

Set of values of population vectors towards which the activity of a neuronal network is attracted during its temporal evolution.

Posterior parietal cortex

(PPC). A region of cortex considered to be at the interface of sensation and action and to participate in evidence accumulation for decision-making, movement planning, spatial navigation and other processes.

information through stimulus-dependent noise correlations⁵⁸. This study also identified redundant hubs that preferentially interact with other redundant hubs, as well as synergistic hubs that preferentially interact with other synergistic hubs, suggesting a non-uniform network-level structure in how noise correlations are organized. This structure combines the advantages of synergy and redundancy by creating a synergistic network that exploits the stimulus dependence of noise correlations to maximize encoded information, and a redundant network that offers a backbone of robustness in the face of the loss of function of some cells.

Together, these studies have begun to reveal that correlations in a population are structured. The structures of correlations may organize synergistic and redundant encoding in specialized subnetworks and mitigate the negative effect of noise correlations on information encoding.

Population-wise and higher-order correlations. The above results address pairwise correlations between neurons and their organization within the network. However, in recent years, studies have begun to move from detailing these relationships at the pairwise level to characterizing the strength, properties and consequences of correlations on encoded information directly at the population level with population-wise correlations^{59,60}. Different mean values of pairwise correlations can have profoundly different effects on population-wise correlation strengths, depending on the variability of the values of pairwise noise correlations across the population⁶¹, again confirming the importance of heterogeneity for population codes.

The effect of noise correlations on population information is measured directly at the population level by comparing the amount of information in the population with correlations present and with correlations removed analytically (BOX 1). These direct quantifications potentially capture all effects of correlations on population encoding, including the stimulus dependence of noise correlations; the presence of small information-limiting correlations that may be difficult to detect, especially for continuous stimulus variables, and yet have a major effect on information for large populations; and all pairwise signal–noise relationships.

One emerging result from direct studies of population-wise correlations is that neurons can interact in larger groups in ways that cannot be decomposed into combinations of pairwise interactions. These higher-order correlations are important to explain the statistics of neural population responses, particularly for nearby neurons, and in executive areas of cortex^{62–69}. For example, for nearby neurons in monkey⁶⁴ and mouse⁶⁷ visual cortex, higher-order correlations make codes more sparse by suppressing periods of intermediate firing, increasing periods of quiescence and concentrating population activity in time. The effects of higher-order correlations are particularly prominent for larger populations^{62,70}, suggesting that future studies of large-scale recordings of hundreds or thousands of neurons should systematically consider these higher-order correlations.

However, the question of how the higher-order structure of noise correlations affects encoding is still largely unexplored. A theoretical study has revealed that the principles of encoding with pairwise noise correlations are maintained for higher-order correlations, in particular with the stimulus dependence of higher-order correlations having a potentially large impact on encoding⁷¹. Current empirical evidence suggests that higher-order correlations, like pairwise correlations, have largely information-limiting effects for information encoding in sensory cortices^{64,65}.

Correlations and encoding timescales

Noise correlations have often been considered as across-neuron noise correlations between the time-averaged spike rates of different neurons. The importance of across-time noise correlations between the activity of populations of neurons at different times is less well understood than that of across-neuron noise correlations. To produce behaviours, the brain needs to represent different types of information with timescales that vary widely. For example, tracking the rapid fluctuations of sensory stimuli requires information encoding at timescales of a few tens of milliseconds, a computation that individual neurons in sensory areas can perform⁷². Producing behavioural choices may instead require accumulating and maintaining consistent information over seconds. The single-neuron and network mechanisms responsible for the creation of long timescales have been debated, and a role for across-time noise correlations is beginning to emerge.

Studies of neural activity during working memory and perceptual decision tasks have proposed that long timescales can be encoded at the single-neuron level. Neural recordings in macaque parietal and frontal cortices have commonly found cells with persistent activity during decision-making and working memory^{73–76}. Other studies have identified neurons with ramping activity during the accumulation of sensory evidence^{73,77}. Computational models have proposed that persistent activity reflects the instantiation of attractor states for short-term memory or categorical choices, and that ramping activity reflects the accumulation of evidence towards a decision threshold⁷⁸. Related activity patterns have been identified in various parts of the mouse cortex^{79–82}.

Other studies have proposed that long-lasting choice and memory signals can arise at the neural population level. For example, individual neurons can each be active for only a fraction of a task period but, across the population, sequences of neural activation can arise for specific choices, memories and navigation plans^{83–85}. When considering tasks that combine a perceptual decision and/or evidence accumulation with memory-guided navigation, individual neurons in mouse posterior parietal cortex (PPC) typically form a sequence of brief (about 100–200 ms long) neuronal activations that tile the entire length of a task trial^{86–88} (FIG. 3a). In these cases, long timescales must be implemented in population codes rather than in single neurons.

It was recently shown⁸⁸ that, in the presence of sequential population activity, across-time noise correlations^{87–89} contribute to the generation of long timescales in

Population vector

Vector in the space of neural population activity whose components represent the activity of individual neurons in the population.

Information consistency timescale

The correlation across time of the instantaneous stimulus or choice signal (for example, the posterior probability of stimulus or choice given the observation of a population vector at a specific time).

population codes. During a sound-localization task performed by mice navigating through a virtual reality maze⁸⁸, PPC population activity showed across-time noise correlations with lags of up to about 1–2 s (FIG. 3b). The information obtained by integrating the population vector across time points increased when across-time noise correlations were disrupted by trial shuffling, suggesting that across-time correlations limit the benefit of integrating noisy information over time. Thus, these noise correlations are key for a speed–accuracy trade-off, because they set the time window after which there is no advantage in further integrating information^{90,91}. Across-time correlations also increase the information consistency timescale of choice information in the PPC population, as considerably longer information consistency timescales were obtained from real population responses than from pseudo-populations (FIG. 3c). Similar results were also present during another task in mice, in which sequences

of neurons represented the accumulation of visual evidence for making navigation decisions⁸⁷. Interestingly, the timescale of choice information consistency in PPC was longer on correct trials than error trials. Long timescales in correlated population codes may therefore be important for conveying signals relevant for accurate behaviour.

The timescales of population codes differ across areas. In the same sound-localization task described above, the auditory cortex had weaker across-time correlations and a shorter information consistency timescale than PPC. These findings raise the possibility that a function of population codes is to set the timescale for computation. In sensory cortices, a short timescale might aid representations of rapidly fluctuating stimuli and high-dimensional sensory features. By contrast, long timescales of information consistency may enable information in temporally separate inputs to be combined; enable the integration of new information with ongoing activity dynamics⁸⁷; facilitate the formation of consistent information for planning and making a choice; and make representations more invariant to categorization-irrelevant transformations of sensory inputs⁹². Across-time correlations may be particularly important when sequential, transient, single-cell activity is present, as might be common in tasks with many spatial and temporal components, natural behaviours and navigation⁸⁶. Across-time correlations may create consistent pathways of neural activity in this high-dimensional space that enable evidence accumulation or decision implementation⁸⁷.

In sum, across-time noise correlations may provide advantages for behaviour by lengthening the timescales of information consistency, even if they concurrently reduce the amount of encoded information. This suggests that noise correlations serve multiple functions and that population codes balance constraints besides optimizing information encoding.

Are codes read out optimally?

The fact that often noise correlations limit the information-encoding capacity of neural populations has been taken to imply that they hinder the ability to perceptually discriminate sensory stimuli. This conclusion holds if perceptual discrimination performance increases proportionally with the amount of sensory information encoded in neural activity^{46,93,94}. The latter would be true when sensory information in population activity is read out optimally to inform behaviour. However, the amount of information that is transmitted and used for behaviour depends on the intersection between the properties of encoding and readout⁹⁵. If the readout of the encoded information is not optimal, and in particular if the readout depends on correlations in a way that differs from encoding, then population codes with higher information may not necessarily generate more accurate behaviour^{13,17,94–99}. For example, a population encoding large amounts of information in stimulus-dependent pairwise correlations or higher-order correlations may carry more information than do other codes that do not rely on correlations, but this extra information can only be accessed downstream in the presence of nonlinear

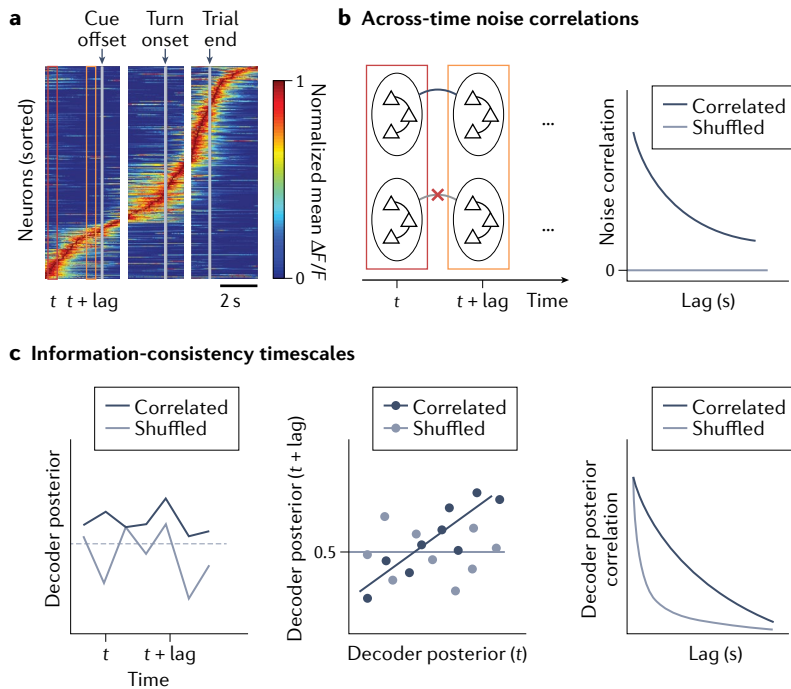


Fig. 3 | Across-time noise correlations can generate long information consistency timescales. **a** | In many experiments, neurons are transiently active for a fraction of the task period. Different neurons are sequentially active, tiling the task time, with specific sequences of activation that can arise for specific choices. Data recorded from posterior parietal cortex in a decision-making task during spatial navigation⁸⁶. **b** | Long information timescales in population codes may be created by correlations across times between the activity of different neurons. Long-lag across-time noise correlations between different neurons and different populations are present in experimental data (depicted here schematically)⁸⁸. **c** | Across-time noise correlations between different neurons lead to across-time correlations between population vectors at different times, which in turn leads to an increase of information-consistency timescales of population codes. As shown here schematically, for simultaneously recorded (correlated) responses⁸⁸, across-time correlations ensure that the time course of the instantaneous posterior probability of stimuli given the population vectors (obtained, for example, using a stimulus decoder based on population activity) is smoother (left), and the correlation of posteriors obtained at different lags is higher for a given lag (middle) and has a longer decay time constant when computed across lags (right). Longer information-consistency timescales generate stable representations that might contribute to accurate behaviour. Part **a** is reproduced from REF.⁸⁶, Springer Nature Limited. Parts **b** and **c** are adapted from REF.⁸⁸, Springer Nature Limited. F, fluorescence; t, time.

Box 2 | Suboptimalities in the readout of the encoded information

Here, we consider methods and evidence to establish whether sensory information encoded in population activity is read out optimally to inform behaviour.

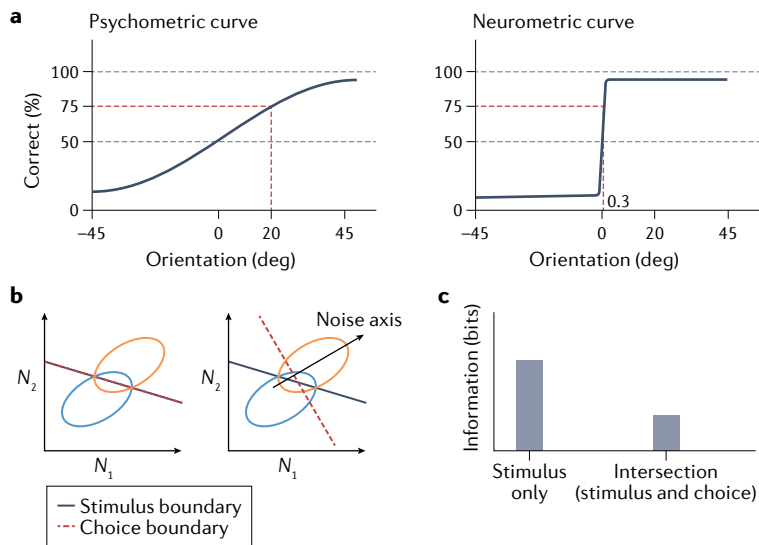
If the discrimination of sensory stimuli from neural activity is more accurate than the behavioural performance, then the information encoded in population activity is not read out optimally (see the figure, part a). A recent study reported that information about stimulus orientation in large populations in mouse primary visual cortex (V1) affords stimulus-discrimination thresholds 100 times better than the behavioural discrimination threshold¹⁰¹, implying that perceptual discrimination may be limited by the suboptimality of downstream decoders and not by noise correlations in sensory areas. However, another study reported that information in large mouse V1 populations saturates at levels similar to those reflected in behavioural discriminations⁴².

A further test of suboptimality relies on comparing the discrimination boundaries of optimal stimulus or choice decoding from the same neural population⁹⁵. For an optimal readout, the two boundaries are aligned (see the figure, part b, left). For a suboptimal readout, the two boundaries are misaligned (see the figure, part b, right). N_1 , feature of neural population activity. Recent studies in sensory cortex have found that the choice-discrimination boundary is not aligned to the optimal stimulus-discrimination boundary^{59,102} and, instead, that it can lie along the noise axis⁵⁹. This suggests both that the behavioural readout of the encoded information is not optimal and that noise correlations may affect the formation of choices.

A related way is to compute the fraction of the sensory information encoded in neural activity that is used to inform choices (see the figure, part c). If this fraction is small, then the readout is suboptimal^{60,103}. In sensory and association cortices, only a relatively small fraction of the sensory information encoded in neural activity is used to inform choices, again indicating that the readout is suboptimal^{60,103}. These evaluations are valid under the assumption that the choice information in population activity causally contributes to behavioural choice. This may not always be true, especially for choice information available after sensory evidence is integrated and a behavioural choice is expressed¹⁶⁹. Thus, care should be used in computing these quantities — as done, for example, in REFS^{60,103} — in post-stimulus time windows before the start of choice expression.

readout mechanisms that learn the joint probabilities of groups of neurons rather than of a simpler, suboptimal readout that linearly weights the activity of individual neurons²¹. Conversely, population codes with information-limiting noise correlations may carry less information than do population codes without noise correlations but the correlations might help downstream readout to aid behaviour¹⁰⁰. In the presence of readout suboptimalities, it is thus theoretically possible that correlated population codes have advantages for perceptual discrimination if such codes are more efficiently read out, despite encoding less information.

Despite the major consequences that a readout suboptimality may have, little is known about whether downstream networks can use all the encoded information they receive. Recent studies^{59,60,101–103} have begun to provide analytical means and empirical evidence to assess if the information encoded in a population is read out optimally to inform behaviour (reviewed in BOX 2). Together, these recent results suggest that the readout of information from a neural population may contain suboptimalities. It is therefore important to study the form of the readout, the features of the downstream



Determining the choice boundary causally using optogenetics may help identify neural choice signals that are consequential to choice¹³⁶.

A further caveat is that these indications often assume that the subject is trying to perform the behavioural task optimally according to the experimenter's design. However, animals might use suboptimal strategies from an experimenter's point of view, because they have not learned the task fully, have goals distinct from those wanted by the experimenter or have shifts in attention and engagement. It thus remains possible that sensory information is read out optimally, but that subjects combine this information with information related to shifting behavioural goals and inaccurate internal models of the task rules.

One possible way to address this caveat is to compare the sensory information encoded at subsequent stages of neural processing — for example, in V1, V2 and V4 — to understand where and how information is lost along a processing hierarchy. These quantifications would require experimental methods to record large populations of neurons and analytical methods to estimate the information plateau in the complete population beyond even the size that can currently be experimentally recorded³⁸.

biological circuitry that introduce readout suboptimalities and how readout suboptimalities are affected by noise correlations.

Information transmission and behaviour


Correlated population codes have typically been studied from the perspective of information encoding but, ultimately, the importance of these codes depends on how they are transmitted to downstream brain regions and used to guide behaviour. This issue is critical if the reading out of population codes is suboptimal. Thus, a major question concerns whether correlations in neural populations help or hinder the propagation of signals to downstream networks. In principle, correlations could aid the transmission of information even if they limit the encoding capacity.

Biophysical factors influencing readout. A major theme in the field of neuronal cell biology has been how the spatial and temporal patterns of synaptic inputs on dendrites influence the spiking output of neurons. Biophysical studies have shown that integration time constants of neurons are often short (5–10 ms)¹⁰⁴. In such conditions,


Optimal stimulus-discrimination boundary
The plane (or surface) in the high-dimensional space of population activity that optimally separates responses elicited by different sensory stimuli, and that thus serves as an indication of how to extract sensory information from neural activity optimally.

a Information transmission

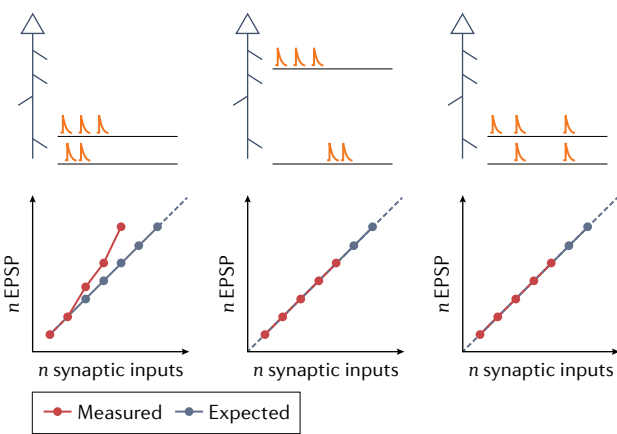
High noise correlation strength



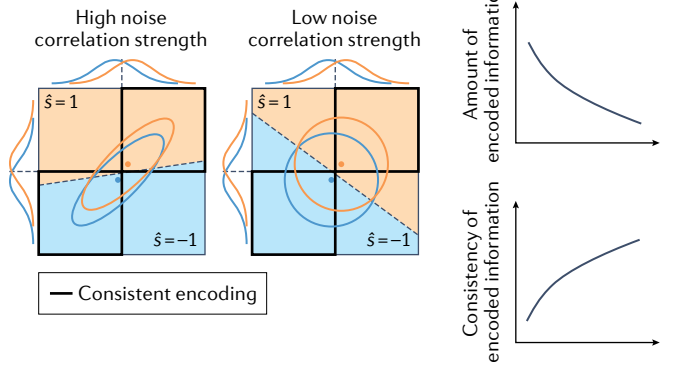
Low noise correlation strength



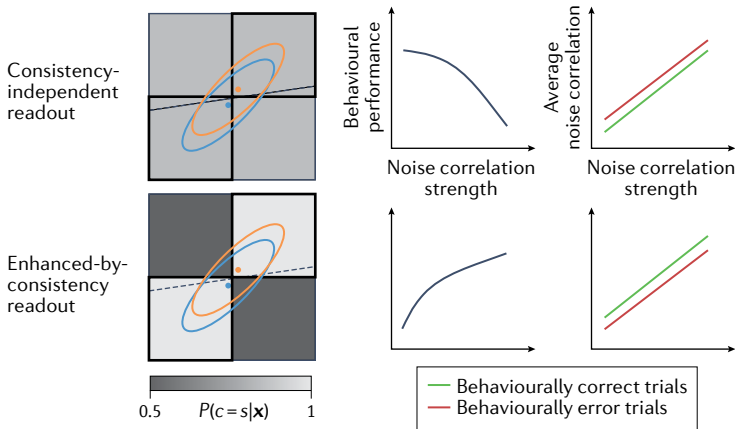
b Spatio-temporal integration




c Information encoding




d Information readout



e Optogenetic perturbation

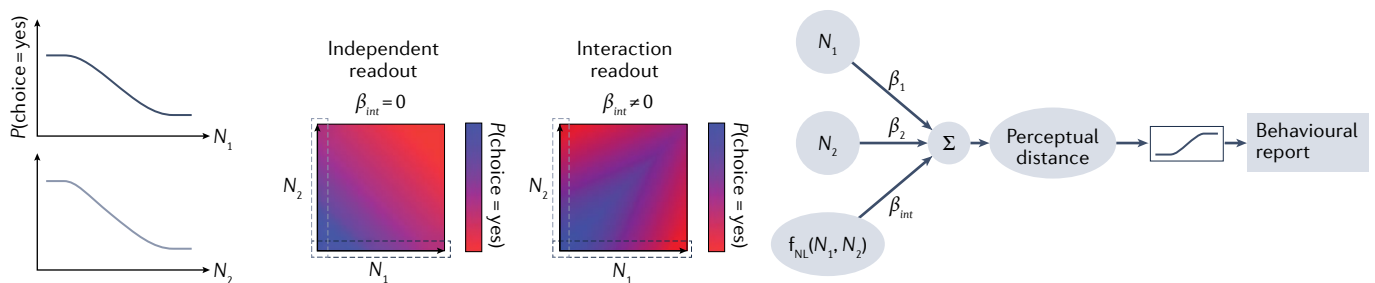


f Across-time correlations favour behaviour



g

Single feature manipulation Multiple feature manipulation



◀ Fig. 4 | **Noise correlations for information transmission and behavioural readout.**

Across-time and across-neuron noise correlations in population codes can enhance the propagation and readout of information. **a** | Transmission of information through coincidence detection. A neuron with a short integration time constant (red window) that receives stimulus-modulated presynaptic input spike trains from different neurons ($n = 2$ in this example) will produce a larger output firing rate when receiving correlated inputs (top left) than when the inputs are not correlated. This is because correlations increase the number of time points with coincident input spikes compared with when receiving inputs with the same individual-neuron characteristics but no correlations (bottom left). Even when input correlations are information limiting (input information decreases with correlation strength; top right), the transmitted information can increase with the level of input correlations (bottom right). **b** | At the dendritic level, spatiotemporal correlations in synaptic inputs generate evoked postsynaptic potentials (EPSPs) in a supra-linear manner (left), thus enhancing the spiking output of the readout neuron, whereas only spatial correlations or only temporal correlations do not generate dendritic supra-linearities (middle and right). **c** | The readout of stimulus information from population activity may be enhanced by correlations. A plot in two-dimensional population activity ($n = 2$ neurons) shows that information-limiting correlations (signal aligned with noise) increase the proportion of trials with consistent information encoding (when both neurons encode the same stimulus, top-right and bottom-left quadrants of the 2D plot). With information-limiting correlations, increasing the correlation strength decreases the information but increases the consistency of information encoding over trials. **d** | Top: if the readout is not affected by consistency (uniform readout efficacy in converting the encoded stimulus into choice, shown as uniform grey in the 2D neural space plot), noise correlations also impair behavioural performance and are higher in error trials. Bottom: if the readout is enhanced by the consistency of the encoded information (the readout leverages an interaction between stimulus information and consistency, leading to higher readout efficacy in the consistent regions of the 2D neural space plot), noise correlations have a net positive effect leading to an increased behavioural performance and higher correlations in correct trials compared with error trials. **e** | Manipulating activity patterns in a select brain area (red box) through two-photon holographic optogenetics while monitoring behavioural outcomes can establish the causal effect of noise correlations on the behavioural readout of population activity. **f** | A recent study¹³⁵ showed that mice better report the presence of a two-photon optogenetic stimulation pattern with a certain firing rate level when it has a higher degree of temporal synchrony, supporting that across-time correlations favour behavioural readout. **g** | Combining independent perturbation of multiple population-activity features with computational modelling can reveal the role of the interaction between features on the readout. Here, two features (timing of activation of two populations) are independently manipulated. The independent manipulation of two features can reveal whether the behavioural readout does or does not depend on interactions between features (middle), whereas single-feature manipulations would be unable to distinguish whether the readout does or does not depend on interactions (left). From these data, a mathematical readout model (right) could be formulated to describe how different features interact to causally generate behaviour. Experimental results in the olfactory bulb show the presence of strong interactions between timing of activation of different populations¹³⁶. β_1 , readout weight of single feature activity; β_{int} , readout weight of interactions between features; c , choice; f_{NL} , function modelling nonlinear interaction between neural features; N_1 , feature of neural population activity; s , presented stimulus; \hat{s} , stimulus encoded in the neural population activity; x , combination of all (neural and non-neural) predictors of behavioural choice. Part **a** is adapted from REF.⁶⁰, Springer Nature Limited. Part **b** is adapted with permission from REF.¹¹⁰, Elsevier. Parts **c** and **d** are adapted from REF.⁶⁰, Springer Nature Limited. Part **f** is adapted with permission from REF.¹³⁵, Elsevier.

spatially and temporally correlated spiking in the presynaptic inputs to a neuron elicit larger firing rates in postsynaptic neurons through coincidence detection^{100,105–109} (FIG. 4a). Correlations in input spike trains, for a given level of average input firing rate, increase the frequency by which enough input spikes are present in the short integration period to reach the threshold for firing, and so increase the output firing rate. In addition, dendrites have nonlinear integration properties. When synaptic inputs enter dendrites close in space and time, they can sum in a supra-linear manner, owing to the generation of a dendritic spike^{110–113}. This supra-linear summation

(FIG. 4b) enhances the spiking output of the neuron and leads to calcium influx that can trigger plasticity events^{110,114}. These factors are thought to contribute to orientation tuning in the primary visual cortex (V1) and to grid cell firing^{115–117}. Thus, from the perspective of a readout neuron, correlated synaptic inputs can have a supra-linear effect on spiking, and thus correlated population codes may enhance the propagation of signals to downstream neurons.

Readout efficacy may depend on how information is encoded. Principles from the field of correlated population codes can be applied to understanding how correlations affect the information contained in the inputs to a postsynaptic neuron, and the field of dendritic integration addresses how a neuron's synaptic inputs result in its spiking output and thus in signal propagation. Although these fields have largely been studied separately, bringing them together offers the opportunity to understand how information is encoded and propagated in a network.

Recent work has begun to bridge these areas. Studies showed that feedforward networks with high convergence faithfully transmit stimulus-dependent across-time correlations and that, as a consequence, populations of mitral cells in the mouse olfactory system can use information encoded through temporal correlations in their inputs to localize odour sources¹¹⁸. Another study showed that correlations in presynaptic inputs make the propagation of information through the postsynaptic neuron more efficient specifically when the correlations are information limiting¹¹⁹. When input correlations are information limiting and output spikes in a readout neuron are generated through coincidence detection, the accuracy of information transmission is higher when the inputs show consistent information encoding⁶⁰. This enhancement is strong enough to offset the decrease of information in the inputs, such that more information is transmitted to the output in the presence of input correlations⁶⁰.

Effects of signal propagation on behaviour. These studies suggest that noise correlations, when they are information limiting, may enhance signal propagation and thus have an effect beyond limiting the encoding of sensory information. However, the extent to which correlation-aided signal propagation affects behaviour has remained debated¹²⁰. A recent study⁶⁰ investigated this issue by analysing population activity in PPC recorded while mice performed perceptual discrimination tasks^{87,88}. In these PPC data, both across-time and across-neuron correlations reduced the sensory information encoded in neural activity. Under the traditional view that choice accuracy is proportional to the amount of information in a neural population and that correlations are detrimental to perceptual behaviours because they decrease information, one would expect noise correlations to be lower when mice make correct choices and higher when mice make errors. However, contrary to this expectation, both across-time and across-neuron noise correlations were higher in PPC in correct trials than in error trials⁶⁰.

These observations can be reconciled by hypothesizing that, as suggested by previous reports indicating that

Coincidence detection

Spike-generation mechanism that, because of the neuron's short integration time constant, requires the near-simultaneous occurrence of several input action potentials to generate an output action potential.

Consistent information encoding

When different elements of a population code (for example, the activity of different pools of neurons) all signal the presence of the same stimulus.

Across-time encoding consistency

When population activity at a given time signals the same stimulus as the population activity at another time.

Across-neuron encoding consistency

When the activity of separate neuronal pools in the same time window signals the same stimulus.

Feature amplification motifs

Motifs of cells with similar tuning that functionally excite one another to increase the signal contained in the neural population as revealed by anatomical connections or influence mapping.

the noise axis aligns with the choice boundary in neural activity space⁵⁹ and by the above-discussed results on biophysical signal propagation, the readout of stimulus information might use aspects of population activity imposed by correlations (FIG. 4c,d). Intuitively, correlations imply greater consistency in information encoding in a population. Modelling a mouse's choices⁶⁰ revealed that these depended not only on the stimulus information decoded from the PPC population activity, as predicted in the case of an optimal readout, but also on the interaction of the stimulus information decoded from PPC activity with across-time encoding consistency or across-neuron encoding consistency. Specifically, when consistent encoding was present, the stimulus information decoded from neural activity had an amplified effect on the mouse's choice. By using the model of mouse choices validated on PPC data to estimate how well the mouse would have performed on the task with and without correlations present, it was predicted that correlations in PPC benefited task performance even when they decreased sensory information⁶⁰. This was because correlations increased encoding consistency, and consistency enhanced the conversion of sensory information into behavioural choices.

Together, these results suggest that noise correlations can benefit behaviour by enhancing signal propagation

and that this can offset the negative effects that they have on encoding.

Higher-order correlations. Current evidence indicates that higher-order correlations aid information transmission and readout beyond the pairwise correlation effects discussed above. In V1, higher-order correlations help to concentrate information into brief periods, which may facilitate signal propagation as discussed above^{64,67}. Further, the strength of higher-order correlations within and across areas correlates with task accuracy, being stronger in behaviourally correct trials than in error trials^{121,122}. These findings are also supported by a theoretical study that proposes a beneficial role of higher-order correlations for input–output information transmission¹²³.

These results corroborate the view that population codes perform multiple functions relevant to behaviour besides information encoding. Ultimately, these results imply that the basis of perceptual abilities can be studied only by considering the intersection of how information is encoded and how it is propagated and read out at the level of single behavioural trials.

Using perturbations to probe function

Most work has studied the potential functions of correlated population codes by focusing on the statistical relationships between the information encoded in neural populations and the behavioural choices of animals. However, to understand the causal functions of correlations in neural populations, it is necessary to manipulate neural activity while monitoring behaviour.

New technology based on two-photon patterned optogenetics is emerging that enables the creation of artificial spatial-temporal patterns of activity in a neural population. This technology raises the possibility of manipulating specific features of the neural code to test how they affect behaviour^{124–131}. This technology has been used for influence mapping, which measures how spikes added to one or a small group of neurons causally affect the spiking of neighbouring neurons, including neurons with similar or dissimilar tuning^{132,133} (BOX 3). These approaches have also been used to create fictive percepts to test the causality of features of neural codes in perception^{126,127,134}.

Two-photon patterned optogenetics presents the opportunity to probe how the readout of population activity to guide choices causally depends on correlations among neurons, by optogenetically manipulating the level of correlations between neurons without changing the response properties of individual neurons. A recent study trained mice to report the activation, in the absence of sensory stimulation, of groups of approximately 30 neurons in the olfactory bulb that were stimulated using two-photon patterned optogenetics¹³⁵. The optogenetic stimulation pattern was varied from trial to trial to alter the level of temporal synchronization between neurons while keeping the firing of individual cells similar across trials. The optogenetically induced activity had a larger effect on behavioural choices when the optogenetic stimulation was more synchronized, providing a causal demonstration that temporal

Box 3 | Probing the relationship between neural interactions and population computations

A key question of population coding is how the activity of one neuron affects the activity of other neurons to shape population-level information processing. This effect is problematic to measure with observations of neural activity because such observations cannot reveal causal relationships between neurons. Recent work has used two-photon optogenetics to trigger action potentials in a targeted neuron, plus calcium imaging to measure the effect on spiking in neighbouring neurons in layer 2/3 of the primary visual cortex (V1) of awake mice viewing visual stimuli¹³². Stimulation of a neuron had a predominantly suppressive influence on the activity of other V1 neurons and, on average, it suppressed more of the activity of neurons with similar tuning than that of neurons with different tuning. The exception to this finding was the presence of a positive influence between a small population of neurons with a strong positive signal and noise correlations. These results support earlier theories¹³⁸ that proposed a 'feature competition' computation, which could reduce noise and amplify signal in single-trial responses^{132,151} and demix information from feedforward inputs^{132,170}. In this computation, negative noise correlations between two neurons would mean that evidence provided by one neuron 'speaks against' that provided by the other neuron, and competition explains away the least likely cause of neuronal inputs.

The power of coupling computational models with neural perturbations and imaging has been further highlighted by other studies that have used neural ablations or two-photon optogenetics to reveal feature amplification motifs in various parts of the cortex. In these studies, photostimulation of only part of a neural population has led to the recruitment of neurons with similar tuning^{126,127,133}. Consistently, the ablation of neurons encoding particular stimuli reduced the activity of neurons with similar stimulus responses¹⁷¹. These findings are consistent with anatomical like-to-like excitatory connections between neurons with similar tuning^{172,173} and have been proposed as evidence for pattern completion and attractor dynamics in cortical circuits¹⁷⁴.

Future studies that couple two-photon optogenetics and imaging with population-coding analyses are needed to fully understand the implications of such results for population codes. For example, feature competition results are compatible with the view that recurrent connections in V1 may create information-enhancing correlations or reduce the effect of information-limiting correlations. Feature amplification results may lead to redundancy in the subnetwork of similarly tuned neurons if such amplification creates stimulus-independent correlations, but could create synergy if recurrent amplification creates stimulus-dependent correlations.

Two-photon patterned optogenetics

The use of light-sculpting, such as with a spatial light modulator, and two-photon excitation to create arbitrary spatial and temporal patterns of light to photostimulate neurons with approximately single-cell resolution.

Influence mapping

The process of measuring how spikes added by two-photon-patterned optogenetic perturbation to one or a few neurons causally affect the spiking of neighbouring neurons.

Multi-objective optimization

An optimization procedure that minimizes multiple cost functions simultaneously.

correlations among neurons enhance the behavioural readout of population activity (FIG. 4e,f).

Given that, during naturalistic sensory stimulation, many features of neural population activity can be characterized, including the strength and timing of neural firing and the correlations between neurons, a key question is how many different features come together to produce perception and whether correlations among these features are consequential to perception. To address this question, mice were trained to report a specific optogenetically induced spatial-temporal activity pattern in the olfactory bulb¹³⁶, again in the absence of sensory stimulation. Activity features were then manipulated independently or in combination by changing the optogenetic stimulus and the effects on the mouse's behavioural reports were observed. A conceptual advance of this study is that it used a common computational metric to assess the causal effects of individual features of the population code as well as how these features interact to shape the readout of information (FIG. 4g). This study found that temporal correlations between the activity of populations at different anatomical sites causally affected behavioural reports, again supporting the idea that correlations shape behavioural readout.

A limitation of these optogenetic studies is that they have not tested whether correlations serve the same functions during sensation, because optogenetic stimulation is applied in the absence of sensory stimuli. The fact that animals can detect the direct artificial activation of few tens of neurons in the absence of sensory stimuli does not imply that perception of naturalistic sensory stimuli relies on similarly small neural populations, and is not informative about the size of the neural population that underlies the perception of real sensory stimuli. Yet, understanding the size of populations needed to represent sensory stimuli is highly relevant, because the information-limiting effects of correlations increase with it.

These caveats, together with the decades of work studying the theoretical principles of correlated population codes, make it compelling to computationally design perturbations that can probe the causal role of correlations during the perception of real sensory stimuli. The ability to monitor neural activity with calcium imaging while simultaneously making optogenetic perturbations enables a direct assessment of the effect of the perturbation on the neural code. As one example of this opportunity, the hypothesis that information consistency enhances the readout of sensory information can be probed by designing optogenetic patterns that perturb responses to elicit different degrees of consistency across neurons or across time, but with comparable amounts of total sensory information in the population. A prediction is that perceptual discriminations will be more accurate in trials in which the optogenetic perturbations elicit higher levels of stimulus information consistency.

Outlook

The work reviewed above has begun to reveal how correlations and their structure shape the multiple functions that a population code performs. Despite this unequivocal progress, we consider there to be various avenues for continued advances in the near future.

Structure of population codes to balance multiple functions. Given that population codes support multiple functions, with possibly competing requirements on the correlation structure^{60,119}, a key goal is to develop a principled theoretical framework for how codes balance the constraints imposed by information encoding, the propagation of signals to downstream networks and the generation of appropriate timescales for computation. A major challenge is that the multi-objective optimization that a neural population has to perform is not known and may vary across the brain. For example, sensory cortices may benefit from weak correlations to maximize information encoding with short timescales for dynamic stimuli, whereas association cortices might optimize for stronger signal propagation and longer timescales, which arise from stronger correlations, to accumulate information and create a consistent code. Indeed, the strength of correlations and their effects on encoding, readout and timescales are smaller in sensory than in association areas⁸⁸, increase along the visual cortical hierarchy⁹² and are weaker in input than in output layers of the visual cortex¹³⁷.

One approach could be to consider different putative multi-objective optimizations that a population code may perform, derive theoretically optimal response properties under each optimization, and compare these properties with single-neuron tuning and noise correlation structures in different areas, similar to how previous work derived neural codes optimal for information encoding alone^{6–8,138,139}. To form empirical hypotheses about how multiple constraints are traded off across areas, the cost function optimized by each area could be inferred from the behavioural losses obtained when inactivating an area (that is, reducing its activity using, for instance, optogenetics). For example, population codes in areas whose inactivation leads to larger behavioural deficits related to the integration of information across time may weigh the benefits of generating long timescales more heavily¹⁴⁰.

Understanding how correlation structure relates to anatomical connections. Network theory has explained how differences in network structures lead to large differences in functions such as the speed and spread of information transfer across the networks^{5,141}. In neuroscience, these ideas have helped researchers to understand the dynamics and function of whole-brain networks measured using, for example, functional MRI¹⁴¹. Key to this progress has been the interpretation of functional data in relation to computational models informed by the underlying anatomical connections¹⁴². At the level of population codes, theoretical and experimental work has begun to elucidate the importance of the structure of signal and noise correlations and of synergistic and redundant pools in a network for information encoding and readout. However, our understanding of the importance of noise correlation structures has remained more limited, partly owing to the relative scarcity of data sets that include both the functional responses of many individual neurons and their anatomical wiring within a network. Ongoing improvements in electron microscopy and in X-ray holographic nano-tomography techniques

Retrograde labelling

Methods based on dyes or viruses that are taken up by axons and transported back to a neuron's cell body.

will make such data more common¹⁴³. We predict that relating the structure of noise correlations computed from simultaneous neural recordings and the structure of causal interactions computed with influence mapping to the underlying anatomical connectivity, especially with the help of computational models, will be key to further our understanding of population coding¹⁴⁴.

Noise correlations can reflect monosynaptic or polysynaptic anatomical connections between cells as well as reflecting or being modulated by common inputs, general fluctuations in behavioural state, network excitability or normalization mechanisms^{145–150}. The functional implications of noise correlations for encoding and readout discussed above are expected to hold regardless of their biophysical origin, because this shared variability in activity would act in the same way regardless of its origin. However, comparing the structure of noise correlations with the anatomical connectivity will help to establish the population-coding computations that anatomical motifs can support^{145,151}.

In addition, noise correlations have not often been investigated for the different cell types within neural circuits. Recent studies have found that correlations can differ between excitatory and inhibitory populations and across inhibitory interneuron types, including with changes over the course of learning^{152–154}. A challenge is therefore to test how the structure of noise correlations in a population code relates to the diversity of molecularly and functionally defined cell types.

Information flow of signals from sensation to action.

Given that the readout of information may not be optimal, it cannot be assumed that all information encoded in a population is used downstream to inform behaviour. It thus becomes crucial to establish tools and concepts to better understand the constraints and principles of information transmission by population codes.

A challenge is that the principles for neural population coding have been established mostly considering populations of neurons recorded in one location, without identifying subpopulations of neurons that send axonal projections to specific downstream targets. This raises the question of the extent to which these principles are relevant for between-area information transmission. This question is now, in principle, addressable using neural recording combined with retrograde labelling or other methods. One possibility is that the structure of population codes is different for populations of neurons that send information to the same downstream network compared with populations of neurons that send projections to distinct targets and might thus not be read out together. For example, the populations of neurons that send information to the same target network may have higher noise correlations to help more robust signal

propagation and/or may have a correlation structure that enhances information encoding. One intriguing possibility is that the routing and propagation of signals is structured in cortical networks^{155–158}, with specialized functional interactions in subpopulations of neurons that project to the same target that enhance information transmission and are helpful to generate accurate behaviour.

A potential approach to understanding which parts of the encoded information are used or transmitted relies on measuring how different features of the population code correlate with the behavioural outcome — for example, whether properties, such as the levels of noise correlations or information consistency, are higher in behaviourally correct trials⁶⁰. A second approach could be to estimate and compare the total amount of stimulus information encoded across areas in a processing hierarchy³⁸ (BOX 2). Experimental methods using electrophysiology and calcium imaging are rapidly emerging to address these questions to simultaneously record large populations of neurons in multiple, synaptically connected brain areas^{159–161}. Coupled with analytical methods and mathematical models for quantifying transmission of stimulus-specific information across neural populations in different cortical areas^{95,162–164}, there is an opportunity to directly measure and test, using perturbations, which features of neural codes are transmitted to downstream networks, and how information propagation relates to sensory perception and behavioural decisions⁹⁵.

Conclusion

Population coding is one of the fundamental levels at which information is processed in the brain. Recent work has made substantial progress in revealing that correlations in neural populations serve multiple potential functions, including encoding information, propagating signals to downstream networks to guide behaviour and setting diverse timescales of computation. Emerging work is discovering how the structure of correlations affects these functions. We propose that a major future direction is to understand how population codes trade off potentially competing constraints imposed by these multiple functions, and how these trade-offs vary for codes across different parts of the nervous system and different cell types. New avenues of discovery will arise from emerging technologies for the recording of large-scale neural activity, measurement of anatomical connectivity, perturbation of specific features of population activity, modelling and analysis of high-dimensional neural populations and, especially, from the combination of these approaches.

Published online: 22 June 2022

1. Wilson, M. A. & McNaughton, B. L. Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055–1058 (1993).
2. Rigotti, M. et al. The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).
3. Bernardi, S. et al. The geometry of abstraction in the hippocampus and prefrontal cortex. *Cell* **183**, 954–967.e21 (2020).
4. Koren, V., Andrei, A. R., Hu, M., Dragoi, V. & Obermayer, K. Pairwise synchrony and correlations depend on the structure of the population code in visual cortex. *Cell Rep.* <https://doi.org/10.1016/j.celrep.2020.108367> (2020).
5. Battiston, F. et al. The physics of higher-order interactions in complex systems. *Nat. Phys.* **17**, 1093–1098 (2021).
6. Barlow, H. B. Possible principles underlying the transformation of sensory messages. *Sens. Commun.* **1**, 217–234 (1961).
7. Attneave, F. Some informational aspects of visual perception. *Psychol. Rev.* **61**, 183–193 (1954).
8. Olshausen, B. A. & Field, D. J. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**, 607–609 (1996).

9. Mlynarski, W. F. & Hermundstad, A. M. Efficient and adaptive sensory codes. *Nat. Neurosci.* **24**, 998–1009 (2021).
10. Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation through neural population dynamics. *Annu. Rev. Neurosci.* **43**, 249–275 (2020).
11. Panzeri, S., Schultz, S. R., Treves, A. & Rolls, E. T. Correlations and the encoding of information in the nervous system. *Proc. Biol. Sci.* **266**, 1001–1012 (1999).
12. Pola, G., Thiele, A., Hoffmann, K. P. & Panzeri, S. An exact method to quantify the information transmitted by different mechanisms of correlational coding. *Network* **14**, 35–60 (2003).
13. Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).
14. Abbott, L. F. & Dayan, P. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* **11**, 91–101 (1999).
15. Gawne, T. J. & Richmond, B. J. How independent are the messages carried by adjacent inferior temporal cortical neurons? *J. Neurosci.* **13**, 2758–2771 (1993).
16. Cohen, M. R. & Kohn, A. Measuring and interpreting neuronal correlations. *Nat. Neurosci.* **14**, 811–819 (2011).
17. Averbeck, B. B. & Lee, D. Effects of noise correlations on information encoding and decoding. *J. Neurophysiol.* **95**, 3633–3644 (2006).
18. Nogueira, R. et al. The effects of population tuning and trial-by-trial variability on information encoding and behavior. *J. Neurosci.* **40**, 1066–1083 (2020).
19. Zohary, E., Shadlen, M. N. & Newsome, W. T. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* **370**, 140–143 (1994).
20. Ecker, A. S., Berens, P., Tolias, A. S. & Bethge, M. The effect of noise correlations in populations of diversely tuned neurons. *J. Neurosci.* **31**, 14272–14283 (2011).
21. Shamir, M. & Sompolinsky, H. Nonlinear population codes. *Neural Comput.* **16**, 1105–1136 (2004).
22. Josic, K., Shea-Brown, E., Doiron, B. & de la Rocha, J. Stimulus-dependent correlations and population codes. *Neural Comput.* **21**, 2774–2804 (2009).
23. Azeredo da Silveira, R. & Rieke, F. The geometry of information coding in correlated neural populations. *Annu. Rev. Neurosci.* **44**, 403–424 (2021).
24. Gray, C. M., König, P., Engel, A. K. & Singer, W. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* **338**, 334–337 (1989).
25. deCharms, R. C. & Merzenich, M. M. Primary cortical representation of sounds by the coordination of action-potential timing. *Nature* **381**, 610–613 (1996).
26. Franke, F. et al. Structures of neural correlation and how they favor coding. *Neuron* **89**, 409–422 (2016).
27. Dan, Y., Alonso, J. M., Usrey, W. M. & Reid, R. C. Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nat. Neurosci.* **1**, 501–507 (1998).
28. Zylberberg, J., Cafaro, J., Turner, M. H., Shea-Brown, E. & Rieke, F. Direction-selective circuits shape noise to ensure a precise population code. *Neuron* **89**, 369–383 (2016).
29. Kohn, A. & Smith, M. A. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* **25**, 3661–3673 (2005).
30. Romo, R., Hernandez, A., Zainos, A. & Salinas, E. Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. *Neuron* **38**, 649–657 (2003).
31. Reich, D. S., Mechler, F. & Victor, J. D. Independent and redundant information in nearby cortical neurons. *Science* **294**, 2566–2568 (2001).
32. Pillow, J. W. et al. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
33. Graf, A. B., Kohn, A., Jazayeri, M. & Movshon, J. A. Decoding the activity of neuronal populations in macaque primary visual cortex. *Nat. Neurosci.* **14**, 239–245 (2011).
34. Rupasinghe, A. et al. Direct extraction of signal and noise correlations from two-photon calcium imaging of ensemble neuronal activity. *eLife* <https://doi.org/10.7554/eLife.68046> (2021).
35. Rothschild, G., Nelken, I. & Mizrahi, A. Functional organization and population dynamics in the mouse primary auditory cortex. *Nat. Neurosci.* **13**, 353–360 (2010).
36. Kwon, S. E., Tsytarev, V., Erzurumlu, R. S. & O'Connor, D. H. Organization of orientation-specific whisker deflection responses in layer 2/3 of mouse somatosensory cortex. *Neuroscience* **368**, 46–56 (2018).
37. Bartolo, R., Saunders, R. C., Mitz, A. R. & Averbeck, B. B. Information-limiting correlations in large neural populations. *J. Neurosci.* **40**, 1668–1678 (2020).
38. Kafashan, M. et al. Scaling of sensory information in large neural populations shows signatures of information-limiting correlations. *Nat. Commun.* **12**, 473 (2021).
39. Petersen, R. S., Panzeri, S. & Diamond, M. E. Population coding of stimulus location in rat somatosensory cortex. *Neuron* **32**, 503–514 (2001).
40. 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).
41. Adibi, M., McDonald, J. S., Clifford, C. W. & Arabzadeh, E. Adaptation improves neural coding efficiency despite increasing correlations in variability. *J. Neurosci.* **33**, 2108–2120 (2013).
42. Rumyantsev, O. I. et al. Fundamental bounds on the fidelity of sensory cortical coding. *Nature* **580**, 100–105 (2020).
43. Sanayei, M. et al. Perceptual learning of fine contrast discrimination changes neuronal tuning and population coding in macaque V4. *Nat. Commun.* **9**, 4238 (2018).
44. Tremblay, S., Pieper, F., Sachs, A. & Martinez-Trujillo, J. Attentional filtering of visual information by neuronal ensembles in the primate lateral prefrontal cortex. *Neuron* **85**, 202–215 (2015).
45. Moreno-Bote, R. et al. Information-limiting correlations. *Nat. Neurosci.* **17**, 1410–1417 (2014).
46. Kanitscheider, I., Coen-Cagli, R. & Pouget, A. Origin of information-limiting noise correlations. *Proc. Natl Acad. Sci. USA* **112**, E6973–E6982 (2015).
47. Ecker, A. S. et al. Decorrelated neuronal firing in cortical microcircuits. *Science* **327**, 584–587 (2010).
48. Okun, M. et al. Diverse coupling of neurons to populations in sensory cortex. *Nature* **521**, 511–515 (2015).
49. Mincses, V., Pinto, L., Dan, Y. & Chiba, A. A. Cholinergic shaping of neural correlations. *Proc. Natl Acad. Sci. USA* **114**, 5725–5730 (2017).
50. Shamir, M. & Sompolinsky, H. Implications of neuronal diversity on population coding. *Neural Comput.* **18**, 1951–1986 (2006).
51. Wilke, S. D. & Eurich, C. W. Representational accuracy of stochastic neural populations. *Neural Comput.* **14**, 155–189 (2002).
52. Cohen, M. R. & Maunsell, J. H. Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* **12**, 1594–1600 (2009).
53. Gutnitsky, D. A. & Dragoi, V. Adaptive coding of visual information in neural populations. *Nature* **452**, 220–224 (2008).
54. Jeanne, J. M., Sharpee, T. O. & Gentner, T. Q. Associative learning enhances population coding by inverting interneuronal correlation patterns. *Neuron* **78**, 352–363 (2013).
55. Ruff, D. A. & Cohen, M. R. Attention can either increase or decrease spike count correlations in visual cortex. *Nat. Neurosci.* **17**, 1591–1597 (2014).
56. Downer, J. D., Niwa, M. & Sutter, M. L. Task engagement selectively modulates neural correlations in primary auditory cortex. *J. Neurosci.* **35**, 7565–7574 (2015).
57. Gu, Y. et al. Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* **71**, 750–761 (2011).
58. Nigam, S., Pojoga, S. & Dragoi, V. Synergistic coding of visual information in columnar networks. *Neuron* **104**, 402–411.e4 (2019).
59. 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).
60. Valente, A. et al. Correlations enhance the behavioral readout of neural population activity in association cortex. *Nat. Neurosci.* **24**, 975–986 (2021).
61. Umakantha, A. et al. Bridging neuronal correlations and dimensionality reduction. *Neuron* **109**, 2740–2754.e12 (2021).
62. Ganmor, E., Segev, R. & Schneidman, E. Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proc. Natl Acad. Sci. USA* **108**, 9679–9684 (2011).
63. Granot-Atedgi, E., Tkacik, G., Segev, R. & Schneidman, E. Stimulus-dependent maximum entropy models of neural population codes. *PLoS Comput. Biol.* **9**, e1002922 (2013).
64. Ohiorhuan, I. E. et al. Sparse coding and high-order correlations in fine-scale cortical networks. *Nature* **466**, 617–621 (2010).
65. Montani, F. et al. The impact of high-order interactions on the rate of synchronous discharge and information transmission in somatosensory cortex. *Philos. Trans. R. Soc. A Phys. Eng. Sci.* **367**, 3297–3310 (2009).
66. Giusti, C., Pastalkova, E., Curto, C. & Itskov, V. Clique topology reveals intrinsic geometric structure in neural correlations. *Proc. Natl Acad. Sci. USA* **112**, 13455–13460 (2015).
67. Froudarakis, E. et al. Population code in mouse V1 facilitates readout of natural scenes through increased sparseness. *Nat. Neurosci.* **17**, 851–857 (2014).
68. Chelaru, M. I. et al. High-order interactions explain the collective behavior of cortical populations in executive but not sensory areas. *Neuron* <https://doi.org/10.1016/j.neuron.2021.09.042> (2021).
69. Yu, S. et al. Higher-order interactions characterized in cortical activity. *J. Neurosci.* **31**, 17514–17526 (2011).
70. Gardner, R. J. et al. Toroidal topology of population activity in grid cells. *Nature* **602**, 123–128 (2022).
71. Cayco-Gajic, N. A., Zylberberg, J. & Shea-Brown, E. Triplet correlations among similarly tuned cells impact population coding. *Front. Comput. Neurosci.* <https://doi.org/10.3389/fncom.2015.00057> (2015).
72. Panzeri, S., Brunel, N., Logothetis, N. K. & Kayser, C. Sensory neural codes using multiplexed temporal scales. *Trends Neurosci.* **33**, 111–120 (2010).
73. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
74. Andersen, R. A. & Cui, H. Intention, action planning, and decision making in parietal-frontal circuits. *Neuron* **63**, 568–583 (2009).
75. Bisley, J. W. & Goldberg, M. E. Attention, intention, and priority in the parietal lobe. *Annu. Rev. Neurosci.* **33**, 1–21 (2010).
76. Curtis, C. E. & Lee, D. Beyond working memory: the role of persistent activity in decision making. *Trends Cogn. Sci.* **14**, 216–222 (2010).
77. Shadlen, M. N. & Newsome, W. T. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* **86**, 1916–1936 (2001).
78. Wang, X. J. Decision making in recurrent neural circuits. *Neuron* **60**, 215–234 (2008).
79. Inagaki, H. K., Fontolan, L., Romani, S. & Svoboda, K. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* **566**, 212–217 (2019).
80. Li, N., Daie, K., Svoboda, K. & Druckmann, S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* **532**, 459–464 (2016).
81. Li, N., Chen, T. W., Guo, Z. V., Gerfen, C. R. & Svoboda, K. A motor cortex circuit for motor planning and movement. *Nature* **519**, 51–56 (2015).
82. Hanks, T. D. et al. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* **520**, 220–223 (2015).
83. Crowe, D. A., Averbeck, B. B. & Chafee, M. V. Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. *J. Neurosci.* **30**, 11640–11653 (2010).
84. Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsáki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322–1327 (2008).
85. Baeg, E. H. et al. Dynamics of population code for working memory in the prefrontal cortex. *Neuron* **40**, 177–188 (2003).
86. Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
87. Morcos, A. S. & Harvey, C. D. History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat. Neurosci.* **19**, 1672–1681 (2016).
88. Runyan, C. A., Piasini, E., Panzeri, S. & Harvey, C. D. Distinct timescales of population coding across cortex. *Nature* **548**, 92–96 (2017).
89. Akrami, A., Kopec, C. D., Diamond, M. E. & Brody, C. D. Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* **554**, 368–372 (2018).

90. Zariwala, H. A., Kepecs, A., Uchida, N., Hirokawa, J. & Mainen, Z. F. The limits of deliberation in a perceptual decision task. *Neuron* **78**, 339–351 (2013).
91. Mazurek, M. E. & Shadlen, M. N. Limits to the temporal fidelity of cortical spike rate signals. *Nat. Neurosci.* **5**, 463–471 (2002).
92. Piasini, E. et al. Temporal stability of stimulus representation increases along rodent visual cortical hierarchies. *Nat. Commun.* **12**, 4448 (2021).
93. Gold, J. I. & Shadlen, M. N. Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* **5**, 10–16 (2001).
94. Beck, J. M., Wei, J. M., Pitkow, X., Peter, E. L. & Pouget, A. Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* **74**, 30–39 (2012).
95. Panzeri, S., Harvey, C. D., Piasini, E., Latham, P. E. & Fellin, T. Cracking the neural code for sensory perception by combining statistics, intervention, and behavior. *Neuron* **93**, 491–507 (2017).
96. Wu, S., Nakahara, H. & Amari, S. Population coding with correlation and an unfaithful model. *Neural Comput.* **13**, 775–797 (2001).
97. Latham, P. E. & Nirenberg, S. Synergy, redundancy, and independence in population codes, revisited. *J. Neurosci.* **25**, 5195–5206 (2005).
98. Nirenberg, S., Carcieri, S. M., Jacobs, A. L. & Latham, P. E. Retinal ganglion cells act largely as independent encoders. *Nature* **411**, 698–701 (2001).
99. Karpas, E. M., Kiani, R. O. & Schneidman, E. Strongly correlated spatiotemporal encoding and simple decoding in the primate cortex. *bioRxiv* <https://doi.org/10.1101/693192> (2019).
100. Salinas, E. & Sejnowski, T. J. Correlated neuronal activity and the flow of neural information. *Nat. Rev. Neurosci.* **2**, 539–550 (2001).
101. Stringer, C., Michaelos, M., Tsyboulski, D., Lindo, S. E. & Pachitariu, M. High-precision coding in visual cortex. *Cell* **184**, 2767 (2021).
102. Zhao, Y., Yates, J. L., Levi, A. J., Huk, A. C. & Park, I. M. Stimulus-choice (mis)alignment in primate area MT. *PLoS Comput. Biol.* **16**, e1007614 (2020).
103. Pica, G. P., et al. In *Advances in Neural Information Processing Systems (NeurIPS)*. (ed. Luxburg G. I., et al.) 3686–3696 (Curran Associates, Inc., 2020).
104. Koch, C., Rapp, M. & Segev, I. A brief history of time (constants). *Cereb. Cortex* **6**, 93–101 (1996).
105. Diesmann, M., Gewaltig, M. O. & Aertsen, A. Stable propagation of synchronous spiking in cortical neural networks. *Nature* **402**, 529–533 (1999).
106. Zandvakili, A. & Kohn, A. Coordinated neuronal activity enhances corticocortical communication. *Neuron* **87**, 827–839 (2015).
107. Alonso, J. M., Usrey, W. M. & Reid, R. C. Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* **383**, 815–819 (1996).
108. Softky, W. R. & Koch, C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* **13**, 334–350 (1993).
109. Reyes, A. D. Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro. *Nat. Neurosci.* **6**, 593–599 (2003).
110. Losonczy, A. & Magee, J. C. Integrative properties of radial oblique dendrites in hippocampal CA1 pyramidal neurons. *Neuron* **50**, 291–307 (2006).
111. Ariav, G., Polsky, A. & Schiller, J. Submillisecond precision of the input-output transformation function mediated by fast sodium dendritic spikes in basal dendrites of CA1 pyramidal neurons. *J. Neurosci.* **23**, 7750–7758 (2003).
112. London, M. & Häusser, M. Dendritic computation. *Annu. Rev. Neurosci.* **28**, 503–532 (2005).
113. Polsky, A., Mel, B. W. & Schiller, J. Computational subunits in thin dendrites of pyramidal cells. *Nat. Neurosci.* **7**, 621–627 (2004).
114. Golding, N. L., Staff, N. P. & Spruston, N. Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature* **418**, 326–331 (2002).
115. Smith, S. L., Smith, I. T., Branco, T. & Häusser, M. Dendritic spikes enhance stimulus selectivity in cortical neurons in vivo. *Nature* **503**, 115–120 (2013).
116. Schmidt-Hieber, C. et al. Active dendritic integration as a mechanism for robust and precise grid cell firing. *Nat. Neurosci.* **20**, 1114–1121 (2017).
117. Wilson, D. E., Whitney, D. E., Scholl, B. & Fitzpatrick, D. Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. *Nat. Neurosci.* **19**, 1003–1009 (2016).
118. Ackels, T. et al. Fast odour dynamics are encoded in the olfactory system and guide behaviour. *Nature* **593**, 558–563 (2021).
119. Zylberberg, J., Pouget, A., Latham, P. E. & Shea-Brown, E. Robust information propagation through noisy neural circuits. *PLoS Comput. Biol.* **13**, e1005497 (2017).
120. Histed, M. H. & Maunsell, J. H. Cortical neural populations can guide behavior by integrating inputs linearly, independent of synchrony. *Proc. Natl Acad. Sci. USA* **111**, E178–E187 (2014).
121. Shahidi, N., Andrei, A. R., Hu, M. & Dragoi, V. High-order coordination of cortical spiking activity modulates perceptual accuracy. *Nat. Neurosci.* **22**, 1148–1158 (2019).
122. Balaguer-Ballester, E., Nogueira, R., Abofalia, J. M., Moreno-Bote, R. & Sanchez-Vives, M. V. Representation of foreseeable choice outcomes in orbitofrontal cortex triplet-wise interactions. *PLoS Comput. Biol.* <https://doi.org/10.1371/journal.pcbi.1007862> (2020).
123. Zylberberg, J. & Shea-Brown, E. Input nonlinearities can shape beyond-pairwise correlations and improve information transmission by neural populations. *Phys. Rev. E* <https://doi.org/10.1103/PhysRevE.92.062707> (2015).
124. Emiliani, V., Cohen, A. E., Deisseroth, K. & Häusser, M. All-optical interrogation of neural circuits. *J. Neurosci.* **35**, 13917–13926 (2015).
125. Rickgauer, J. P., Deisseroth, K. & Tank, D. W. Simultaneous cellular-resolution optical perturbation and imaging of place cell firing fields. *Nat. Neurosci.* **17**, 1816–1824 (2014).
126. Marshel, J. H. et al. Cortical layer-specific critical dynamics triggering perception. *Science* <https://doi.org/10.1126/science.aaw5202> (2019).
127. Carrillo-Reid, L., Han, S., Yang, W., Akrouh, A. & Yuste, R. Controlling visually guided behavior by holographic recalling of cortical ensembles. *Cell* **178**, 447–457.e5 (2019).
128. Dalgleish, H. W. et al. How many neurons are sufficient for perception of cortical activity? *eLife* <https://doi.org/10.7554/eLife.58889> (2020).
129. Packer, A. M., Russell, L. E., Dalgleish, H. W. & Häusser, M. Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo. *Nat. Methods* **12**, 140–146 (2015).
130. Pegard, N. C. et al. Three-dimensional scanless holographic optogenetics with temporal focusing (3D-SHOT). *Nat. Commun.* **8**, 1228 (2017).
131. Forli, A. et al. Two-photon bidirectional control and imaging of neuronal excitability with high spatial resolution in vivo. *Cell Rep.* **22**, 3087–3098 (2018).
132. Chettih, S. N. & Harvey, C. D. Single-neuron perturbations reveal feature-specific competition in V1. *Nature* **567**, 334–340 (2019).
133. Daie, K., Svoboda, K. & Druckmann, S. Targeted photostimulation uncovers circuit motifs supporting short-term memory. *Nat. Neurosci.* **24**, 259–265 (2021).
134. Robinson, N. T. M. et al. Targeted activation of hippocampal place cells drives memory-guided spatial behavior. *Cell* **183**, 2041–2042 (2020).
135. Gill, J. V. et al. Precise holographic manipulation of olfactory circuits reveals coding features determining perceptual detection. *Neuron* **108**, 382–393.e5 (2020).
136. Chong, E. et al. Manipulating synthetic optogenetic odors reveals the coding logic of olfactory perception. *Science* <https://doi.org/10.1126/science.aba2357> (2020).
137. Hansen, B. J., Chelaru, M. I. & Dragoi, V. Correlated variability in laminar cortical circuits. *Neuron* **76**, 590–602 (2012).
138. Olshausen, B. A. & Field, D. J. Sparse coding of sensory inputs. *Curr. Opin. Neurobiol.* **14**, 481–487 (2004).
139. Simoncelli, E. P. & Olshausen, B. A. Natural image statistics and neural representation. *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
140. Pinto, L., Tank, D. W. & Brody, C. D. Multiple timescales of sensory-evidence accumulation across the dorsal cortex. *bioRxiv* <https://doi.org/10.1101/2020.12.28.424600> (2021).
141. Lynn, C. W. & Basset, D. S. The physics of brain network structure, function and control. *Nat. Rev. Phys.* **1**, 318–332 (2019).
142. Sporns, O., Tononi, G. & Kötter, R. The human connectome: a structural description of the human brain. *PLoS Comput. Biol.* **1**, e42 (2005).
143. Kuan, A. T. et al. Dense neuronal reconstruction through X-ray holographic nano-tomography. *Nat. Neurosci.* **23**, 1637–1643 (2020).
144. Ocker, G. K. et al. From the statistics of connectivity to the statistics of spike times in neuronal networks. *Curr. Opin. Neurobiol.* **46**, 109–119 (2017).
145. Rosenbaum, R., Smith, M. A., Kohn, A., Rubin, J. E. & Doiron, B. The spatial structure of correlated neuronal variability. *Nat. Neurosci.* **20**, 107–114 (2017).
146. Goris, R. L., Movshon, J. A. & Simoncelli, E. P. Partitioning neuronal variability. *Nat. Neurosci.* **17**, 858–865 (2014).
147. Kohn, A., Coen-Cagli, R., Kanitscheider, I. & Pouget, A. Correlations and neuronal population information. *Annu. Rev. Neurosci.* **39**, 237–256 (2016).
148. Ostojic, S., Brunel, N. & Hakim, V. How connectivity, background activity, and synaptic properties shape the cross-correlation between spike trains. *J. Neurosci.* **29**, 10234–10253 (2009).
149. Stringer, C. et al. Spontaneous behaviors drive multidimensional, brainwide activity. *Nature* **364**, aav7893 (2019).
150. Verhoef, B. E. & Maunsell, J. H. R. Attention-related changes in correlated neuronal activity arise from normalization mechanisms. *Nat. Neurosci.* **20**, 969–977 (2017).
151. Sadeh, S. & Clopath, C. Theory of neuronal perturbome in cortical networks. *Proc. Natl Acad. Sci. USA* **117**, 26966–26976 (2020).
152. Khan, A. G. et al. Distinct learning-induced changes in stimulus selectivity and interactions of GABAergic interneuron classes in visual cortex. *Nat. Neurosci.* **21**, 851–859 (2018).
153. Poort, J. et al. Learning enhances sensory and multiple non-sensory representations in primary visual cortex. *Neuron* **86**, 1478–1490 (2015).
154. Bittner, S. R. et al. Population activity structure of excitatory and inhibitory neurons. *PLoS One* **12**, e0181773 (2017).
155. Economo, M. N. et al. Distinct descending motor cortex pathways and their roles in movement. *Nature* **563**, 79–84 (2018).
156. Chen, J. L., Carta, S., Soldado-Magraner, J., Schneider, B. L. & Helmchen, F. Behaviour-dependent recruitment of long-range projection neurons in somatosensory cortex. *Nature* **499**, 336–340 (2013).
157. Huda, R. et al. Distinct prefrontal top-down circuits differentially modulate sensorimotor behavior. *Nat. Commun.* **11**, 6007 (2020).
158. Itokazu, T. et al. Streamlined sensory motor communication through cortical reciprocal connectivity in a visually guided eye movement task. *Nat. Commun.* **9**, 338 (2018).
159. Sofroniew, N. J., Flickinger, D., King, J. & Svoboda, K. A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *eLife* <https://doi.org/10.7554/eLife.14472> (2016).
160. Jun, J. J. et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).
161. Urai, A. E., Doiron, B., Leifer, A. M. & Churchland, A. K. Large-scale neural recordings call for new insights to link brain and behavior. *Nat. Neurosci.* **25**, 11–19 (2022).
162. Kohn, A. et al. Principles of corticocortical communication: proposed schemes and design considerations. *Trends Neurosci.* **43**, 725–737 (2020).
163. Semedo, J. D., Gokcen, E., Machens, C. K., Kohn, A. & Yu, B. M. Statistical methods for dissecting interactions between brain areas. *Curr. Opin. Neurobiol.* **65**, 59–69 (2020).
164. 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).
165. Schneidman, E., Berry, M. J. II, Segev, R. & Bialek, W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* **440**, 1007–1012 (2006).
166. Onken, A., Grunewald, S., Munk, M. H. & Obermayer, K. Analyzing short-term noise dependencies of spike-counts in macaque prefrontal cortex using copulas and the flashlight transformation. *PLoS Comput. Biol.* **5**, e1000577 (2009).
167. Berkes, P., Wood, F. & Pillow, J. Characterizing neural dependencies with copula models. *Adv. Neural Inf. Process. Syst.* **21**, 129–136 (2009).
168. Safaai, H., Onken, A., Harvey, C. D. & Panzeri, S. Information estimation using nonparametric copulas.

- Phys. Rev. E* <https://doi.org/10.1103/PhysRevE.98.053302> (2018).
169. Nienborg, H. & Cumming, B. G. Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* **459**, 89–92 (2009).
170. Moreno-Bote, R. & Drugowitsch, J. Causal inference and explaining away in a spiking network. *Sci. Rep.* **5**, 17531 (2015).
171. Peron, S. et al. Recurrent interactions in local cortical circuits. *Nature* **579**, 256–259 (2020).
172. Cossell, L. et al. Functional organization of excitatory synaptic strength in primary visual cortex. *Nature* **518**, 399–403 (2015).
173. Lee, W. C. et al. Anatomy and function of an excitatory network in the visual cortex. *Nature* **532**, 370–374 (2016).

174. Carrillo-Reid, L. & Yuste, R. Playing the piano with the cortex: role of neuronal ensembles and pattern completion in perception and behavior. *Curr. Opin. Neurobiol.* **64**, 89–95 (2020).

Acknowledgements

The authors thank members of their laboratories for helpful discussions, and B. Babadi, C. Becchio, G. Bondanelli, J. Drugowitsch, M. Histed, G. Iurilli, S. Lemke, J. Maunsell and E. Piasini for feedback. This work was supported by the US National Institutes of Health (NIH) grants DP1 MH125776 (C.D.H.); the US National Institute of Neurological Disorders and Stroke (NINDS) R01 NS089521 (C.D.H.); the Brain Research Through Advancing Innovative Neurotechnologies (BRAIN) Initiative R01 NS108410 (C.D.H. and S.P.), U19 NS107464 (S.P.), R01 NS109961 (S.P.); and the Fondation Bertarelli (S.P.).

Author contributions

The authors contributed to all aspects of the article.

Competing interests

The authors declare no competing interests.

Peer review information

Nature Reviews Neuroscience thanks V. Dragoi and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Publisher's note

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

© Springer Nature Limited 2022