REVIEW SUMMARY

Analysis methods for large-scale neuronal recordings

Carsen Stringer* and Marius Pachitariu

BACKGROUND: It is now easier than ever to record the activity of many neurons in awake, behaving animals. These recordings can be obtained either with multichannel electrodes or with two-photon microscopy, each with their own advantages. The latter breaks records for the number of simultaneously recorded neurons, which can reach into the hundreds of thousands. The former allows for unmatched temporal resolution and easy access to deep brain areas. Spoiled with this embarrassment of riches, neuroscientists must now demonstrate that they can use such data to answer longstanding questions and formulate new ones in the effort to understand the basic principles of brain function. The primary obstacle appears to be conceptual: What can we even infer from large-scale recordings that we could not have learned from smaller ones?

ADVANCES: Some of the recent advances in this field have come by formulating new analysis methods that probe specific properties of

neural computations. For example, a longstanding question has been whether neural circuits use a dedicated set of neural patterns or "neural code" to represent information regardless of whether that activity is externally or internally driven. Specifically, one may ask whether the neural activity patterns observed when an animal explores one environment are similar, up to a temporal permutation, to the activity patterns during rest or during sleep. The answer turned out to be a resounding yes for grid cells in the rat entorhinal cortex, and a clear no for a similar question in primary visual cortex. The activity patterns evoked by stimuli in visual cortex were almost as different as possible from the patterns during spontaneous activity, and the latter were strongly related to orofacial behaviors instead. In further studies, movements in general were found to influence activity across the entire mouse brain, even during the execution of sensory-guided motor tasks. Multiple studies have now given stark



B A "rastermap" of neural activity coordinated across spatial scales



An abundance of data available for analysis. (**A**) Two-photon calcium imaging from multiple visual cortical areas in a behaving mouse. (**B**) Raster plot of neural activity from the recording in (A), showing coordinated activity across timescales and spatial scales.

warnings to neuroscientists who thought Check for

were studying decision-making, working normal ory, and other internal processes: What you study might be direct reflections of movement instead—for example, through motor efference or proprioception.

All of these statements required the development of bespoke mathematical models that were carefully applied to large-scale neural data together with a well-thought-out set of statistical controls. The most successful applications of models to neural data can be described intuitively in words such as we did above, even when the underlying mathematical techniques may be quite complex.

In this Review, we synthesize some of the most promising approaches and try to offer insights into what motivates current studies and how the methods developed may be further used in the future. We emphasize the need to answer concrete scientific questions and the importance of performing appropriate controls and cross-validation to avoid some the most common pitfalls when analyzing big data.

OUTLOOK: Newly uncovered scaling laws over the number of recorded neurons demonstrate that in many tasks, such as neural decoding or encoding, the amount of information or unique variance continues to increase as the number of neurons increases even beyond 10,000. Neural activity patterns are thus rich with information, although it remains to be seen whether that information is rich with meaning. To find its meaning, neuroscientists will need to combine largescale recordings with other methods and paradigms such as complex behavior, two-photon optogenetics, or anatomical methods. For the example of grid cells above, the persistence of the neural code suggests a dynamical constraint imposed by recurrent circuitry; connectomic reconstruction could confirm this in the coming years.

Perhaps equally important, experimental neuroscientists will need to quickly become experts in computational and statistical analysis methods so that they can fully take advantage of the overwhelming amount of new data. Innovative online summer courses such as the Neuromatch Academy can get many neuroscientists up to speed quickly and even familiarize them with the rapid advances in artificial intelligence with the new "NeuroAI" course launched this summer.

READ THE FULL ARTICLE AT

https://doi.org/10.1126/science.adp7429

Howard Hughes Medical Institute (HHMI) Janelia Research Campus, Ashburn, VA, USA,

^{*}Corresponding author. Email: stringerc@janelia.hhmi.org Cite this article as C. Stringer, M. Pachitariu, *Science* **386**, eadp7429 (2024). DOI: 10.1126/science.adp7429

REVIEW

NEUROSCIENCE Analysis methods for large-scale neuronal recordings

Carsen Stringer* and Marius Pachitariu

Simultaneous recordings from hundreds or thousands of neurons are becoming routine because of innovations in instrumentation, molecular tools, and data processing software. Such recordings can be analyzed with data science methods, but it is not immediately clear what methods to use or how to adapt them for neuroscience applications. We review, categorize, and illustrate diverse analysis methods for neural population recordings and describe how these methods have been used to make progress on longstanding questions in neuroscience. We review a variety of approaches, ranging from the mathematically simple to the complex, from exploratory to hypothesis-driven, and from recently developed to more established methods. We also illustrate some of the common statistical pitfalls in analyzing large-scale neural data.

he past decade has brought a large array of new tools for investigating the brain. Among these are tools that enable functional recordings of large numbers of neurons with electrophysiological or optical imaging methods (1-5). These tools have been needed to make progress on a variety of questions in neuroscience. For example, single neurons respond with different numbers of action potentials to repeated presentations of the same stimuli (6), but it has not been clear whether that variability is coordinated at a macroscopic level across entire brain areas. Similarly, neurons in sensory cortical areas fire action potentials in the absence of sensory inputs (7), but it has not been known how their so-called "spontaneous" activity is correlated at the level of entire neural populations or what that activity may represent. As a final example, individual grid cells in the entorhinal cortex represent periodic spatial features of the environment (8), but it has not been clear how the individual lattices align and coordinate across an entire population.

More generally, large-scale neural recordings can be useful for at least three reasons: (i) they substantially accelerate data collection; (ii) they enable the study of coordinated neural activity; and (iii) they enable the study of simultaneous computations happening across multiple spatial scales. We briefly review the technological progress that has enabled largescale recordings then devote the majority of this Review to analysis techniques for making sense of large-scale neural data.

A brief history of technological progress

On the side of electrophysiology, the development of Neuropixels constituted a major leap

*Corresponding author. Email: stringerc@janelia.hhmi.org

in the number of simultaneously recorded channels on a single electrode shank (3). This advance was possible because of engineering efforts toward miniaturization and electronic integration. Subsequent iterations of the probe have led to further miniaturization for use in freely moving animals and more channels over multiple shanks for use in flatter brain areas such as the cortex and hippocampus (4), as well as thicker and longer versions of the probe for use in nonhuman primates (5) and humans (9, 10). Other probes with many channels have also been developed separatelynotably, the flat arrays of thousands of electrodes used in retinal recordings (11), and multiple types of flexible electrodes that promise improved stability for chronic recordings (12, 13). To take advantage of these new devices, an overhaul of classical data processing methods was needed because existing methods required substantial human manual curation, which is not effective for data at this scale. This led to the development of automated spike-sorting methods such as Kilosort, MountainSort, JRClust, and others (14-16).

On the optical imaging side, progress has mainly relied on two-photon calcium imaging (17). Although first developed much earlier (18), two-photon calcium imaging became widespread with the introduction of the genetically encoded calcium indicator GCaMP6 (1). Other factors that havedriven the adoption of calcium imaging have been the availability of easy-touse commercial microscopes, innovations in high-power laser technology, and a shift in research interests toward smaller animal models for which molecular tools such as GCaMP are easier to develop and test. Very-large-scale recordings of tens of thousands to hundreds of thousands of neurons were demonstrated in both zebrafish and mouse, although these required sacrifices in temporal resolution (an example recording is provided in movie S1) (19-21). Specialized headmounted "miniscope"

devices have also been developed for use in freely moving animals, by using either onephoton (22) or more recently two-photon excitation (23). Recent iterations of the calcium sensors (24) more accurately represent action potentials, and the sensors can now be targeted toward specific neuronal compartments such as somas and axons (25, 26). Calcium imaging has also become effective in nonhuman primates (27, 28). Similar to the developments in electrophysiology, the large amount of data collected required the development of automated pipelines such as Suite2p, Caiman, and others for identifying regions of interest such as cells, extracting their time-varying activity traces and converting these traces to an estimate of spike times through use of deconvolution methods (20, 29, 30).

Over the past few years, these technological improvements have spread through many neuroscience laboratories. Large amounts of data are being generated, and the question of what to do with this data arises often. We review recent analysis methods for largescale neural recordings and the discoveries they generated. This Review will proceed as a series of questions, methods, and possible answers that someone might encounter when analyzing large-scale neural recordings. Perhaps you collected this data yourself, or a collaborator has collected it for you. Or perhaps you found the data shared on the internet, there being increasingly more open datasets online. Here is a possible plan for how you might approach this data: (i) To start, you might use scaled-up versions of classical methods to look for neurons with specific kinds of tuning properties and study the distribution of these properties under varying conditions. (ii) Then, you might look for coordinated structure in the activity of all the neurons recorded simultaneously, which would typically be done with dimensionality reduction methods. (iii) Last, you might attempt to connect the neural populations found in steps (i) and (ii) with taskrelevant variables that can either be explicitly defined (stimuli or behaviors) or implicitly defined through models (such as reward expectation or attentional modulation). We illustrate such analyses with toy examples throughout the figures.

Single-neuron properties at scale

The easiest way to start analyzing large-scale recordings is to use existing methods for analyzing single neurons. This can be particularly convenient as a starting point because it allows one to replicate previous results and ensure that the new recording method does not introduce artifacts or warp the statistics of the data in unknown ways. Conversely, it allows one to validate results found in previous studies with many fewer neurons and potentially to identify special cases such as particular

Howard Hughes Medical Institute (HHMI) Janelia Research Campus, Ashburn, VA, USA.

populations in which previous results do not hold. Taking advantage of the relatively quick pace of collecting neural recordings with modern devices, one might be able to experimentally acquire multiple datasets in the same animal from different brain areas, from different cell types, over multiple days of learning, or across various manipulations. This enables the conducting of survey-type studies, in which the distribution of single-neuron properties is measured and compared between neural populations or between conditions, something that would have been difficult with previous datasets of relatively few neurons (Fig. 1A). This approach has been effective in studying the response properties of sensory systems such as the visual, auditory, and somatosensory cortices (31-33) or the thalamic and midbrain sensory areas (34-37) and has also been used in decision-making tasks and navigation tasks (38-41). In this data regime, exploratory analyses can be performed with simple scatter plots between pairs of parameters estimated from single-neuron response properties (Fig. 1A). Many response properties can typically be measured-for example, selectivity indices (42), tuning curve shape (43), latency (44), and behavior correlations (45)-across trial variability or Fano factor (46). Thus, a large number of exploratory scatter plots can be generated and visualized, which may allow for unexpected relations to be discovered. This can be followed by investigations of how the relations change across brain areas or as a result of manipulations (Fig. 1A).

Population averaging

Having potentially found some tuning properties of interest, the next step may be to combine or "average" the neurons in some suitable way to clearly illustrate the tuning of the entire population. In a recent study, we used this approach to find populations of neurons tuned to some learned, familiar visual stimuli and tested whether that population responds to new stimuli according to their visual categories (47). Other applications include identifying a neural population that is tuned to the preparatory movement activity or to the execution of a motor behavior, and analyzing the dynamics of that population on test trials with global optogenetic perturbations (48, 49).

A simple example of population averaging is the "coding direction" analysis (Fig. 1B). To obtain a coding direction, similarly tuned neurons can be averaged, which creates a less noisy population average, and opposite-tuned population averages may be subtracted from this average (47, 48, 50-53). Multiple coding directions can be obtained by use of different variables. Decoding approaches can also be used to infer the weighting of neurons across the coding direction-for example, with linear regression, reduced rank regression, and other models (54-58). One possible caveat of such methods is that they rely on single-neuron properties, which may be noisy. To improve on this, one may start by denoising or "cosmoothing" the neurons by using either simple methods such as principal components analysis (PCA) or non-negative matrix factorization (NMF). or more complex methods that use temporal information such as Gaussian process factor analysis (GPFA) (59).

The structure of neural population activity

A more ambitious goal may be to use largescale recordings to identify structure and coordinated patterns in the neural populations (60-66). This structure may be generated intrinsically by the dynamics of neural circuits, such as in the precisely wired systems of the *Drosophila* ring attractor (67) or in the central



Fig. 1. Single-neuron analyses at scale. (**A**) Sequence of operations in a typical analysis of single-neuron properties. (Left) Response properties such as tuning curves can have multiple quantifiable properties—param1, param2, and so on. (Left middle) Scatter plots between these properties can help identify unexpected relations. (Middle) Scatter plots can be transformed into 2D density plots for quantification. (Middle right) Multiple density plots can be obtained from different populations or the same population under different manipulations. (Right) The difference in densities can reveal differences between neural populations or the effect of manipulations.

(**B**) Illustration of population averaging approach. Neurons in different brain areas are found that are selective to (left) some particular stimulus and (middle left) some particular behavior. (Middle right) Neurons with the same selectivity are grouped together and averaged to form multiple coding directions. (Right) The dynamics of the coding directions can be visualized as a function of stimuli and behavior, possibly across brain areas and conditions. (**C**) Illustration of population smoothing or denoising approach that increases the signal-to-noise ratio of single neurons to help in better identifying single-neuron properties.



Fig. 2. Population vectors and dimensionality reduction. (**A**) Geometrical approaches. (Left) Measuring angles between neural population vectors. (Middle left) Measuring the variances and eigenvalues of the principal components of a neural population. (Middle right) Comparing representations patterns between neural populations and models or between two different neural populations. (Right) Aligning representations from one neural population to another. (**B**) Topological methods can identify nonlinear underlying structure in neural population recordings, such as a torus

pattern generators characterized in other invertebrates (68). In mammalian brains, it is much less clear which patterns are generated intrinsically in a circuit and which are inherited from the inputs to that circuit (69, 70). As we shall see, useful clues can emerge by extracting and analyzing the structure of the neural population activity.

These analyses typically start by converting the data into a sequence of population vectors, with each of these vectors representing the pattern of neural firing across the entire population at a specific time point. This is commonly followed by some type of dimensionality reduction method to identify simpler, latent representations of the high-dimensional neural data. Although the representations can be studied in isolation, they can also be related back to properties of the task or of the behavior.

Characterizing population vectors

To start, one may ask whether the structure and geometry of the neural population vectors are maintained across various task conditions and in spontaneous activity or during sleep (Fig. 2A). If the geometry was maintained, that could be consistent with the hypothesis that the neural coordination emerges from intrinsic dynamics rather than inputs. For example, the population vectors in entorhinal grid cells, which represent space, maintain their structure between wakefulness and sleep and across multiple different environments, suggesting a role for intrinsic attractor dynamics (71). By contrast, in sensory brain regions we and others have found a near-orthogonal representation of sensory-evoked and spontaneous activity in both mice and fish (21, 72, 73).

entire population at once.

Another property of interest is the linear dimensionality of the possible subspace spanned by neural population vectors. If neural activity in a specific circuit was restricted to a lowdimensional subspace, across many stimuli or behaviors, that could be consistent with the hypothesis that neuronal network dynamics force the activity to that subspace. Although earlier studies had suggested that neural activity in many brain areas was restricted to a low-dimensional subspace, theoretical analyses showed that this may be have been due to the low number of recorded neurons and the small number of task conditions (74). Large-scale recordings found that both the stimulus-evoked and the movement-driven population activity in mouse visual cortex is high-dimensional and appears to obey a power-law decay of variance across principal components (21, 75). Similarly high-dimensional structure was found in the cerebellum and in monkey V1 (76, 77). By contrast, the linear dimensionality of entorhinal population activity appears to be six, regardless of task condition (77).

for grid cells. (C) Dynamical systems methods can identify dynamical

reduction for exploratory analyses. (Left) Schematic of the relation

motifs, such as fixed points and line attractors. (D) Using dimensionality

between neurons and components. (Middle left) Finding relations between

trajectories over time for the components of interest. (Right) A different

approach called Rastermap for visualizing neural components across the

components and stimuli or behaviors. (Middle right) Scatter plots to visualize

Another class of geometric methods focuses on pairwise comparisons between the geometries of different datasets such as, for example, two recordings in different animals, or in the same animal on two different days or across experimental conditions, or between a neuronal dataset and a model (78–82). When a direct comparison is not possible or desirable, one may instead rotate a subspace of the neural population activity from one dataset to match

A Encoding models



Fig. 3. Encoding, decoding, and cross-validation. (**A**) Two approaches for building encoding models of neural responses. (Left) Schematic of reduced rank regression. (Right) Schematic of encoding models with basis function representations of discrete temporal events. (**B**) Decoding-based approaches. (Left) Schematic of decoder. (Middle left) Comparison of decoder prediction and the true variables on test trials. (Middle) Scaling of decoding error with the number of recorded neurons or trials. (Middle right) Finding a relation between the decoding error and another experimental variable such as a behavior or a second neural population (Right) Using the decoder to identify neural events corresponding to memory replay or planning. (**C**) Even single-

the structure of the equivalent subspace in a different dataset (*83, 84*). Such subspace alignment approaches may be relevant for brainmachine interfaces or for studies of across-day learning that may result in "representation drift" in some brain areas such as the hippocampus (*85*).

Nonlinear dimensionality reduction

The intrinsic dynamics of neural circuits may result in low-dimensional population dynamics that span a nonlinear, curved manifold and therefore may not be found with linear methods. Nonlinear dimensionality reduction methods can uncover such a structure that is low-dimensional in a topological sense (86). Methods to identify such structures often start by constructing a graph of population vectors connected according to distance in Euclidean space (Fig. 2B) and use an optimization method such as Isomap, uniform manifold approximation and projection (UMAP), or *t*-distributed stochastic neighbor embedding (t-SNE) to map these graphs onto a low-dimensional embedding (71, 87, 88). Last, the embedding dimensions are often identified as relevant behavioral variables, such as head direction or two-dimensional (2D) position, and the topology of the neural activity is often

neuron analyses require careful statistical cross-validation. (Left) Apparent sequences of neural activity may only be a consequence of ordering on train trials if they do not persist on test trials. (Middle left) Tuning curves from real data are often much more noisy than the assumed ideal. (Middle and right) If the tuning curves are calculated on the training trials used to obtain the preferred stimulus, data with no tuning may appear tuned (example 1), data with tuning may have distorted tuning curve shapes (example 2), and experimental manipulations may appear to produce complex tuning curve changes (example 3). These biases are not present when tuning curves are computed on test trials.

compared with the topology of the task or environment (89, 90). In some cases, the analysis of perturbations around the identified manifold (87) or the persistence of the manifold across behavioral states (71) can further be indicative of attractor-like dynamics around fixed points. The next barrier in this field is in identifying topological invariants from more noisy and mixed representations such as those found in the cortex during active behavior.

Dimensionality reduction of neural dynamics

Yet another possible clue for intrinsic neuronal dynamics comes from considering the behavior



Fig. 4. Frameworks for analyzing large-scale recordings. An overview of analysis methods for large-scale recordings corresponding to the different sections of this Review.

of neurons across time. Neurons do not change their firing rates randomly from one moment to the next; instead, their dynamics can be predictive of future neural activity, and correlations exist between neurons not just instantaneously but also at different timelags, on multiple timescales. Could these spatiotemporal correlations be the result of an intrinsic dynamical system, and can we uncover the rules of the dynamics from the observed neural patterns? This question has motivated much research into comparing neural network simulations with neural data (91-94) or inferring directly from data temporal structure and complex dynamical systems (62-66, 95). However, fully reviewing that work is beyond the scope of this Review. Further progress may take inspiration from the relatively new field of neural ordinary differential equations (ODEs) (96), which has recently found applicability in neuroscience (97).

Relating neural structure to task variables

This Review has described methods that try to link intrinsically generated neural dynamics to neural population activity through various forms of dimensionality reduction and modeling. However, dimensionality reduction can also be used for exploratory analyses—for example, to help in identifying the effect of observed or latent variables such as attentional modulation or reward anticipation, which may not be immediately evident otherwise. Regardless of what specific dimensionality reduction method is used (98-104), the steps following its application are similar: The components are typically compared with stimuli, behaviors, or inferred latent variables to see which components may be of interest. Once the interesting components are found (not necessarily the "top" ones output by the algorithm), further exploratory analyses consist of plotting trajectories in pairwise component plots and annotating those trajectories with various information (105-108). Success is typically defined by the ability to identify scientifically meaningful patterns in such plots and is usually followed by quantifications such as those described in the next section.

One caveat of this approach is that the components are often not distinctly identifiable across animals or sessions. For example, component 2 in one animal may actually be component 3 in another, or may not even be found at all. This is not necessarily attributable to variations across animals but rather to the similarity of most dimensionality reduction methods: Similar model performance can be obtained when the components are rotated or otherwise mixed together, which prevents a single solution from dominating the cost function landscape (59, 109). Various approaches exist to improve identifiability, usually by placing constraints such as sparsity, orthogonality, maximum variance ordering, or eigenvalue decompositions. Another caveat of dimensionality reduction is that it requires checking potentially very many components to find ones with interesting structure and may also require pairwise scatter plots between these components. A more global approach for this exploratory step was recently proposed as Rastermap, which is a technique for embedding the main components or clusters of neural activity into a single plot that can be more easily visualized on a single-trial basis (*110*).

Encoding and decoding models

In the previous sections, we discussed several exploratory analyses and various unsupervised models that can be fit directly to neural data and potentially related back to stimuli and behaviors. For the rest of this Review, we discuss encoding and decoding models, which can more directly relate neural activity to observed and latent variables. For example, one may ask whether neurons across the mouse brain encode sensory, motor, and/or feedback information, and this evaluation can be performed on a per-region basis (40). Alternatively, one may ask how much information a sensory neuronal population such as V1 contains about the orientation of a drifting grating, as done in our study (56). These types of goals can be achieved with encoding and decoding models, respectively, which we discuss next, followed by a section on the statistical caveats that arise from doing such analyses at scale.

Encoding models

Suppose we wanted to show that a particular population of neurons strongly encodes some variables of interest, such as behaviors and stimuli. We could construct an appropriate mathematical function that is based on those

variables-an encoding model-fit it to some training data from the neural responses, and use the fitted model to predict responses on test trials (Fig. 3A). In recent work, we did just that to predict neural responses in superficial cortex on the basis of orofacial behaviors and found that we could predict ~55% of the explainable variance (not accounting for singleneuron variability) (111). This was only possible because we could fit encoding models to a large population of simultaneously recorded neurons, thus allowing us to fit bigger models, with more parameters and more predictive power. In this case, performance continued to improve until reaching around 10,000 simultaneously recorded neurons. It is also possible to use encoding models for linking neural data to latent variables such as decision variables, confidence, expectations, beliefs, behavioral states, policies, or internal time estimates (65, 80, 95, 103, 112). To do so, one must first construct a behavioral model to estimate the latent variables from observable ones then use these estimates as predictors (112).

Encoding models can also be used to disambiguate the effect of correlated variables on neural activity. For example, decision-making signals are often highly correlated to movement signals, which makes it difficult to disambiguate whether a neuron in the brain correlates with the former or the latter. An encoding model can separate the contributions of these signals, as long as the signals have sufficient variability over trials to provide enough statistical power. This was done in multiple studies to show that whereas movement signals were widely distributed across the brain, decision-making signals were much more spatially restricted (40, 113). Thus, one must be careful not to wrongly interpret these movement correlates as correlates of decisions (113-115) or as correlates of behavior-evoking stimuli such as sounds (116, 117). It is also not possible without perturbation studies to know whether these neural representations are causally driving movements or if they just represent motor efference and proprioception.

Encoding models can take a wide range of mathematical forms. For large-scale neural activity, a particularly effective class of encoding models can be built by using reduced rank regression or canonical correlation analysis, which can capture the relation between a large number of inputs and outputs by using a small number of intermediate linear or nonlinear combinations (21, 118, 119). A different class of encoding models can capture the relations between neural data and discrete temporal events such as sensory cues or reward times by using a set of predefined basis functions. When appropriately weighted, these basis functions define the temporally extended nature of neural responses to discrete events (39, 40, 120-123). Encoding models can also be more powerful nonlinear models such as convolutional neural networks (*111, 124–126*) or even transformers, which are typically applied in a self-supervised manner (*127, 128*).

Decoding methods

Another way to show that a neural population encodes an external variable is to decode that variable from the population activity on test trials (Fig. 3B). Either actions or stimuli can be decoded, in continuous-time or on a trial-by-trial basis (112, 129-134). Performance is always measured on a test set of time points or trials not used for training the decoder. The degree of decodability of a variable may be informative in itself, especially when analyzed across time or spatial locations or when compared across brain areas, cell types, or experimental manipulations (106, 135-138). Linear decoding of abstract variables such as image categories may be used as evidence for neural computations that perform input untangling and increase representational invariance (139). The scaling of the error as the number of neurons and trials increases can also be informative. It has been shown that in some cortical sensory systems, the stimulus decoding error approaches zero, which is inconsistent with some information theory-based explanations for the origins of perceptual errors (Fig. 3A) (56). When the decoding error does asymptote at a nonzero value, the deviations of the decoder from the true stimulus may be indicative of internal variability and can be interpreted with respect to behaviors such as those reporting a perceptual state in trained animals (140). Last, the decoder may be used to "decode" fully internally generated neural activity such as that which occurs during replay or during planning in the hippocampus or HVC (Fig. 3B) (141, 142).

Statistical pitfalls for large-scale analyses

Large-scale analyses have several pitfalls that can be overcome by using more rigorous methodology. To start, it is important to reduce the statistical dependence between testing and training data points as much as possible, which may not be easy to achieve (143). For example, slowly varying behavioral variables (such as arousal or pose) may appear to be decodable on test data if both the neural activity and the behavior contain slow timescales, which thus create "nonsense" correlations (144). To reduce the dependence between train and test data, time points or trials can be split into temporally extended blocks, and each block may be assigned randomly to either train or test data (21, 111, 143).

The interpretation of the decoding error can also be a major pitfall because this error is sometimes assumed to represent neuronal uncertainty about a decoded stimulus, about spatial locations, or about latent reward states in an environment. However, if an insufficient number of neurons or trials is recorded, the error is more likely to simply represent decoder noise, which retains the structure of the noise on the training trials. When similarities in noise structure exist between a test and a train trial, the decoder will output predictions on that test trial that resemble the training trial. This type of noise cannot be distinguished from real signal, unless a very large number of trials and neurons is available. However, with sufficiently many neurons and trials, the decoding error may approach zero in at least some scenarios (75), in which case the decoder variability can only be noise rather than representing internal uncertainty. Thus, a high bar needs to be cleared before a decoding error can be interpreted as perceptual uncertainty; further, perceptual uncertainty may instead be encoded and used in different ways by the brain (145). Generally, large-scale recordings may allow any variable of interest to be decoded with sufficiently many neurons, and this may diminish the role of decoding analyses for analyzing neural data.

Simple analyses are also not immune from statistical pitfalls, especially when applied to large-scale neural data. These pitfalls generally arise from the variability of single neurons, which makes the estimation of their parameters noisy. This estimation thus becomes a statistical problem and must be treated with statistical rigor. As a simple example, consider the case of sorting neurons in a raster plot according to their latency of firing after a stimulus (Fig. 3C). When the data are displayed on the trials that were used to obtain the sorting (training trials), it appears that the neurons fire in an orderly sequence. However, this can be a consequence of overfitting to noisy data, and the same ordering may give completely random latencies or positions on test trials from the same experimental condition. Few studies perform this analysis in a crossvalidated way; an example of the correct approach is provided in (146).

Many such scenarios exist, and they may not be easy to recognize. Consider the case of constructing population tuning curves by aligning and averaging single-neuron tuning curves (Fig. 3C). Implicit in this process is the selection of the preferred stimulus for each neuron, which is similar to the process of selecting the best latency in the previous example. Thus, a similar bias exists on training trials, in which an apparent sensory response may simply be due to overfitting. When a genuine sensory response is present, it may appear larger on the training trials than it really is. These biases are typically eliminated or much reduced on test trials. The biases might be particularly difficult to detect in the context of manipulations. In a hypothetical scenario, if the preferred stimulus is found by using the control

condition, the tuning curves after a manipulation may appear wider and smaller in amplitude relative to the training trials, when their width could be unchanged and their amplitude increased when correctly compared with test trials of the "control" condition. Methods to detect and address such scenarios are unfortunately not sufficiently widespread.

Discussion

This Review can be used as a guide for scientists with large-scale neural data (Fig. 4, graphical summary). Many of the methods described above can be used in an exploratory fashion to identify new relations between neural activity and behavior or stimuli. Quantification can typically be done with encoding or decoding analyses, which can be applied directly to neural data or to the reduced components found by use of unsupervised analysis methods. There are numerous pitfalls that must be avoided when quantifying high-dimensional data with complex analyses, and these pitfalls may be difficult to identify without sufficient experience. Future theoretical and computational studies may be especially helpful to address these problems, but they must target an audience of experimental scientists. Some examples exist of statistical studies with a combination of simple, intuitive explanations (147, 148), but an expansion of the statistical neuroscience field will be necessary to avoid the numerous pitfalls that largescale data introduces.

Looking forward, animal neuroscience is gaining a renewed focus on more complex and more ethological behaviors-for example, in freely moving animals in large environments (149). Large-scale neural recordings in these contexts are becoming possible, potentially with continuous 24/7 monitoring of neural activity and behavior. The large variability in these scenarios may create insurmountable difficulties for most of the analysis methods we describe above and will require innovation. The large number of potential nuisance variables (such asenvironmental variables or spontaneous movements) are very difficult to control for in such scenarios. Model-based approaches will be needed for refining and testing hypotheses iteratively through manipulations of the environment, task, or the neural activity.

For more standard neuroscience experiments, there will still likely be substantial future progress in analysis methods. The obstacles in the acquisition of such data have been substantially reduced, and many laboratories are finding themselves with more data than they know what to do with. Establishing collaborations with computational and theoretical scientists will be key to taking advantage of this data. To drive faster progress in this field, experimental neuroscientists will need to adopt new mathematical concepts and analysis techniques. Summer schools such as the Neuromatch Academy offer good opportunities for training in data science methods specifically with neuroscience applications (*150*). Now may be a good time for both newcomers and established neuroscientists to refresh and update their knowledge of mathematical concepts and techniques.

REFERENCES AND NOTES

- T.-W. Chen *et al.*, Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295–300 (2013). doi: 10.1038/nature12354; pmid: 23868258
- N. J. Sofroniew, D. Flickinger, J. King, K. Svoboda, A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *eLife* 5, e14472 (2016). doi: 10.7554/eLife.14472; pmid: 27300105
- J. J. Jun et al., Fully integrated silicon probes for high-density recording of neural activity. *Nature* 551, 232–236 (2017). doi: 10.1038/nature24636; pmid: 29120427
- N. A. Steinmetz *et al.*, Neuropixels 2.0: A miniaturized highdensity probe for stable, long-term brain recordings. *Science* **372**, eabf4588 (2021). doi: 10.1126/science.abf4588; pmid: 33859006
- E. M. Trautmann *et al.*, Large-scale high-density brain-wide neural recording in nonhuman primates. bioRxiv 526664 [Preprint] (2023); https://doi.org/10.1101/2023.02.01.526664.
- D. J. Tolhurst, J. A. Movshon, A. F. Dean, The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* 23, 775–785 (1983). doi: 10.1016/ 0042-6989(83)90200-6; pmid: 6623937
- J. Fiser, C. Chiu, M. Weliky, Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* 431, 573–578 (2004). doi: 10.1038/nature02907; pmid: 15457262
- T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, E. I. Moser, Microstructure of a spatial map in the entorhinal cortex. *Nature* 436, 801–806 (2005). doi: 10.1038/nature03721; pmid: 15965463
- A. C. Paulk et al., Large-scale neural recordings with single neuron resolution using Neuropixels probes in human cortex. Nat. Neurosci. 25, 252–263 (2022). doi: 10.1038/s41593-021-00997-0; omid: 35102333
- J. E. Chung *et al.*, High-density single-unit human cortical recordings using the Neuropixels probe. *Neuron* **110**, 2409–2421.e3 (2022). doi: 10.1016/j.neuron.2022.05.007; pmid: 35679860
- J. Müller et al., High-resolution CMOS MEA platform to study neurons at subcellular, cellular, and network levels. Lab Chip 15, 2767–2780 (2015). doi: 10.1039/C5LC00133A; pmid: 25973786
- S. Zhao *et al.*, Tracking neural activity from the same cells during the entire adult life of mice. *Nat. Neurosci.* 26, 696–710 (2023). doi: 10.1038/s41593-023-01267-x; pmid: 36804648
- C.-H. Chiang *et al.*, Development of a neural interface for high-definition, long-term recording in rodents and nonhuman primates. *Sci. Transl. Med.* **12**, eaay4682 (2020). doi: 10.1126/scitranslmed.aay4682; pmid: 32269166
- M. Pachitariu, S. Sridhar, J. Pennington, C. Stringer, Spike sorting with Kilosort4. *Nat. Methods* 21, 914–921 (2024). doi: 10.1038/s41592-024-02232-7; pmid: 38589517
- J. E. Chung *et al.*, A fully automated approach to spike sorting. *Neuron* **95**, 1381–1394.e6 (2017). doi: 10.1016/ i.neuron.2017.08.030; pmid: 28910621
- J. J. Jun et al., Real-time spike sorting platform for highdensity extracellular probes with ground-truth validation and drift correction. bioRxiv 101030 [Preprint] (2017); https://doi.org/10.1101/101030.
- C. Grienberger, A. Giovannucci, W. Zeiger, C. Portera-Cailliau, Two-photon calcium imaging of neuronal activity. *Nat. Rev. Methods Primers* 2, 67 (2022). doi: 10.1038/s43586-022-00147-1; pmid: 38124998
- W. Denk, J. H. Strickler, W. W. Webb, Two-photon laser scanning fluorescence microscopy. *Science* 248, 73–76 (1990). doi: 10.1126/science.2321027; pmid: 2321027
- M. B. Ahrens, M. B. Orger, D. N. Robson, J. M. Li, P. J. Keller, Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nat. Methods* **10**, 413–420 (2013). doi: 10.1038/nmeth.2434; pmid: 23524393

- M. Pachitariu et al., Suite2p: Beyond 10,000 neurons with standard two-photon microscopy, bioRxiv 061507 [Preprint] (2016): https://doi.org/10.1101/061507.
- C. Stringer *et al.*, Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019). doi: 10.1126/science.aav7893; pmid: 31000656
- D. Aharoni, B. S. Khakh, A. J. Silva, P. Golshani, All the light that we can see: A new era in miniaturized microscopy. *Nat. Methods* 16, 11–13 (2019). doi: 10.1038/s41592-018-0266-x; pmid: 30573833
- W. Zong et al., Large-scale two-photon calcium imaging in freely moving mice. Cell 185, 1240–1256.e30 (2022). doi: 10.1016/j.cell.2022.02.017; pmid: 35305313
- Y. Zhang *et al.*, Fast and sensitive GCaMP calcium indicators for imaging neural populations. *Nature* **615**, 884–891 (2023). doi: 10.1038/s41586-023-05828-9; pmid: 36922596
- G. J. Broussard, L. Petreanu, Eavesdropping wires: Recording activity in axons using genetically encoded calcium indicators. J. Neurosci. Methods 360, 109251 (2021). doi: 10.1016/j.jneumeth.2021.109251; pmid: 34119572
- Y. Chen *et al.*, Soma-targeted imaging of neural circuits by ribosome tethering. *Neuron* **107**, 454–469.e6 (2020). doi: 10.1016/j.neuron.2020.05.005; pmid: 32574560
- T. Ebina *et al.*, Two-photon imaging of neuronal activity in motor cortex of marmosets during upper-limb movement tasks. *Nat. Commun.* 9, 1879 (2018). doi: 10.1038/s41467-018-04286-6; pmid: 29760466
- A. K. Garg, P. Li, M. S. Rashid, E. M. Callaway, Color and orientation are jointly coded and spatially organized in primate primary visual cortex. *Science* **364**, 1275–1279 (2019). doi: 10.1126/science.aaw5868; pmid: 31249057
- P. Zhou *et al.*, Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *eLife* 7, e28728 (2018). doi: 10.7554/eLife.28728; pmid: 29469809
- A. Giovannucci et al., CalmAn an open source tool for scalable calcium imaging data analysis. eLife 8, e38173 (2019). doi: 10.7554/eLife.38173; pmid: 30652683
- S. E. J. de Vries *et al.*, A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nat. Neurosci.* 23, 138–151 (2020). doi: 10.1038/s41593-019-0550-9; pmid: 31844315
- E. H. Lyall et al., Synthesis of a comprehensive population code for contextual features in the awake sensory cortex. eLife 10, e62687 (2021). doi: 10.7554/eLife.62687; pmid: 34723796
- L. Feigin, G. Tasaka, I. Maor, A. Mizrahi, Sparse coding in temporal association cortex improves complex sound discriminability. *J. Neurosci.* 41, 7048–7064 (2021). doi: 10.1523/JNEUROSCI.3167-20.2021; pmid: 34244361
- J. H. Siegle et al., Survey of spiking in the mouse visual system reveals functional hierarchy. Nature 592, 86–92 (2021). doi: 10.1038/s41586-020-03171-x; pmid: 33473216
- S. Schröder *et al.*, Arousal modulates retinal output. *Neuron* 107, 487–495.e9 (2020). doi: 10.1016/j.neuron.2020.04.026; pmid: 32445624
- A. Fratzl et al., Flexible inhibitory control of visually evoked defensive behavior by the ventral lateral geniculate nucleus. *Neuron* 109, 3810–3822.e9 (2021). doi: 10.1016/ j.neuron.2021.09.003; pmid: 34614420
- Y. Liu et al., Mapping visual functions onto molecular cell types in the mouse superior colliculus. *Neuron* 111, 1876–1886.e5 (2023). doi: 10.1016/j.neuron.2023.03.036; pmid: 37086721
- G. Matteucci et al., Cortical sensory processing across motivational states during goal-directed behavior. Neuron 110, 4176–4193.e10 (2022). doi: 10.1016/ i.neuron.2022.09.032; pmid: 36240769
- W. E. Allen *et al.*, Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science* 364, 253 (2019). doi: 10.1126/science.aav3932; pmid: 30948440
- N. A. Steinmetz, P. Zatka-Haas, M. Carandini, K. D. Harris, Distributed coding of choice, action and engagement across the mouse brain. *Nature* 576, 266–273 (2019). doi: 10.1038/ s41586-019-1787-x; pmid: 31776518
- J. B. Issa, B. A. Radvansky, F. Xuan, D. A. Dombeck, Lateral entorhinal cortex subpopulations represent experiential epochs surrounding reward. *Nat. Neurosci.* 27, 536–546 (2024). doi: 10.1038/s41593-023-01557-4: pmid: 38272968
- D. A. Aponte et al., Recurrent network dynamics shape direction selectivity in primary auditory cortex. *Nat. Commun.* 12, 314 (2021). doi: 10.1038/s41467-020-20590-6; pmid: 33436635
- 43. A. J. Duszkiewicz et al., Local origin of excitatory-inhibitory tuning equivalence in a cortical network. Nat. Neurosci. 27,

782-792 (2024). doi: 10.1038/s41593-024-01588-5; pmid: 38491324

- W. Xiao, S. Sharma, G. Kreiman, M. S. Livingstone, Featureselective responses in macaque visual cortex follow eye movements during natural vision. *Nat. Neurosci.* 27, 1157–1166 (2024). doi: 10.1038/s41593-024-01631-5; pmid: 38684892
- S. Chen *et al.*, Brain-wide neural activity underlying memoryguided movement. *Cell* **187**, 676–691.e16 (2024). doi: 10.1016/j.cell.2023.12.035; pmid: 38306983
- Q. Gaucher *et al.*, Complexity of frequency receptive fields predicts tonotopic variability across species. *eLife* 9, e53462 (2020). doi: 10.7554/eLife.53462; pmid: 32420865
- L. Zhong *et al.*, Distinct streams for supervised and unsupervised learning in the visual cortex. bioRxiv 581990 [Preprint] (2024); https://doi.org/10.1101/ 2024.02.25.581990.
- N. Li, K. Daie, K. Svoboda, S. Druckmann, Robust neuronal dynamics in premotor cortex during motor planning. *Nature* 532, 459–464 (2016). doi: 10.1038/nature17643; pmid: 27074502
- D. J. O'Shea *et al.*, Direct neural perturbations reveal a dynamical mechanism for robust computation. bioRxiv 520768 [Preprint] (2022); https://doi.org/10.1101/ 2022.12.16.520768.
- A. J. Zimnik, M. M. Churchland, Independent generation of sequence elements by motor cortex. *Nat. Neurosci.* 24, 412–424 (2021). doi: 10.1038/s41593-021-00798-5; pmid: 33619403
- S. Han, F. Helmchen, Behavior-relevant top-down crossmodal predictions in mouse neocortex. *Nat. Neurosci.* 27, 298–308 (2024). doi: 10.1038/s41593-023-01534-x; pmid: 38177341
- W. Yang, S. L. Tipparaju, G. Chen, N. Li, Thalamus-driven functional populations in frontal cortex support decisionmaking. *Nat. Neurosci.* 25, 1339–1352 (2022). doi: 10.1038/ s41593-022-01171-w; pmid: 36171427
- A.-K. Eiselt *et al.*, Hunger or thirst state uncertainty is resolved by outcome evaluation in medial prefrontal cortex to guide decision-making. *Nat. Neurosci.* 24, 907–912 (2021). doi: 10.1038/s41593-021-00850-4; pmid: 33972802
- D. Kobak et al., Demixed principal component analysis of neural population data. eLife 5, e10989 (2016). doi: 10.7554/ eLife.10989; pmid: 27067378
- O. G. Sani, H. Abbaspourazad, Y. T. Wong, B. Pesaran, M. M. Shanechi, Modeling behaviorally relevant neural dynamics enabled by preferential subspace identification. *Nat. Neurosci.* 24, 140–149 (2021). doi: 10.1038/s41593-020-00733-0; pmid: 33169030
- C. Stringer, M. Michaelos, D. Tsyboulski, S. E. Lindo, M. Pachitariu, High-precision coding in visual cortex. *Cell* 184, 2767–2778.e15 (2021). doi: 10.1016/j.cell.2021.03.042; pmid: 33857423
- W.-C. Jiang, S. Xu, J. T. Dudman, Hippocampal representations of foraging trajectories depend upon spatial context. *Nat. Neurosci.* 25, 1693–1705 (2022). doi: 10.1038/ s41593-022-01201-7; pmid: 36446934
- R. Sasaki et al., Balancing risk-return decisions by manipulating the mesofrontal circuits in primates. *Science* 383, 55–61 (2024). doi: 10.1126/science.adj6645; pmid: 38175903
- B. M. Yu et al., Gaussian-process factor analysis for lowdimensional single-trial analysis of neural population activity. J. Neurophysiol. **102**, 614–635 (2009). doi: 10.1152/ jn.90941.2008; pmid: 19357332
- A. H. Williams *et al.*, Unsupervised discovery of demixed, low-dimensional neural dynamics across multiple timescales through tensor component analysis. *Neuron* 98, 1099–1115.e8 (2018). doi: 10.1016/j.neuron.2018.05.015; pmid: 29887338
- L. Duncker, M. Sahani, Temporal alignment and latent gaussian process factor inference in population spike trains. bioRxiv 331751 [Preprint] (2018); https://doi.org/10.1101/331751.
- M. Genkin, O. Hughes, T. A. Engel, Learning non-stationary Langevin dynamics from stochastic observations of latent trajectories. *Nat. Commun.* 12, 5986 (2021). doi: 10.1038/ s41467-021-26202-1; pmid: 34645828
- O. Karniol-Tambour *et al.*, "Modeling state-dependent communication between brain regions with switching nonlinear dynamical systems" in *The Twelfth International Conference on Learning Representations* (2023).
- 64. N. E. Bush, J.-M. Ramirez, Latent neural population dynamics underlying breathing, opioid-induced respiratory depression

and gasping. *Nat. Neurosci.* **27**, 259–271 (2024). doi: 10.1038/s41593-023-01520-3; pmid: 38182835

- S. Wang, R. Falcone, B. Richmond, B. B. Averbeck, Attractor dynamics reflect decision confidence in macaque prefrontal cortex. *Nat. Neurosci.* 26, 1970–1980 (2023). doi: 10.1038/ s41593-023-01445-x; pmid: 37798412
- A. R. Khanna et al., Single-neuronal elements of speech production in humans. *Nature* 626, 603–610 (2024). doi: 10.1038/s41586-023-06982-w; pmid: 38297120
- B. K. Hulse *et al.*, A connectome of the *Drosophila* central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. *eLife* **10**, e66039 (2021). doi: 10.7554/eLife.66039; pmid: 34696823
- E. Marder, D. Bucher, D. J. Schulz, A. L. Taylor, Invertebrate central pattern generation moves along. *Curr. Biol.* 15, R685–R699 (2005). doi: 10.1016/j.cub.2005.08.022; pmid: 16139202
- A. R. Galgali, M. Sahani, V. Mante, Residual dynamics resolves recurrent contributions to neural computation. *Nat. Neurosci.* 26, 326–338 (2023). doi: 10.1038/s41593-022-01230-2; pmid: 36635498
- J. Soldado-Magraner, V. Mante, M. Sahani, Inferring contextdependent computations through linear approximations of prefrontal cortex dynamics. bioRxiv 527389 [Preprint] (2023); https://doi.org/10.1101/2023.02.06.527389.
- R. J. Gardner *et al.*, Toroidal topology of population activity in grid cells. *Nature* **602**, 123–128 (2022). doi: 10.1038/s41586-021-04268-7; pmid: 35022611
- L. Avitan et al., Spontaneous and evoked activity patterns diverge over development. eLife 10, e61942 (2021). doi: 10.7554/eLife.61942; pmid: 33871351
- L. Avitan, C. Stringer, Not so spontaneous: Multi-dimensional representations of behaviors and context in sensory areas. *Neuron* 110, 3064–3075 (2022). doi: 10.1016/ j.neuron.2022.06.019; pmid: 35863344
- P. Gao *et al.*, A theory of multineuronal dimensionality, dynamics and measurement. bioRxiv 214262 [Preprint] (2017); https://doi.org/10.1101/214262.
- C. Stringer, M. Pachitariu, N. Steinmetz, M. Carandini, K. D. Harris, High-dimensional geometry of population responses in visual cortex. *Nature* 571, 361–365 (2019). doi: 10.1038/s41586-019-1346-5; pmid: 31243367
- F. Lanore, N. A. Cayco-Gajic, H. Gurnani, D. Coyle, R. A. Silver, Cerebellar granule cell axons support high-dimensional representations. *Nat. Neurosci.* 24, 1142–1150 (2021). doi: 10.1038/s41593-021-00873-x; pmid: 34168340
- N. C. L. Kong, E. Margalit, J. L. Gardner, A. M. Norcia, Increasing neural network robustness improves match to macaque VI eigenspectrum, spatial frequency preference and predictivity. *PLOS Comput. Biol.* **18**, e1009739 (2022). doi: 10.1371/journal.pcbi.1009739; pmid: 34995280
- C. F. Cadieu *et al.*, Deep neural networks rival the representation of primate IT cortex for core visual object recognition. *PLOS Comput. Biol.* **10**, e1003963 (2014). doi: 10.1371/journal.pcbi.1003963; pmid: 25521294
- B. Dudok *et al.*, Recruitment and inhibitory action of hippocampal axo-axonic cells during behavior. *Neuron* **109**, 3838–3850.e8 (2021). doi: 10.1016/j.neuron.2021.09.033; pmid: 34648750
- V. Samborska, J. L. Butler, M. E. Walton, T. E. J. Behrens, T. Akam, Complementary task representations in hippocampus and prefrontal cortex for generalizing the structure of problems. *Nat. Neurosci.* 25, 1314–1326 (2022). doi: 10.1038/s41593-022-01149-8; pmid: 36171429
- S. L. Pashkovski et al., Structure and flexibility in cortical representations of odour space. *Nature* 583, 253–258 (2020). doi: 10.1038/s41586-020-2451-1; pmid: 32612230
- H. C. Barron *et al.*, Neuronal computation underlying inferential reasoning in humans and mice. *Cell* **183**, 228–243. e21 (2020). doi: 10.1016/j.cell.2020.08.035; pmid: 32946810
- A. H. Williams, E. Kunz, S. Kornblith, S. W. Linderman, Generalized shape metrics on neural representations. *Adv. Neural Inf. Process. Syst.* **34**, 4738–4750 (2021). pmid: 38170102
- M. Safaie et al., Preserved neural dynamics across animals performing similar behaviour. Nature 623, 765–771 (2023). doi: 10.1038/s41586-023-06714-0; pmid: 37938772
- L. N. Driscoll, L. Duncker, C. D. Harvey, Representational drift: Emerging theories for continual learning and experimental future directions. *Curr. Opin. Neurobiol.* **76**, 102609 (2022). doi: 10.1016/j.conb.2022.102609; pmid: 35939861
- C. Giusti, E. Pastalkova, C. Curto, V. Itskov, Clique topology reveals intrinsic geometric structure in neural correlations.

Proc. Natl. Acad. Sci. U.S.A. **112**, 13455–13460 (2015); doi: 10.1073/pnas.1506407112; pmid: 26487684

- R. Chaudhuri, B. Gerçek, B. Pandey, A. Peyrache, I. Fiete, The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat. Neurosci.* 22, 1512–1520 (2019). doi: 10.1038/s41593-019-0460-x; pmid: 31406365
- E. H. Nieh *et al.*, Geometry of abstract learned knowledge in the hippocampus. *Nature* **595**, 80–84 (2021). doi: 10.1038/ s41586-021-03652-7; pmid: 34135512
- S. A. Koay, A. S. Charles, S. Y. Thiberge, C. D. Brody, D. W. Tank, Sequential and efficient neural-population coding of complex task information. *Neuron* **110**, 328–349.e11 (2022). doi: 10.1016/j.neuron.2021.10.020; pmid: 34776042
- W. Yang et al., Selection of experience for memory by hippocampal sharp wave ripples. Science 383, 1478–1483 (2024). doi: 10.1126/science.adk8261; pmid: 38547293
- L. Duncker, M. Sahani, Dynamics on the manifold: Identifying computational dynamical activity from neural population recordings. *Curr. Opin. Neurobiol.* **70**, 163–170 (2021). doi: 10.1016/j.conb.2021.10.014; pmid: 34837752
- G. Bondanelli, T. Deneux, B. Bathellier, S. Ostojic, Network dynamics underlying OFF responses in the auditory cortex. *eLife* 10, e53151 (2021). doi: 10.7554/eLife.53151; pmid: 33759763
- B. Akitake et al., Amplified cortical neural responses as animals learn to use novel activity patterns. *Curr. Biol.* 33, 2163–2174.e4 (2023). doi: 10.1016/j.cub.2023.04.032; pmid: 37148876
- 94. B. Wang et al., Unsupervised restoration of a complex learned behavior after large-scale neuronal perturbation. Nat. Neurosci. 27, 1176–1186 (2024). doi: 10.1038/s41593-024-01630-6; pmid: 38684893
- A. Nair *et al.*, An approximate line attractor in the hypothalamus encodes an aggressive state. *Cell* 186, 178–193.e15 (2023). doi: 10.1016/j.cell.2022.11.027; pmid: 36608653
- R. T. Chen, Y. Rubanova, J. Bettencourt, D. K. Duvenaud, Neural ordinary differential equations. *Adv. Neural Inf. Process. Syst.* **31**, 6572 (2018).
- T. D. Kim, T. Z. Luo, J. W. Pillow, C. D. Brody, "Inferring latent dynamics underlying neural population activity via neural differential equations" in *International Conference on Machine Learning* (PMLR, 2021), p. 5551.
- B. A. Sauerbrei et al., Cortical pattern generation during dexterous movement is input-driven. Nature 577, 386–391 (2020). doi: 10.1038/s41586-019-1869-9; pmid: 31875851
- M. K. Leonard *et al.*, Large-scale single-neuron speech sound encoding across the depth of human cortex. *Nature* **626**, 593–602 (2024). doi: 10.1038/s41586-023-06839-2; pmid: 38093008
- C. Pandarinath et al., Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat. Methods* 15, 805–815 (2018). doi: 10.1038/s41592-018-0109-9; pmid: 30224673
- A. Pellegrino, H. Stein, N. A. Cayco-Gajic, Dimensionality reduction beyond neural subspaces with slice tensor component analysis. *Nat. Neurosci.* 27, 1199–1210 (2024). doi: 10.1038/s41593-024-01626-2; pmid: 38710876
- 102. F. Zhu et al., A deep learning framework for inference of single-trial neural population dynamics from calcium imaging with subframe temporal resolution. *Nat. Neurosci.* 25, 1724–1734 (2022). doi: 10.1038/s41593-022-01189-0; pmid: 36424431
- E. L. Sylwestrak *et al.*, Cell-type-specific population dynamics of diverse reward computations. *Cell* **185**, 3568–3587.e27 (2022). doi: 10.1016/j.cell.2022.08.019; pmid: 36113428
- S. Schneider, J. H. Lee, M. W. Mathis, Learnable latent embeddings for joint behavioural and neural analysis. *Nature* 617, 360–368 (2023). doi: 10.1038/s41586-023-06031-6; pmid: 37138088
- J. Brown et al., Spatial integration during active tactile sensation drives orientation perception. *Neuron* 109, 1707–1720.e7 (2021). doi: 10.1016/j.neuron.2021.03.020; pmid: 33826906
- C. A. Zimmerman et al., A neural mechanism for learning from delayed postingestive feedback. bioRxiv 561214 [Preprint] (2023); https://doi.org/10.1101/2023.10.06.561214.
- L. Petrucco *et al.*, Neural dynamics and architecture of the heading direction circuit in zebrafish. *Nat. Neurosci.* 26, 765–773 (2023). doi: 10.1038/s41593-023-01308-5; pmid: 37095397

- M. H. Plitt, L. M. Giocomo, Experience-dependent contextual codes in the hippocampus. *Nat. Neurosci.* 24, 705–714 (2021). doi: 10.1038/s41593-021-00816-6; pmid: 33753945
- J. H. Macke et al., Empirical models of spiking in neural populations. Adv. Neural Inf. Process. Syst. 24, 1350 (2011).
- C. Stringer *et al.*, Rastermap: A discovery method for neural population recordings. *Nat. Neurosci.* 10.1038/ s41593-024-01783-4 (2024).
- A. Syeda et al., Facemap: A framework for modeling neural activity based on orofacial tracking. Nat. Neurosci. 27, 187–195 (2024). doi: 10.1038/s41593-023-01490-6; pmid: 37985801
- 112. F. Cazettes *et al.*, A reservoir of foraging decision variables in the mouse brain. *Nat. Neurosci.* **26**, 840–849 (2023). doi: 10.1038/s41593-023-01305-8; pmid: 37055628
- S. Musall, M. T. Kaufman, A. L. Juavinett, S. Gluf, A. K. Churchland, Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* 22, 1677–1686 (2019). doi: 10.1038/s41593-019-0502-4; pmid: 31551604
- E. Zagha et al., The importance of accounting for movement when relating neuronal activity to sensory and cognitive processes. J. Neurosci. 42, 1375–1382 (2022). doi: 10.1523/ JNEUROSCI.1919-21.2021; pmid: 35027407
- S. Tremblay, C. Testard, R. W. DiTullio, J. Inchauspé, M. Petrides, Neural cognitive signals during spontaneous movements in the macaque. *Nat. Neurosci.* 26, 295–305 (2023). doi: 10.1038/s41593-022-01220-4; pmid: 36536242
- C. Bimbard *et al.*, Behavioral origin of sound-evoked activity in mouse visual cortex. *Nat. Neurosci.* 26, 251–258 (2023). doi: 10.1038/s41593-022-01227-x; pmid: 36624279
- M. N. Oude Lohuis, P. Marchesi, U. Olcese, C. M. A. Pennartz, Triple dissociation of visual, auditory and motor processing in mouse primary visual cortex. *Nat. Neurosci.* 27, 758–771 (2024). doi: 10.1038/s41593-023-01564-5; pmid: 38307971
- N. Shahidi et al., Population coding of strategic variables during foraging in freely moving macaques. *Nat. Neurosci.* 27, 772–781 (2024). doi: 10.1038/s41593-024-01575-w; pmid: 38443701
- J. Singh Alvarado et al., Neural dynamics underlying birdsong practice and performance. Nature 599, 635–639 (2021). doi: 10.1038/s41586-021-04004-1; pmid: 34671166
- I. M. Park, M. L. Meister, A. C. Huk, J. W. Pillow, Encoding and decoding in parietal cortex during sensorimotor decisionmaking. *Nat. Neurosci.* 17, 1395–1403 (2014). doi: 10.1038/ nn.3800; pmid: 25174005
- P. R. L. Parker, E. T. T. Abe, E. S. P. Leonard, D. M. Martins, C. M. Niell, Joint coding of visual input and eye/head position in V1 of freely moving mice. *Neuron* **110**, 3897–3906.e5 (2022). doi: 10.1016/j.neuron.2022.08.029; pmid: 36137549
- P. M. Goltstein, S. Reinert, T. Bonhoeffer, M. Hübener, Mouse visual cortex areas represent perceptual and semantic features of learned visual categories. *Nat. Neurosci.* 24, 1441–1451 (2021). doi: 10.1038/s41593-021-00914-5; pmid: 34545249
- C. Buetfering et al., Behaviorally relevant decision coding in primary somatosensory cortex neurons. *Nat. Neurosci.* 25, 1225–1236 (2022). doi: 10.1038/s41593-022-01151-0; pmid: 36042310

- K. Franke et al., State-dependent pupil dilation rapidly shifts visual feature selectivity. Nature 610, 128–134 (2022). doi: 10.1038/s41586-022-05270-3; pmid: 36171291
- 125. S. Sabesan, A. Fragner, C. Bench, F. Drakopoulos, N. A. Lesica, Large-scale electrophysiology and deep learning reveal distorted neural signal dynamics after hearing loss. *eLife* 12, e85108 (2023). doi: 10.7554/eLife.85108; pmid: 37162188
- J. W. Lovelace *et al.*, Vagal sensory neurons mediate the Bezold-Jarisch reflex and induce syncope. *Nature* **623**, 387–396 (2023). doi: 10.1038/s41586-023-06680-7; pmid: 37914931
- Ye, C. Pandarinath, Representation learning for neural population activity with neural data transformers. *Neurons Behav. Data Anal. Theory* 5, 1–18 (2021). doi: 10.51628/ 001c.27358
- Y. Zhang et al., Towards a "universal translator" for neural dynamics at single-cell, single-spike resolution. arXiv:2407.14668 [q-bio.NC] (2024).
- N. Kriegeskorte, P. K. Douglas, Interpreting encoding and decoding models. *Curr. Opin. Neurobiol.* 55, 167–179 (2019). doi: 10.1016/j.conb.2019.04.002; pmid: 31039527
- J. I. Glaser *et al.*, Machine learning for neural decoding. *eNeuro* **7**, ENEURO.0506-19.2020 (2020). doi: 10.1523/ ENEURO.0506-19.2020; pmid: 32737181
- Y. J. Kim et al., Nonlinear decoding of natural images from large-scale primate retinal ganglion recordings. *Neural Comput.* 33, 1719–1750 (2021). doi: 10.1162/neco_a_01395; pmid: 34411268
- M. Azabou et al., A unified, scalable framework for neural population decoding. Adv. Neural Inf. Process. Syst. 36, 44937 (2024).
- N. L. Pettit, X. C. Yuan, C. D. Harvey, Hippocampal place codes are gated by behavioral engagement. *Nat. Neurosci.* **25**, 561–566 (2022). doi: 10.1038/s41593-022-01050-4; pmid: 35449355
- 134. Y. Xie et al., Geometry of sequence working memory in macaque prefrontal cortex. Science 375, 632–639 (2022). doi: 10.1126/science.abm0204; pmid: 35143322
- S. Keshavarzi et al., Multisensory coding of angular head velocity in the retrosplenial cortex. *Neuron* **110**, 532–543.e9 (2022). doi: 10.1016/j.neuron.2021.10.031; pmid: 34788632
- A. Nagelhus, S. O. Andersson, S. G. Cogno, E. I. Moser, M.-B. Moser, Object-centered population coding in CA1 of the hippocampus. *Neuron* **111**, 2091–2104.e14 (2023). doi: 10.1016/i.neuron.2023.04.008. pmid: 37148872
- J. A. Charlton, R. L. T. Goris, Abstract deliberation by visuomotor neurons in prefrontal cortex. *Nat. Neurosci.* 27, 1167–1175 (2024). doi: 10.1038/s41593-024-01635-1; pmid: 38684894
- A. Bellafard, G. Namvar, J. C. Kao, A. Vaziri, P. Golshani, Volatile working memory representations crystallize with practice. *Nature* 629, 1109–1117 (2024). doi: 10.1038/ s41586-024-07425-w; pmid: 38750359
- J. J. DiCarlo, D. Zoccolan, N. C. Rust, How does the brain solve visual object recognition? *Neuron* 73, 415–434 (2012). doi: 10.1016/j.neuron.2012.01.010; pmid: 22325196

- 140. V. B. McGinty, S. M. Lupkin, Behavioral read-out from population value signals in primate orbitofrontal cortex. *Nat. Neurosci.* 26, 2203–2212 (2023). doi: 10.1038/s41593-023-01473-7; pmid: 37932464
- 141. M. Elmaleh, D. Kranz, A. C. Asensio, F. W. Moll, M. A. Long, Sleep replay reveals premotor circuit structure for a skilled behavior. *Neuron* **109**, 3851–3861.e4 (2021). doi: 10.1016/ j.neuron.2021.09.021; pmid: 34626637
- 142. M. Wang, D. J. Foster, B. E. Pfeiffer, Alternating sequences of future and past behavior encoded within hippocampal theta oscillations. *Science* **370**, 247–250 (2020). doi: 10.1126/ science.abb4151; pmid: 33033222
- 143. F. Nevjen, B. A. Dunn, Insights in neuronal tuning: Navigating the statistical challenges of autocorrelation and missing variables. *eLife* 13, RP92517 (2024).
- K. D. Harris, Nonsense correlations in neuroscience. bioRxiv 402719 [Preprint] (2020); https://doi.org/10.1101/ 2020.11.29.402719.
- M. Salmasi, M. Sahani, "Learning neural codes for perceptual uncertainty" in 2022 IEEE International Symposium on Information Theory (ISIT) (IEEE, 2022), pp. 2463–2468. doi: 10.1109/ISIT50566.2022.9834606
- 146. L. F. Fischer, R. Mojica Soto-Albors, F. Buck, M. T. Harnett, Representation of visual landmarks in retrosplenial cortex. *eLife* 9, e51458 (2020). doi: 10.7554/eLife.51458; pmid: 32154781
- 147. G. F. Elsayed, J. P. Cunningham, Structure in neural population recordings: An expected byproduct of simpler phenomena? *Nat. Neurosci.* **20**, 1310–1318 (2017). doi: 10.1038/nn.4617; pmid: 28783140
- 148. R. Kass, The two cultures: Statistics and machine learning in science. Obs. Stud. 7, 135–144 (2021). doi: 10.1353/ obs.2021.0000
- 149. E. J. Dennis et al., Systems neuroscience of natural behaviors in rodents. J. Neurosci. 41, 911–919 (2021). doi: 10.1523/ JNEUROSCI.1877-20.2020; pmid: 33443081
- 150. T. van Viegen et al., Neuromatch academy: Teaching computational neuroscience with global accessibility. *Trends Cogn. Sci.* 25, 535–538 (2021). doi: 10.1016/ j.tics.2021.03.018; pmid: 33994097

ACKNOWLEDGMENTS

We thank S. Romani and T. Ahamed for helpful suggestions. **Funding:** C.S. and M.P. are funded by the Howard Hughes Medical Institute at the Janelia Research Campus. **Competing interests:** The authors declare no competing interests. **License information:** Copyright © 2024 the authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original US government works. https://www.science.org/ about/science-licenses-journal-article-reuse

SUPPLEMENTARY MATERIALS

science.org/doi/10.1126/science.adp7429 Movie S1

Submitted 8 June 2024; accepted 27 September 2024 10.1126/science.adp7429