






Topological insights into the neural basis of flexible behavior

Tevin C. Rouse^a, Amy M. Ni^b , Chengcheng Huang^{b,c} , and Marlene R. Cohen^{a,1} 

Edited by Doris Tsao, University of California, Berkeley, CA; received November 29, 2022; accepted March 28, 2023

It is widely accepted that there is an inextricable link between neural computations, biological mechanisms, and behavior, but it is challenging to simultaneously relate all three. Here, we show that topological data analysis (TDA) provides an important bridge between these approaches to studying how brains mediate behavior. We demonstrate that cognitive processes change the topological description of the shared activity of populations of visual neurons. These topological changes constrain and distinguish between competing mechanistic models, are connected to subjects' performance on a visual change detection task, and, via a link with network control theory, reveal a tradeoff between improving sensitivity to subtle visual stimulus changes and increasing the chance that the subject will stray off task. These connections provide a blueprint for using TDA to uncover the biological and computational mechanisms by which cognition affects behavior in health and disease.

neurophysiology | visual attention | computational neuroscience | topological data analysis

Perhaps the most remarkable hallmark of the nervous system is its flexibility. Cognitive processes including visual attention have long been known to affect both behavior (e.g., performance on visual tasks) and virtually every measure of neural activity in the visual cortex and beyond (1, 2). The diversity of changes associated with cognitive processes like attention makes it unsurprising that very simple, common measures of neural population activity provide limited accounts of how those neural changes affect behavior.

Arguably, the most promising simple link between sensory neurons and behavior is correlated variability [often quantified as noise or spike count correlations, or r_{SC} , which measure correlations between trial-to-trial fluctuations in the responses of a pair of neurons to repeated presentations of the same stimulus; (3)]. Correlated variability in visual cortex is related to the anatomical and functional relationships between neurons (3, 4). We demonstrated previously that the magnitude of correlated variability predicts performance on a difficult but simple visual task (Fig. 1*D*) across experimental sessions and on individual trials (5). This early success relating neural activity to simple behaviors means that correlated variability is a foundation on which to build efforts to explain more complex aspects of flexible behavior and the concomitant neural computations.

However, our efforts to relate correlated variability to a wider variety of sensory and cognitive phenomena and to constrain mechanistic models reveal a need for more sophisticated ways to characterize neuronal population activity. For example, although low correlations are associated with better performance in the case of attention and learning (5, 6), they are associated with worse performance when modulated by adaptation or contrast (7). Even in the case of cognitive processes like attention or task switching, good performance is associated with increases in correlation among particular subgroups of neurons (7, 8). And although mean correlated variability places much stronger constraints on cortical circuit models of cognition than measures of single-neuron responses (9, 10), these models remain underconstrained.

These results highlight the need to use holistic methods to investigate the relationship between noise correlations and behavior. We focused on topological data analysis, TDA; (11, 12), which is an emerging area in mathematics and data science that leverages groundbreaking advances in computational topology to summarize, visualize, and discriminate complex data based on topological data summaries. These approaches, which have mostly been used in fields like astrophysics or large scale neural measurements (e.g., refs. 13–15), are able to identify features in the data that are qualitatively distinct from those highlighted using traditional analytic methods.

1. Results

A. Topological Signatures of Correlated Variability. We used TDA specifically persistent homology; (16) to quantify the higher-order structure in the pairwise interactions between simultaneously recorded neurons from area V4 of rhesus monkeys performing

Significance

As the fields of systems, computational, and cognitive neuroscience strive to establish links between computations, biology, and behavior, there is an increasing need for an analysis framework to bridge levels of analysis. We demonstrate that topological data analysis (TDA) of the shared activity of populations of neurons is an effective link. TDA allows us to distinguish between competing mechanistic models and to answer longstanding questions in cognitive neuroscience, such as why there is a tradeoff between visual sensitivity and staying on task. These results and analysis framework have applications to many systems within neuroscience and beyond.

Author affiliations: ^aDivision of Biological Sciences, Department of Neurobiology, University of Chicago, Chicago, IL 60637; ^bDietrich School of Arts and Sciences, Department of Neuroscience, University of Pittsburgh, Pittsburgh, PA 15260; and ^cDepartment of Mathematics, University of Pittsburgh, Pittsburgh, PA 15260

Author contributions: T.C.R. and M.R.C. designed research; T.C.R., A.M.N., and C.H. performed research; T.C.R., A.M.N., C.H., and M.R.C. analyzed data; and T.C.R., A.M.N., C.H., and M.R.C. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2023 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed. Email: marlenecohen@uchicago.edu.

This article contains supporting information online at <http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2219557120/-DCSupplemental>.

Published June 6, 2023.

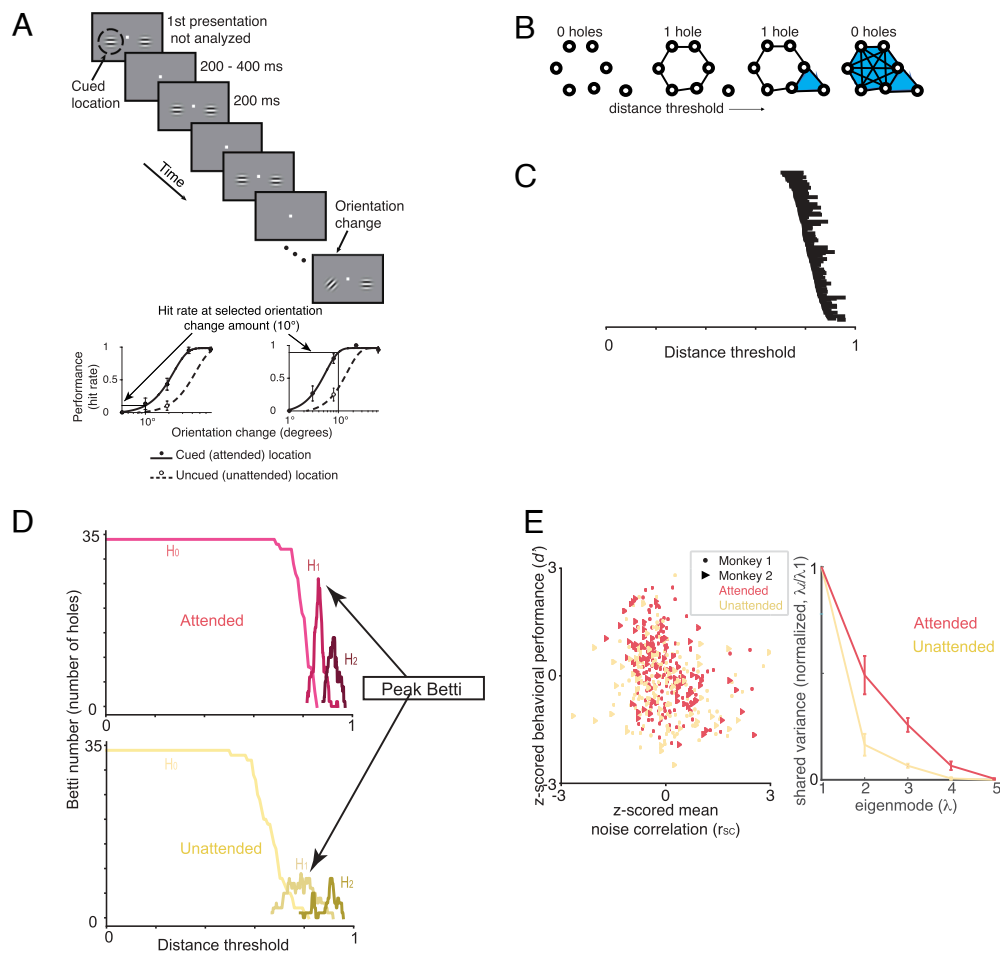


Fig. 1. Experimental and topological methods. (A) Orientation change detection task with cued attention (5). The lower panels are psychometric curves (hit rate as a function of orientation change amount) for two example recording sessions to illustrate how we calculated performance at one selected orientation change amount on every recording session. (B) Illustration of TDA methods. Each circle represents a neuron, and the distance between each pair is $1 - r_{SC}$, where r_{SC} is their noise correlation (note that in real networks, more than two dimensions are typically required to represent all of the pairwise interactions). This analysis method iterates through difference thresholds (going from small to large from *Left* to *Right*). When the difference between two points is less than the threshold, they are considered connected. The shaded regions indicate groups of points that are fully interconnected, which indicates a higher order interaction between that subgroup of neurons. We summarize this structure by counting the number of holes at each threshold [which constructs the Betti curve in (D)]. In this example, the holes are one-dimensional (topologically equivalent to a circle) as opposed to higher dimension (equivalent to a sphere) or connected components (equivalent to a point). For a more detailed description, see *SI Appendix* (TDA Example). (C) Topological description (persistence barcode) of an example recording session showing the difference thresholds (x axis) at which each one-dimensional hole (equivalent to a circle) exists (holes are ordered by the threshold at which they appear). Many datasets are characterized by a small number of persistent topological features, which would show up as long horizontal lines in this plot. Instead, our neural data are characterized by a large number of holes that persist only for a small range of difference thresholds (many short horizontal lines in this plot). (D) Example Betti curves [plots of the number of holes as a function of difference threshold, which corresponds to the number of lines present at each threshold in (C)] for an example recording session. H_0 (lightest color; 0th dimension) curve keeps track of the number of connected components ('holes' equivalent to points) in the graph as the threshold is varied. H_1 (middle color; 1st dimension) curve tracks the number of circular features with changes in the threshold. H_2 (darkest color; 2nd dimension) curve tracks the number of spherical features with changes in the threshold. Our analyses focus on the peak of the Betti curve for the 1st and 2nd dimensions (*SI Appendix* for other topological descriptions). The peak of the Betti curve for the 0th dimension is uninformative because it is always equal to the number of neurons. We include the curve for H_0 to illustrate attention-related difference its shape and because analyzing the Betti curve for H_0 , which is the most tractable to compute with limited data, may be useful for future neuroscience applications. (E) We focus our topological analyses of neural populations on correlated variability because there is a strong relationship between r_{SC} and performance on our psychophysical task (quantified as sensitivity or d^-). The plot shows d^- as a function of r_{SC} (both values are z-scored for each animal and computed from responses to the stimulus before the orientation change). The colors represent the trials when attention was directed inside ('attended', red) or outside the receptive fields of the recorded V4 neurons ('unattended', yellow). The correlation between d^- and r_{SC} was significant for each attention condition (attended: $r = -0.11$, $P = 0.11$; unattended $r = -0.25$, $P = 4.58e - 5$). (E) Factor analysis is a common linear method to assess the dimensionality of the correlated variability. The plot shows the shared variance (first five eigenvalues of the shared covariance matrix with private variance removed using Factor analysis) normalized by the shared variance in the first (dominant mode).

a difficult visual detection task with an attention cue (Fig. 1A; different aspects of these data have been presented previously; (5)). We analyzed the structure of noise correlations in a population of neurons in the visual cortex using a metric in which the difference measure between a pair of neurons is $1 - r_{SC}$, where r_{SC} is their noise correlation (Fig. 1B). Using this metric, highly correlated neurons have high similarity (low difference), and anticorrelated neurons have low similarity.

As is typical of the persistence homology approach, we iterate through a difference threshold (*Left* to *Right* in Fig. 1B) to understand the topological features of the correlation structure. For each threshold, we consider a pair of neurons to be functionally connected if their difference is less than the threshold. As the threshold increases, we thus include functional connections between pairs that are less correlated. For each difference threshold, we use established TDA methods to identify

“holes” in the correlation structure, which correspond to the organization of a lack of functional connections between a subset of neurons. The organization of existing and missing functional connections has implications for the organization and function of the network (17, 18).

We use TDA in a slightly different way than in most past work. The most common uses of TDA focus on persistent features holes that persist through a large range of difference thresholds; (19–22). For example, a beautiful recent study used TDA to analyze the structure of the signals represented by a population of neurons (22). Those authors focused on the persistent features of that dataset, which reflect the quantities encoded by that population of neurons. In contrast, we analyzed noise, which is not thought to have any particular structure (much less one characterized by holes of different dimensionality). In our datasets, we simply did not observe persistent features (Fig. 1C). Instead, we observed large numbers of holes that did not persist (23), and the number and difference threshold of those holes flexibly depended on attention and other cognitive processes. Our observations support the idea that there is information that can be found in features that do not persist (24).

We therefore summarize the topology of the correlation matrix as the peak Betti number, which is the maximal number of holes of a given type (called homology group) that appeared at any threshold, Fig. 1D; (23). We focus here on holes that are equivalent to circles (those detected by the first homology group) and spheres (detected by the second homology group), because these can be estimated using datasets of experimentally tractable size (16, 25, 26). For simplicity, we refer to these as circular and spherical features, respectively. In our data and models, focusing on the peak Betti curve led to qualitatively similar conclusions as other common topological summaries (similar conclusions in refs. 27 and 28; *SI Appendix*, Figs. S1 and S2).

Here, we demonstrate that topological descriptions of correlated variability are an effective bridge between behavioral, physiological, and theoretical approaches to studying neuronal populations. The peak Betti number is flexibly modulated by cognition, is related to performance on a visually guided task, and provides insights into mechanistic models and the function of real and artificial neural networks in different cognitive conditions.

B. Topology as a Bridge to Behavior. The primary reason for focusing on noise correlations is that the magnitude of noise correlations in visual cortex has been strongly linked to performance on visually guided tasks (5, 29). To justify our use of persistent homology to study neuronal networks, we tested the hypothesis that topological signatures of network activity capture key properties of the relationship between correlated variability and behavior.

Four observations suggest that the peak Betti number captures the aspects of noise correlations that are related to performance. First, across recording sessions, there was a negative relationship between the peak Betti number and the average noise correlation (Fig. 2C and D), meaning that sessions in which the average noise correlation was low tended to have a higher peak Betti number. Second, consistent with the observation that attention reduces noise correlations (2, 5, 30, 31), attention changes the peak Betti number (Fig. 2A and B). Third, the peak Betti number was higher on trials in which the animal correctly detected a change in a visual stimulus compared to trials in which the animal missed the stimulus change (Attended condition: average peak Betti number in H_1 for correct trials: 14.47, incorrect trials: 13.46; in H_2 for

correct trials: 8.38, incorrect trials: 7.41; paired t -test (peak Betti number H_1) $P = 0.014$, (peak Betti number H_2) $P = 0.015$). Finally, there was a positive correlation between the peak Betti number and behavioral performance (Fig. 2E and F). Together, these results show that peak Betti number is a good description of the aspects of correlated variability that correspond to changes in behavior.

C. Topology as a Bridge to Mechanism. The magnitude of correlated variability places strong constraints on circuit models of the neuronal mechanisms underlying attention, Fig. 3A; (9, 10). In particular, network models are constrained by the observation that attention changes correlated variability in essentially a single dimension of neuronal population space in area V4, Fig. 2; (7, 10, 32, 33).

Topological descriptions of simulated networks can distinguish between competing models in situations when the magnitude of shared variability, even in the most relevant dimension, fails to do so. We analyzed the outputs of our spatially extended network of spiking neuron models, which internally generate correlated variability through spatiotemporal dynamics (10). In the model, the magnitude of correlated variability can be changed by modulating inhibition in two distinct ways: either increasing the input drive to the inhibitory neurons (μ_i in Fig. 3B) or decreasing the timescale of inhibition (τ_i in Fig. 3B) changes correlated variability in a low rank way.

These two mechanisms have very different effects on the topology of the correlated variability, even when the mean variability is equivalent. For most parameter values, changing the input drive to the inhibitory neurons has a much greater effect on the peak Betti number than changing the timescale (Fig. 3C–E). While changing the timescale of inhibition is extremely common in circuit models (for review, ref. 10), in real neural networks, the timescale of inhibition is longer than excitation and is inflexible (34–36). Both the biology and the topology are consistent with the idea that attention instead acts by increasing the input drive to the inhibitory neurons (9, 10).

These results demonstrate that topological signatures of correlated variability provide constraints on mechanistic models that are unavailable using linear measures of neural activity. Changes in the mean or dimensionality of correlated variability are not necessarily coupled with changes in the topological signatures of the network. Together, our results highlight the value of using circuit models as a platform on which to test and generate hypothesized mechanisms underlying perception and cognition.

D. Topology as a Bridge to Network Function. The past two decades have seen an explosion in the number of studies demonstrating that correlated variability depends on a wide range of sensory, cognitive, and motor conditions that change behavior (for review, ref. 2). Despite much effort from the experimental and theoretical neuroscience communities (9, 10, 32, 33, 37–39), how changes in correlated variability might affect behavior remains unclear. The observations that topological summaries of the noise correlation matrix are related to behavior suggest that, via known connections to network control theory (40–42), TDA can provide insight into the relationship between correlated variability, the function of a network, and behavior.

TDA has known connections to network control theory because both measure the structure (or lack thereof) in a functional connectivity matrix (42). We reasoned that network control theory, which seeks to quantify the ability of interventions (in our

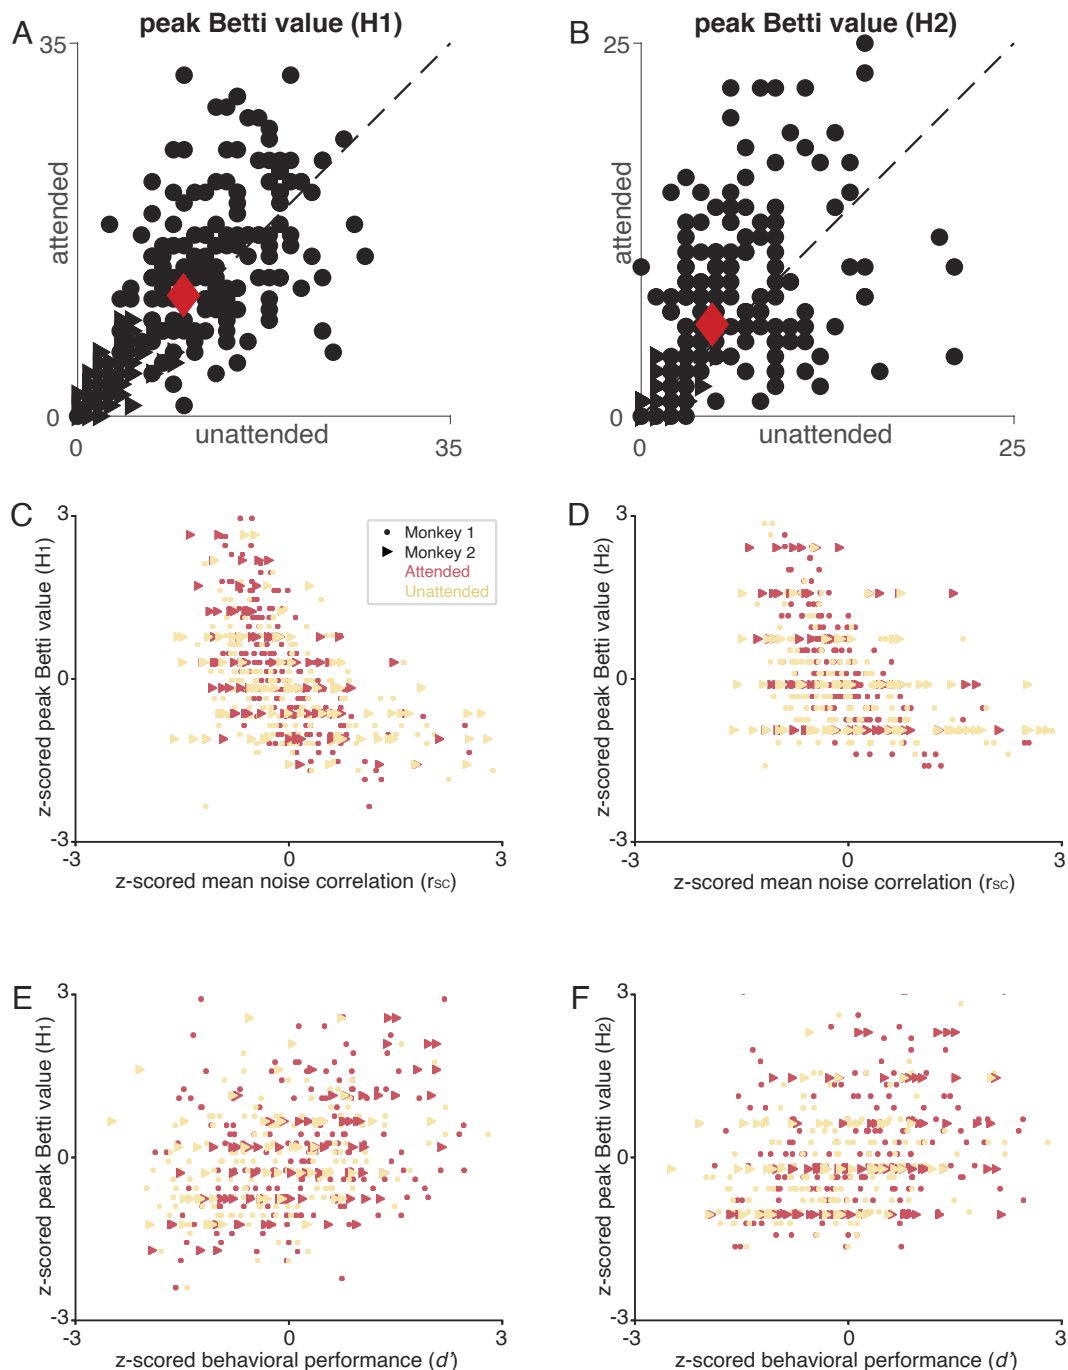


Fig. 2. TDA reveals the relationship between neurons and behavior. (A, B) Attention increases the maximum number of circular features (i.e., 1st dimension features; A) and spherical features (i.e., 2nd dimension features; B) over the range of threshold values. Each point represents one experimental session, and the red points are the mean values, which are significantly greater for the attended condition (y axes) than the unattended condition (x axes; paired t -tests, $P < 0.01$). (C) The maximum number of circular features is correlated with the mean r_{sc} in both attention conditions both are z-scored for each animal; (Attended: $r = -0.52$, $P = 1.3e - 18$; unattended: $r = -0.42$, $P = 7.9e - 12$; paired t -test (attended and unattended, peak of the Betti curve): $P = 8.41e - 5$). (D) Same, for the maximum number of spherical features 2nd homology group (attended: $r = -0.44$, $P = 4.9e - 13$; unattended: $r = -0.36$, $P = 5.3e - 9$; paired t -test (attended and unattended, peak of the Betti curve): $P = 3.09e - 6$). (E) There is a strong relationship between the maximum number of circular features and behavioral performance (d^- or sensitivity calculated for a single orientation change for each session; both measures are z-scored; Attended: $r = 0.32$, $P = 2.34e - 7$; Unattended: $r = 0.3$, $P = 1.41e - 5$). (F) Same, for the maximum number of spherical features (Attended: $r = 0.27$, $P = 2.06e - 5$; Unattended: $r = -0.28$, $P = 4.67e - 5$).

case, visual stimuli, cognitive processes, or random fluctuations) to alter the state of a network (43) could provide intuition about the relationship between TDA and the function of our neuronal network. While network control theory is primarily used in engineering, recent work has used controllability to quantify the flexibility of large neural systems, constrained by fMRI data (44).

These methods focus on quantifying the energy required to move between states of the neural population. We define a state as the vector of neural population activity on a given trial, and the energy required to move between states is constrained by the noise correlation matrix (e.g., in Fig. 4A). For example, if the responses of all the neurons are highly positively correlated, then reaching a state in which the response of some is high while

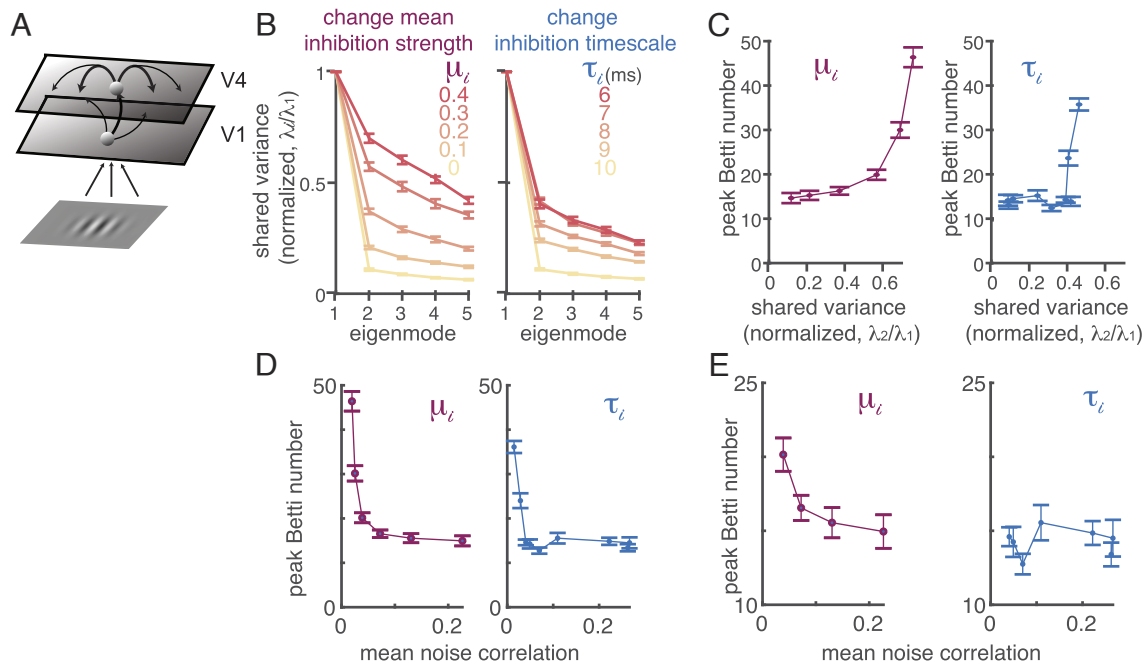


Fig. 3. TDA can distinguish between different mechanistic models of attention. (A) Model schematic of a two-layer network of spatially ordered spiking neurons modeling primary visual cortex (V1) and area V4, respectively. The visual inputs to the model are the same Gabor stimuli used in our experiments. (B) Two distinct attention mechanisms can decrease correlated variability in a low rank way that is similar to linear descriptions of our data. We can reduce correlations either by increasing the currents to all inhibitory neurons (μ_i) or decreasing the decay timescale of inhibitory currents (τ_i^d). The plots depict the shared variance in each mode (the top five eigenvalues from the shared covariance, with private variance removed using Factor analysis), normalized by the shared variance in the first mode, for different values of μ_i (left) or τ_i^d (right; error bars represent SEM). The two mechanisms appear indistinguishable using linear methods. (C) The two mechanisms cause different changes in the topological descriptions of the modeled V4 populations. As μ_i increases (Left panel), so does the shared variance present outside the first mode (x axis) as well as the peak Betti number (shown for the circular (i.e., 1st dimension) features in the y axis). Changes in τ_i (Right panel) result in a different relationship between the peak Betti number and the shared variance in higher dimensions, affecting the peak Betti number only at very short (unrealistic) timescales (those with the greatest shared variance; red lines in B). The peak Betti number is computed from the same simulated responses as in (B). Error bars represent SEM. (D) The peak Betti number (y axis) has a different relationship with average noise correlation (x axis) when modulated by changing the mean current to the inhibitory neurons (μ_i , Left panel) or decreasing the decay timescale (τ_i^d , Right panel). Error bars represent SEM. (E) Same as (D), except zoomed in to exclude parameter values that result in an unrealistically low mean noise correlation (<0.03). In this physiologically realistic range, changing μ_i is associated with monotonic changes to the peak Betti number, while changing τ_i^d does not appreciably change peak Betti number.

the others are low is unlikely and therefore requires significant energy.

If our starting point (e.g., the horizontal Gabor in Fig. 4B) is the population response to a horizontal Gabor stimulus presented before the orientation change in our task (see Fig. 1A), a nearby state might be the population response to the changed stimulus (e.g., the oblique Gabor in Fig. 4B). A distant state might be a population response when the monkey is concentrating on something very different and task-irrelevant (e.g., thinking about the banana in Fig. 4B). Average controllability quantifies how readily the population moves from the starting point to nearby states while modal controllability quantifies how readily the population moves to distant states.

There is no mathematical relationship between average and modal controllability. Indeed, average and modal controllability were uncorrelated across sessions in our data ($R = 0.008$; $P = 0.9$).

However, we found that the topological descriptions of neuronal population are strongly related to both average and modal controllability, and both are related to attention. High peak Betti value (which occurs more readily in the attended state) is associated with decreases in the energy required to drive the system to nearby states (high-average controllability; Fig. 4C and E). In contrast, there is a negative relationship between the peak Betti value and the energy required to drive the system to distant states (modal controllability; Fig. 4D and F) but changing attention conditions increases modal controllability (compare the red and yellow points in Fig. 4D and F).

The different relationships between topology and average and modal controllability observations provide insight into the tradeoffs associated with attention. Attending to a stimulus improves the network's ability to respond to subtle interventions, which is consistent with the attention-related improvements in the animal's ability to detect a subtle change in the visual stimulus (Fig. 1A), but it has complex effects on the ability of the network to change states dramatically, which may mean that attention reduces cognitive flexibility. In future work, it would be interesting to study whether changes in controllability can account for change blindness and other behavioral demonstrations that attention reduces the ability of observers to notice very unexpected stimuli, such as the classic example of failing to notice a gorilla walking through a basketball game (45).

Indeed, average and modal controllability have distinct relationships with the monkeys' performance in our task. We sorted the experimental sessions by average controllability (colors in Fig. 4G) or modal controllability (colors in Fig. 4H). Increased average controllability was associated with improvements in the monkeys' ability to detect all orientation change amounts (except the smallest changes in which a floor effect meant that they were rarely detected). In contrast, modal controllability was unrelated to the monkeys' ability to detect subtle orientation changes and was anticorrelated with the ability to detect large, easy orientation changes. One interpretation is that when modal controllability is high, the monkeys' minds wander more easily to distant, potentially task-irrelevant states, increasing the lapse rate on easy

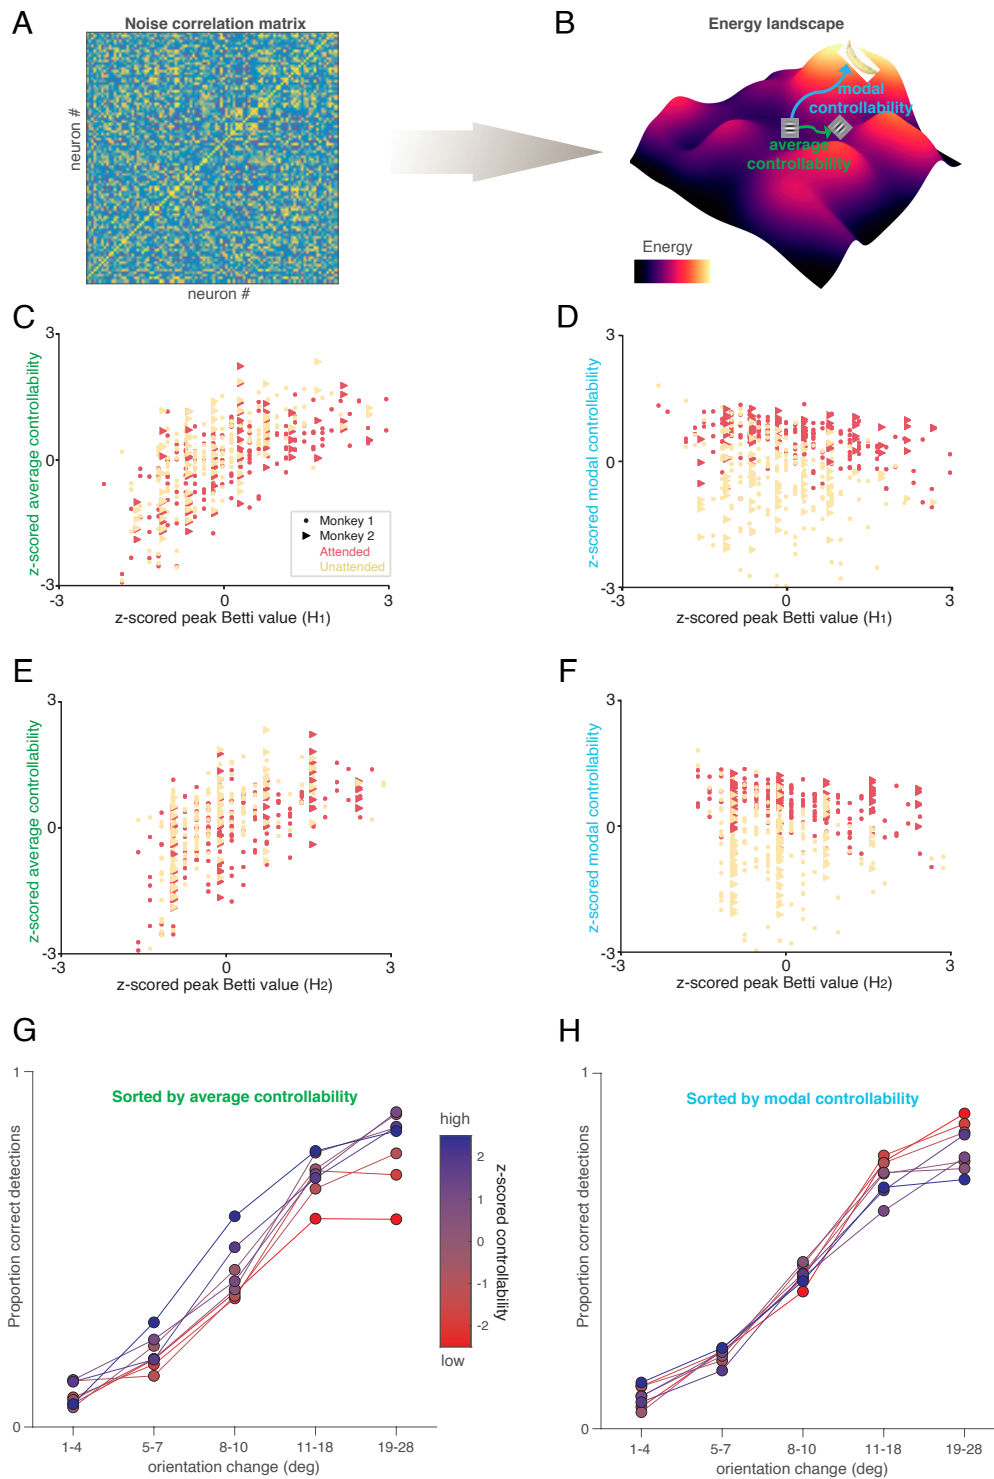


Fig. 4. TDA and controllability provide insight into network function. (A and B) Illustration of our controllability calculation. We consider the noise correlation matrix (A) as a functional connectivity matrix, and use this to calculate an energy landscape (illustrated for a hypothetical situation in B; colors indicate energy). Average controllability is defined as the energy required to move from a starting point (e.g., a response to a horizontal Gabor stimulus) to nearby states (e.g., a response to an oblique Gabor), and modal controllability is defined as the energy required to move to distant states (e.g., thinking about a banana). (C) High-average controllability is associated with maximum number of circular (i.e., 1st dimension) features. Both measures were z-scored for each animal, and the lines were fit for each attention condition; attended: $r = 0.65$, $P = 1.03e - 31$; unattended: $r = 0.68$, $P = 1.28e - 33$; paired t -test (attended and unattended, average controllability): $P = 7.5e - 61$. (D) High modal controllability is associated with lower number of circular (i.e., 1st dimension) features (Conventions as in A; attended: $r = -0.37$, $P = 7.7e - 10$; unattended: $r = -0.23$, $P = 3.2e - 4$; paired t -test (attended and unattended, modal controllability): $P = 1.5e - 38$). (E) Relationship between average controllability and maximum number of spherical features (i.e., 2nd dimension) (attended: $r = 0.59$, $P = 3.5e - 25$; unattended: $r = 0.59$, $P = 9.5e - 24$). Conventions as in (C). (F). Relationship between modal controllability and number of spherical (2nd dimension) features (attended: $r = -0.38$, $P = 3.12e - 10$; unattended: $r = -0.18$, $P = 4.7e - 3$). Conventions as in (D). (G) High-average controllability is associated with better performance at all orientation change amounts. Colors represent z-scored average controllability (the experimental sessions were split into six equally sized bins by average controllability), and the plot shows proportion correct detections (hit rate) as a function of orientation change amount. (H) High modal controllability (bluer colors) is associated with a worse lapse rate (worse performance on easier trials). Conventions as in (H).

trials. Together, these results demonstrate that in addition to linking to behavior and mechanism, topological signatures of the structure of noise correlations can provide insight into the function of the network and the behavioral trade-offs associated with changes in correlated variability.

E. Comparing Topological Summaries and Mean Noise Correlations. The results in Figs. 1–4 demonstrate that topological summaries of noise correlation matrices are related to many quantities of interest, including behavior, average and modal controllability, and the mean pairwise noise correlation. If all of these quantities are related to each other, what is the added value of TDA over the simpler and more common mean noise correlation metric?

Others have written about the value of TDA for many applications, including analyzing signals from groups of neurons as opposed to noise as we have done here; 22. Our results suggest two key advantages of TDA over mean noise correlations for understanding the mechanisms by which populations of neurons guide behavior. First, we demonstrated that TDA can distinguish between models in which the mean noise correlations (and even the dimensionality of the noise correlation matrix) were indistinguishable (Fig. 3).

Second, we found that the relationships between topological summaries and other quantities of interest (including performance on our change detection task, average controllability, and modal controllability) are stronger and dissociable from the relationships between those quantities and mean noise correlations (Fig. 5). To establish this, we first computed the raw session-to-session correlation coefficient between each quantity of interest (Fig. 5A and D, which are identical except that Fig. 5A contains the peak Betti value in the first dimension as a representative topological summary and Fig. 5D contains mean noise correlations). Next, we computed partial correlations between those same quantities while controlling for the effect of

mean noise correlation (Fig. 5B) or the peak Betti value (Fig. 5E). The difference between the raw and partial correlations reflects the extent to which, for example, the relationship between behavior and average controllability, can be attributed to the relationship between each of those and mean noise correlation (Fig. 5C) or the peak Betti value (Fig. 5F).

Controlling for the peak Betti value had a bigger impact on every pairwise relationship than controlling for mean noise correlation (and this difference was statistically significant for all three common off diagonals in Fig. 5C and F; $P < 0.05$ with a Bonferroni). This result indicates that statistically speaking, the peak Betti value provides more and independent insight into behavioral and control theory measures of circuit function than mean noise correlation.

It is worth noting that in all likelihood, none of the measures discussed here (e.g., from TDA, control theory, or other descriptions of the noise correlation matrix) are quantities that are directly used by the brain in neural computations. Neural computations are performed to guide behavior at individual moments or on individual trials, and noise correlations, or any derivative of them, are computed over many ostensibly identical trials. The value of any of these metrics is that they provide insight into the underlying computations. Our results (especially in Fig. 5) demonstrate that TDA provides insight into some key quantities (perhaps most importantly into behavior) that are distinct from the insights that can be gleaned from mean noise correlations alone.

2. Discussion

A. Implications for Topology. Although TDA has been used for many scientific applications (18), our use of TDA differs from most previous work. The prevailing paradigm used in most TDA applications, including many in neuroscience (13, 17), focuses on identifying persistent topological features, such as holes that




Fig. 5. A partial correlation analysis reveals that topological summaries of noise correlation matrices provide insights that are distinct from the mean noise correlation. (A) Raw correlation coefficients summarizing the session-to-session correspondence between modal controllability, average controllability, behavior (defined as perceptual sensitivity as in Fig. 2), and the peak Betti value (maximum number of circular features). The relationships between the controllability and behavioral metrics and other topological descriptions (including the maximum number of spherical features and total persistent) are qualitatively similar, so for simplicity, we focus on the maximum number of circular features here. Diagonals represent self-correlations, and are therefore 1 by definition. (B) Same as A, but the colors reflect partial correlations that control for session-to-session-variability in mean noise correlation. (C) Difference between A and B, showing that the raw and partial correlations are qualitatively similar. (D) Same as A, but the last row is mean noise correlation instead of the peak Betti value. (E) Partial correlations controlling for session-to-session variability in peak Betti value. (F) Difference between D and E.

persist across many thresholds (Fig. 1*B*). These persistent features are appealing because they can reveal the structure of a simple network. However, applying these methods to analyze neural circuits may not lead to any scientific discoveries, since persistent features are not expected in neural response variability, which is thought to arise from complex network properties that make it relatively unstructured (4).

However, we demonstrated that using TDA to analyze correlated variability in neuronal responses is useful, even in the absence of persistent features. The link that we demonstrated between the topology of noise correlations, which have been shown to reflect both cognition and the anatomy of the system (3), and the controllability of the network on individual trials (which are what matter for guiding behaviors) therefore has implications far beyond neuroscience. Throughout the natural and physical sciences, natural systems are complex and call for sophisticated data analytics pipelines. In astronomy, for example, TDA has been used to understand the relationship between planets, stars, and galaxies on a huge range of spatial scales. Our use of TDA to analyze nonpersistent topological features (ref. 23) will be a bridge between neuroscience and other fields. These tools for analyzing and interpreting complex networks can be deployed in many other scientific domains.

B. Implications for Neuroscience. We demonstrated here that using TDA to analyze the variability neural populations can illuminate interesting links between behavior, neurons, computations, and mechanisms. This sort of bridge between different levels of investigation has the potential to be broadly transformative. In an age of massive improvements in experimental technologies and tools for measuring the activity of large numbers of neurons, perhaps the greatest barrier to success understanding the neural basis of behavior is that it is different to compare and integrate results from experiments using different methods in different model systems. TDA can reveal relationships between neural networks, computations, and behaviors that are robust to the differences in neuronal responses that occur between every different experimental system (46). These analytical links make it possible to leverage the complementary strengths of each approach.

A holistic view of neuronal populations is necessary for understanding any neural computation. Essentially every normal behavioral process or disorder of the nervous system is thought to involve the coordinated activity of large groups of neurons spanning many brain areas. Tools for understanding and interpreting large populations have lagged far behind tools for measuring their activity. Standard linear methods have provided a limited view, and the field is in dire need of a new, holistic window into population activity. Our results demonstrate a hopeful future for using the topology of neural networks to fulfill that need.

3. Materials and Methods

C. Experimental Methods. Different analyses of these data have been presented previously (5). Briefly, two adult rhesus monkeys (*Macaca mulatta*) performed an orientation change detection task with a spatial attention component (30). The monkeys fixated a central spot while two peripheral Gabor stimuli flashed on (for 200 ms) and off (for a randomized period between 200 and 400 ms). At a random and un signaled time, the orientation of one stimulus changed, and the monkey received a liquid reward for making an eye movement to the changed stimulus within 500 ms. We cued attention in blocks of 125 trials, and the orientation change occurred at the cued location on 80% of trials. Our analyses are based on responses to the stimulus presentation before

the change, which was the same on every trial within a recording session. The location, contrast, and spatial frequency of the Gabor stimuli were the same during every recording session, but the orientation differed across sessions. The location of one stimulus was within the receptive fields of the recorded neurons and the other stimulus was in the opposite hemifield.

We presented the stimuli on a CRT monitor (calibrated to linearize intensity; 120-Hz refreshed rate) placed 52 cm from the monkey. We monitored the animals' eye position using an infrared eye tracker (Eyelink 1000; SR Research) and recorded eye position, neuronal responses (30,000 samples/s) using Ripple Hardware.

While the monkey performed the task, we recorded simultaneously from a chronically implanted 96-channel microelectrode array (Blackrock Microsystems) in the left hemisphere of visual area V4. We include both single units and sorted multiunit clusters (mean 34 and 15 units for Monkeys 1 and 2, respectively). The average number of simultaneously recorded pairs of units (for computing noise correlations) was 561 for Monkey 1 and 105 for Monkey 2. The data presented are from 42 recording sessions from Monkey 1 and 28 recording sessions from Monkey 2.

D. Data Preparation. To examine how the topology of networks of neurons in visual cortex or outputs of spiking models depend on attention, we constructed difference matrices from noise correlation matrices (3). We defined the noise correlation for each pair of neurons [also known as spike count correlation; (3)] as the correlation coefficient between the spike count responses of the two neurons in response to repeated presentations of the same stimulus. We based our analyses on spike count responses between 60 and 260 ms after the onset of the visual stimulus to allow for the latency of visual responses in area V4. We used responses to the stimulus before the orientation change because those are the same on every trial. We focused on trials when the monkey correctly identified the changed stimulus and compared responses in the two attention conditions.

Many measures of neuronal activity depend on experimental details like the number of recorded neurons or their mean firing rates, which were different for the two monkeys (ref. 5 for details). To allow us to combine across animals, we z-scored the results for each animal (across both attention conditions) and plot those normalized measures in the figures.

E. Behavioral Measures. To analyze the relationship between neuronal responses and behavior, we adopted a signal detection framework (47, 48) to assess how behavior depends on neurons and attention (49–55). Criterion is defined as

$$c = -\frac{1}{2} [\Phi^{-1}(\text{Hit Rate}) + \Phi^{-1}(\text{False Alarm Rate})], \quad [1]$$

where Φ^{-1} is the inverse normal cumulative distribution function. Negative values of c indicate that the subject has a liberal criterion (bias toward reporting changes), and positive values indicate a conservative criterion (bias toward reporting nonchanges).

Sensitivity is defined as

$$d' = \Phi^{-1}(\text{Hit Rate}) - \Phi^{-1}(\text{False Alarm Rate}). \quad [2]$$

Larger values of d' indicate better perceptual sensitivity.

Different orientation change amounts were used in different recording sessions. To compare across sessions, we fit the psychometric curve using a Weibull function and computed performance at a single, fixed orientation change for each session, Fig. 1*A* and (5).

F. Topological Measures. We examine the topology of the noise correlations in each attention condition using a Vietoris Rips construction. This consists of defining a difference matrix (better understood as a weighted adjacency matrix), which we constructed from the noise correlation matrix r_{SC} , to define pairwise (and higher order) connections between the vertices (representing neurons) in the simplicial complex. The difference metric was chosen to be $1 - r_{SC}$, so that

higher weighted interactions (i.e., those neurons which are strongly correlated) are defined as more similar and therefore entered the simplicial complex first.

We consider a difference threshold which defines those pairwise interactions that are permitted to be considered in the simplicial complex. Such a process allows us to examine the evolution of the simplicial complex across different difference thresholds. We assess several properties of the simplicial complex at each threshold value, including the existence of holes (or higher dimensional voids) of a given dimensionality (termed homology dimension). A hole signifies a lack of connections (i.e., differences in the degree of correlation) between a subset of neurons at the current difference threshold. We focus our analysis on the first and second homology dimensions (which correspond to holes that are topologically equivalent to circles and spheres, respectively) because they can be estimated reliably given the size of our datasets.

In most applications of TDA, researchers focus on persistent features that imply a nontrivial structure. For example, imagine a set of people seated around an oval-shaped table. We could use as a difference metric the physical distance between them, and we would consider two people to be "connected" if they are sitting closer than some distance threshold. At very small distance thresholds, no pair of people would be close enough to be considered connected. At very large thresholds (e.g., longer than the length of the room), all pairs of people would be connected. But for a large range of intermediate distance thresholds, each person would be connected to at least their nearest neighbor but would not be connected to everyone else, and the resulting graph would have a "hole", corresponding to the center of the table. In TDA, this hole would represent a persistent feature and would indicate that the seating arrangement has a particular structure. The presence of a single, persistent hole in the first homology dimension (equivalent to a circle), would imply that the people had arranged themselves around a table, but it would not specify whether that table was a circle, a rectangle, or another topologically equivalent shape.

Next, consider a situation where the same people simply sat in a haphazard arrangement on the floor. By chance, there would at some distance thresholds be holes around which, for example, a subgroup of people were arranged. Those chance holes would not persist for a very long range of distance thresholds. An intermediate situation, in which there is some structure to the seating arrangement (e.g., people sitting in small clusters), might have a smaller number of holes that persist for only a small range of thresholds.

Some recent studies have used TDA to investigate the signals encoded by populations of neurons or brain areas (19–22). In these studies, the vertices represent trials, stimuli, or time periods, and the authors construct a distance metric to relate population responses at those different times. Because neural signals have structure, those authors were able to analyze features that persist over a long range of distance thresholds.

Our approach was orthogonal. We took each neuron to be a vertex, and the difference between them was given by the pairwise noise correlation. Although they are typically low rank, noise correlation matrices are in general thought to be unstructured within that small number of dimensions (3, 10). For this reason, most previous studies focus on their mean or linear dimensionality (10). Consistent with the idea of low rank but unstructured correlation structures, we did not observe notable persistent features (Fig. 1C for an example of observed features). Therefore, the 'holes' in our data should be thought of as topological noise. The relationship between this topological noise and other quantities of interest (e.g., behavior) indicates that although no individual hole is particularly important, the distribution of them can provide insight into neurobiological processes.

We therefore adopted the approach of focusing on the distribution of topological features rather than on looking for long persistent cycles (24). We examined how properties of the generated Betti curves (such as the peak or total persistence) relate to common measures of attention like average noise correlations and behavioral performance.

G. Topological Data Analysis Example. We provide here a detailed walk-through of the schematic in Fig. 1B. Suppose that we have a group of vertices. You can think of these vertices as a collection of neurons where each vertex is a

neuron. Indeed, this is the view that we take in this work. Along with this group of vertices, we have an underlying weight matrix that expresses not only what vertices are connected to one another but the strength of these connections. A neuroscience interpretation of this weight matrix is the connectivity matrix of a population of neurons. If we look at a pair of vertices, say n_1 and n_2 and observe an entry of 0.5 in the weight matrix, then we know that not only are n_1 and n_2 connected but the strength of their connection is 0.5.

With both a group of vertices and a weight matrix expressing how the vertices are connected, we can now apply topological methods. These methods will allow us to examine relationships between subgroups of vertices. We define a threshold value that will determine which connections in the weight matrix are allowed. Allowed connections are those whose value is at most the threshold value. However, given that the "optimal" threshold value is unknown, the typical approach is to vary the threshold value over a range while simultaneously tracking the properties of the evolving graph. We choose to track the number of holes (i.e., an empty space in the graph due to a lack of connections). As we continuously increase the threshold, certain connections will enter the graph and the connections that are lacking may form holes within the graph.

We examine this process in the context provided in Fig. 1B. Initially there are seven disconnected vertices. We can assume that this is the case because the difference threshold is zero at the beginning. As we increase the difference threshold, we have one isolated vertex and a ring structure. Given that there is a lack of connections between different subsets of these vertices, a hole is formed. Thus, the number of holes for our structure (i.e., referred to as the Betti Number in the literature) is one. As we increase the difference threshold further, the isolated vertex becomes connected to two vertices and we signify this all-to-all connected subgraph by a shaded region. Observe that although the difference threshold has increased, the hole in the center of the ring is still present, and thus, our Betti Number is still one. Finally, we increase the threshold, and all vertices in the ring become connected. Our Betti number is now zero. This has occurred because the weights of those connections are at most the value of the difference threshold. We apply this same process to our neural data where the weight matrix is one minus the noise correlation value between a pair of neurons in the overall noise correlation matrix.

H. Spatial Model Construction. Our spatial model is a variation of the two-layer network of neurons discussed in ref. 10. Neurons in this network are arranged uniformly on a $[0, 1] \times [0, 1]$ grid. The first layer (i.e., the feed-forward layer) consists of $N_x = 2,500$ excitatory neurons that behave as independent Poisson processes. The second layer consists of 40,000 excitatory and 10,000 inhibitory neurons that are recurrently coupled. The second layer receives input from the first layer. The network's connectivity is probabilistic but dependent on a Gaussian of width σ_* . Thus, neurons that are further away from each other on the grid are less likely to connect.

The parameters are the same as the two-layer network in Huang et al. (10), and are chosen to approximate known biology of cortical circuits (ref. 10 for details). Specifically, the synaptic strengths are scaled by $1/\sqrt{N}$, where N is the total number of neurons in the network, as used in the so-called balanced networks (56) such that the recurrent network can internally generate variability in neural spiking for large N . The projection widths of the excitatory and inhibitory neurons are chosen to be the same ($\sigma_e = \sigma_i = 0.1$), which is consistent with anatomical findings from visual cortex (57, 58). Each neuron is modeled as an exponential integrate-and-fire neuron model, following standard formulation in past work (59).

The only differences between the published model and one here are the following. The feedforward connection strength from layer 1 to layer 2 is $J_{ex} = 1.0$ and $J_{ix} = 0$ for excitatory and inhibitory neurons, respectively. Fig. 3 B, Left: $\mu_j = 0, 0.1, 0.2, 0.3, 0.4$ and τ_j is 10. Fig. 3 B, Right: $\mu_j = 0$, and $\tau_j = 6, 7, 8, 9, 10$ ms. These parameters were chosen so that the manipulations of μ_j and τ_j begin with the same parameter set at high correlation value, which is unstable with turbulent wave dynamics. There were a total of 15 simulations of 20 s each for each parameter conditions. The first 1 s in each simulation was removed. The spike counts were computed using 140-ms time window to mimic the data.

We implemented this model using EIF neurons. The voltage dynamics of these neurons are governed by the following equation (10):

$$C_m \frac{dV_j^\alpha}{dt} = -g_L(V_j^\alpha - E_L) + g_L \Delta_T e^{\frac{(V_j^\alpha - V_T)}{\Delta_T}} + I_j^\alpha(t), \quad [3]$$

where $\tau_m = \frac{C_m}{g_L} = 15$ ms, $E_L = -60$ mV, $V_T = -50$ mV, $V_{th} = -10$ mV, $\Delta_T = 2$ mV, $V_{re} = -65$ mV, $\tau_{ref} = 1.5$ ms and the total current $I_j^\alpha(t)$ obeys the following equation:

$$\frac{I_j^\alpha(t)}{C_m} = \sum_{k=1}^{N_f} \frac{J^{\alpha F}}{\sqrt{N}} \sum_n \eta_F(t - t_n^{F,k}) + \sum_{\beta=e,i} \sum_{k=1}^{N_\beta} \frac{J^{jk}}{\sqrt{N}} \sum_n \eta_\beta(t - t_n^{\beta,k}) + \mu_\alpha, \quad [4]$$

where N is the total number of neurons in the second layer and μ_α is the static current to the $\alpha (\in \{E, I\})$ population. η_β is the postsynaptic current given by the following equation:

$$\eta_\beta(t) = \frac{1}{\tau_{\beta d} - \tau_{\beta r}} \begin{cases} e^{-\frac{t}{\tau_{\beta d}}} - e^{-\frac{t}{\tau_{\beta r}}}, & t \geq 0. \\ 0, & t < 0 \end{cases} \quad [5]$$

where the rise time constant $\tau_{\beta r} = 5$. We consider multiple values of the decay time constant $\tau_{\beta d}$. For both dimensionality (*Factor Analysis*) and topological (*Topological Measures*) comparisons, we also considered a range of values of the μ_I parameter, which correspond to the overall depolarization of the inhibitory population and which has been shown to affect the dimensionality of the generated data.

I. Factor Analysis. To assess the dimensionality of the population simulated using the spatial model with different parameters, we used factor analysis (39). We based our analysis on a number of neurons by the number of trial matrix of spike counts of the simulated excitatory neurons. We then used that matrix to compute a spike count covariance matrix. Factor analysis separates the spike count covariance matrix into a shared component that represents how neurons covary together and an independent component that captures neuron-specific variance. Following the notation of ref. 39, use L to refer to the loading matrix relating m latent variables to the matrix of neural activity. In this way, the rank, m , of the shared component LL^T is the number of latent variables that describes the covariance. We refer to the independent component as Ψ , which is a diagonal matrix of independent variances for each neuron. We then assess the dimensionality of the network activity by analyzing the eigenvalues of LL^T . To focus on dimensionality rather than the total amount of independent or shared variance (which depends on many model parameters), we normalized each eigenvalue by dividing by the largest eigenvalue.

We performed this analysis on the spike count responses of a randomly sampled 500 simulated neurons. All analyses were cross-validated. Error bars in the figures come from analyzing many instances of the network generated using fixed model parameters.

J. Controllability Measures. The goal of our controllability measures is to understand how the noise correlation matrix in each attention state constrains

estimates of the function of the network. We consider a hypothetical (possibly nonlinear) dynamical system whose dynamics can be linearized and whose effective connectivity is defined by the noise correlation matrix. We analyze the properties of the system to assess the amount of effort it takes to change the system's state using external input. We summarize these calculations using two standard measures of controllability (43, 60): average controllability, which relates to the ability to push the system into nearby states or states with little energy and modal controllability, which relates to the ability to push the system into distant states or states that require more energy.

We take the effective connectivity matrix A to be the noise correlation matrix generated from the spike count responses to repeated presentations of the same visual stimulus as described above. To align with the controllability methods in previous cognitive neuroscience studies (e.g., refs. 60–63) and to remove the influence of self-connections (which are defined as 1 for a correlation matrix), we set the diagonals of the effective connectivity matrix A to 0. (Leaving the diagonals as 1 did not qualitatively change our results).

We then consider the energy required to steer the network from an initial state x_0 to a target state $x(T) = x_T$.

The average controllability ρ_c is defined as

$$\rho_c = \text{Trace}(W_c), \quad [6]$$

where $W_c = \int_0^{t_f} e^{A\tau} B B^T e^{A^T \tau} d\tau$ is the controllability Gramian matrix in which B represents a matrix of nodes (neurons) in which we could inject hypothetical inputs to change the network state (the full matrix in our case) and τ represents the fact that the input could be in principle time varying. $\text{Trace}()$ is the trace of the matrix (i.e., the sum of the diagonal elements of the matrix).

The modal controllability ϕ_i is defined as

$$\phi_i = \sum_{j=1}^N = (1 - e^{\lambda_j(A)}) v_{ij}^2, \quad [7]$$

where λ_j is the j th eigenvalue of the effective connectivity matrix A . v_{ij} corresponds to the eigenvectors of the effective connectivity matrix A .

We used the equations given above to compute the average and modal controllability of the recorded population of neurons. To determine the relationship between controllability and behavior (Fig. 4 G and H), we computed average and modal controllability for each session, z-scored those measures for each monkey, and divided the sessions into six equally sized bins for each controllability measure.

Data, Materials, and Software Availability. Electrophysiology and behavior data have been deposited in Open Science Framework (<https://doi.org/10.17605/OSF.IO/RN7TU> (64); <https://github.com/TCR23/TopologyForAttention.git> (65)). Previously published data were used for this work (5).

ACKNOWLEDGMENTS. We thank Douglas Ruff, Joshua Alberts, and Jen Symmonds for contributions to the previously published experiments and Alessandro Rinaldo and Ralph Cohen for helpful conversations. This work was supported by NIH grants 2R01EY022930 (MRC, TCR), RF1NS121913 (MRC, CH), R01EY034723 (MRC), and K99NS118117 (AMN), and grant 542961SPI from the Simons Foundation (MRC).

1. M. Carrasco, Visual attention: The past 25 years. *Vision Res.* **51**, 1484–1525 (2011).
2. J. H. Maunsell, Neuronal mechanisms of visual attention. *Ann. Rev. Vision Sci.* **1**, 373 (2015).
3. M. R. Cohen, A. Kohn, Measuring and interpreting neuronal correlations. *Nat. Neurosci.* **14**, 811–819 (2011).
4. H. Ko *et al.*, Functional specificity of local synaptic connections in neocortical networks. *Nature* **473**, 87–91 (2011).
5. A. M. Ni, D. A. Ruff, J. J. Alberts, J. Symmonds, M. R. Cohen, Learning and attention reveal a general relationship between population activity and behavior. *Science* **359**, 463–465 (2018).
6. Y. Gu *et al.*, Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* **71**, 750–761 (2011).
7. D. A. Ruff, C. Xue, L. E. Kramer, F. Baqai, M. R. Cohen, Low rank mechanisms underlying flexible visual representations. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 29321–29329 (2020).
8. D. A. Ruff, M. R. Cohen, Attention can either increase or decrease spike count correlations in visual cortex. *Nat. Neurosci.* **17**, 1591–1597 (2014).

9. T. Kanashiro, G. K. Ocker, M. R. Cohen, B. Doiron, Attentional modulation of neuronal variability in circuit models of cortex. *Elife* **6**, e23978 (2017).
10. C. Huang *et al.*, Circuit models of low-dimensional shared variability in cortical networks. *Neuron* **101**, 337–348 (2019).
11. G. Carlsson, Topological pattern recognition for point cloud data. *Acta Numer.* **23**, 289–368 (2014).
12. P. Y. Lum *et al.*, Extracting insights from the shape of complex data using topology. *Sci. Rep.* **3**, 1–8 (2013).
13. A. E. Sizemore *et al.*, Cliques and cavities in the human connectome. *J. Comput. Neurosci.* **44**, 115–145 (2018).
14. M. Saggat *et al.*, Towards a new approach to reveal dynamical organization of the brain using topological data analysis. *Nat. Commun.* **9**, 1–14 (2018).
15. C. T. Ellis, M. Lesnick, G. Henselman-Petrusek, B. Keller, J. D. Cohen, Feasibility of topological data analysis for event-related fMRI. *Network Neurosci.* **3**, 695–706 (2019).

16. H. Edelsbrunner, J. L. Harer, *Computational Topology: An Introduction* (American Mathematical Society, 2022).
17. 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).
18. C. Giusti, R. Ghrist, D. Bassett, Two-s company, three (or more) is a simplex: 594 algebraic-topological tools for understanding higher-order structure in neural data. arXiv [Preprint] (2016). <http://arxiv.org/abs/1601.01704> (Accessed 18 May 2023).
19. G. Singh *et al.*, Topological analysis of population activity in visual cortex. *J. Vision* **8**, 11 (2008).
20. 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).
21. E. Rybakken, N. Baas, B. Dunn, Decoding of neural data using cohomological feature extraction. *Neural Comput.* **31**, 68–93 (2019).
22. R. J. Gardner *et al.*, Toroidal topology of population activity in grid cells. *Nature* **602**, 123–128 (2022).
23. J. B. Bardin, G. Spreemann, K. Hess, Topological exploration of artificial neuronal network dynamics. *Network Neurosci.* **3**, 725–743 (2019).
24. H. Adams *et al.*, Persistence images: A stable vector representation of persistent homology. *J. Mach. Learn. Res.* **18**, 1–35 (2017).
25. H. Edelsbrunner, D. Morozov, "Persistent homology" in *Handbook of Discrete and Computational Geometry* (Chapman and Hall/CRC, 2017), pp. 637–661.
26. S. Y. Oudot, *Persistence Theory: From Quiver Representations to Data Analysis* (American Mathematical Society, 2017), vol. 209.
27. A. Adcock, E. Carlsson, G. Carlsson, The ring of algebraic functions on persistence bar codes. arXiv [Preprint] (2013). <http://arxiv.org/abs/1304.0530> (Accessed 18 May 2023).
28. M. Rucco, F. Castiglione, E. Merelli, M. Pettini, "Characterisation of the idiotypic immune network through persistent entropy" in *Proceedings of ECCS 2014* (Springer, 2016), pp. 117–128.
29. D. A. Ruff, A. M. Ni, M. R. Cohen, Cognition as a window into neuronal population space. *Ann. Rev. Neurosci.* **41**, 77 (2018).
30. M. R. Cohen, J. H. Maunsell, Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* **12**, 1594–1600 (2009).
31. J. F. Mitchell, K. A. Sundberg, J. H. Reynolds, Spatial attention decorrelates intrinsic activity fluctuations in macaque area v4. *Neuron* **63**, 879–888 (2009).
32. N. C. Rabinowitz, R. L. Goris, M. Cohen, E. P. Simoncelli, Attention stabilizes the shared gain of v4 populations. *Elife* **4**, e08998 (2015).
33. A. S. Ecker, A. S. Tolias, Is there signal in the noise? *Nat. Neurosci.* **17**, 750–751 (2014).
34. J. R. Geiger, J. Lübke, A. Roth, M. Frotscher, P. Jonas, Submillisecond ampa receptor-mediated signaling at a principal neuron-interneuron synapse. *Neuron* **18**, 1009–1023 (1997).
35. Z. Xiang, J. R. Huguenard, D. A. Prince, Gabaa receptor-mediated currents in interneurons and pyramidal cells of rat visual cortex. *J. Physiol.* **506**, 715–730 (1998).
36. M. C. Angulo, J. Rossier, E. Audinat, Postsynaptic glutamate receptors and integrative properties of fast-spiking interneurons in the rat neocortex. *J. Neurophysiol.* **82**, 1295–1302 (1999).
37. R. L. Goris, J. A. Movshon, E. P. Simoncelli, Partitioning neuronal variability. *Nat. Neurosci.* **17**, 858–865 (2014).
38. J. D. Semedo, A. Zandvakili, C. K. Machens, M. Y. Byron, A. Kohn, Cortical areas interact through a communication subspace. *Neuron* **102**, 249–259 (2019).
39. R. C. Williamson *et al.*, Scaling properties of dimensionality reduction for neural populations and network models. *PLoS Comput. Biol.* **12**, e1005141 (2016).
40. S. F. Muldoon *et al.*, Stimulation-based control of dynamic brain networks. *PLoS Comput. Biol.* **12**, e1005076 (2016).
41. E. Tang, D. S. Bassett, Colloquium: Control of dynamics in brain networks. *Rev. Mod. Phys.* **90**, 031003 (2018).
42. D. H. Serrano, J. Hernández-Serrano, D. S. Gómez, Simplicial degree in complex networks. Applications of topological data analysis to network science. *Chaos, Solitons Fractals* **137**, 109839 (2020).
43. K. P. Szymula, F. Pasqualetti, A. M. Graybiel, T. M. Desrochers, D. S. Bassett, Habit learning supported by efficiently controlled network dynamics in naive macaque monkeys. arXiv [Preprint] (2020). <http://arxiv.org/abs/2006.14565> (Accessed 18 May 2023).
44. S. Gu *et al.*, Controllability of structural brain networks. *Nat. Commun.* **6**, 1–10 (2015).
45. B. T. Hutchinson, Toward a theory of consciousness: A review of the neural correlates of inattentional blindness. *Neurosci. Biobehav. Rev.* **104**, 87–99 (2019).
46. F. J. Luongo *et al.*, Mice and primates use distinct strategies for visual segmentation. *bioRxiv* (2021). <https://doi.org/10.1101/2021.07.04.451059> (Accessed 18 May 2023).
47. D. M. Green *et al.*, *Signal Detection Theory and Psychophysics* (Wiley, New York, NY, 1966), vol. 1.
48. N. A. Macmillan, C. D. Creelman, *Detection Theory: A User's Guide* (Psychology Press, 2004).
49. H. S. Bashinski, V. R. Bacharach, Enhancement of perceptual sensitivity as the result of selectively attending to spatial locations. *Percept. Psychophys.* **28**, 241–248 (1980).
50. H. J. Müller, J. M. Findlay, Sensitivity and criterion effects in the spatial cuing of visual attention. *Percept. Psychophys.* **42**, 383–399 (1987).
51. C. J. Downing, Expectancy and visual-spatial attention: Effects on perceptual quality. *J. Exp. Psychol.: Hum. Percept. Perform.* **14**, 188 (1988).
52. H. L. Hawkins *et al.*, Visual attention modulates signal detectability. *J. Exp. Psychol.: Hum. Percept. Perform.* **16**, 802 (1990).
53. V. Wyart, S. Dehaene, C. Tallon-Baudry, Early dissociation between neural signatures of endogenous spatial attention and perceptual awareness during visual masking. *Front. Hum. Neurosci.* **6**, 16 (2012).
54. T. Z. Luo, J. H. Maunsell, Attentional changes in either criterion or sensitivity are associated with robust modulations in lateral prefrontal cortex. *Neuron* **97**, 1382–1393 (2018).
55. T. Z. Luo, J. H. Maunsell, Attention can be subdivided into neurobiological components corresponding to distinct behavioral effects. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 26187–26194 (2019).
56. C. Van Vreeswijk, H. Sompolinsky, Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).
57. J. Mariño *et al.*, Invariant computations in local cortical networks with balanced excitation and inhibition. *Nat. Neurosci.* **8**, 194–201 (2005).
58. L. F. Rossi, K. D. Harris, M. Carandini, Spatial connectivity matches direction selectivity in visual cortex. *Nature* **588**, 648–652 (2020).
59. N. Fourcaud-Trocmé, D. Hansel, C. Van Vreeswijk, N. Brunel, How spike generation mechanisms determine the neuronal response to fluctuating inputs. *J. Neurosci.* **23**, 11628–11640 (2003).
60. T. M. Karrer *et al.*, A practical guide to methodological considerations in the controllability of structural brain networks. *J. Neural Eng.* **17**, 026031 (2020).
61. A. Sizemore, C. Giusti, D. S. Bassett, Classification of weighted networks through mesoscale homological features. *J. Complex Networks* **5**, 245–273 (2017).
62. A. E. Sizemore, J. E. Phillips-Cremins, R. Ghrist, D. S. Bassett, The importance of the whole: Topological data analysis for the network neuroscientist. *Network Neurosci.* **3**, 656–673 (2019).
63. S. P. Patankar, J. Z. Kim, F. Pasqualetti, D. S. Bassett, Path-dependent connectivity, not modularity, consistently predicts controllability of structural brain networks. *Network Neurosci.* **4**, 1091–1121 (2020).
64. A. M. Ni, M. R. Cohen, NiHuangDoironCohen2022. Open Science Framework. <https://doi.org/10.17605/OSF.IO/RN7TU>. Deposited 2 May 2022.
65. C. Huang, TopologyForAttention. Open Science Framework. <https://github.com/TCR23/TopologyForAttention.git>. Deposited 21 July 2022.