

Geometry of spiking activity in early visual cortex: a Topological Data Analytic approach

Andrea Guidolin* Mathieu Desroches[†] Jonathan D. Victor[‡] Keith P. Purpura[§]
Serafim Rodrigues[¶]

April 20, 2022

Abstract

In the brain, spiking patterns live in a high-dimensional space of neurons and time. Thus, determining the intrinsic structure of this state space for neural activity presents a theoretical and experimental challenge. To address this challenge, we introduce a new framework for applying topological data analysis (TDA) to spike train data to determine the geometry of neural activity space. Key to our approach is a parametrized family of distances based on the timing of spikes to determine the dissimilarity between neuronal responses. We applied TDA to visually driven single-unit and multiple single-unit spiking activity in macaque V1 and V2. TDA across timescales reveals a common geometry in V1 and V2 that is most consistent with a low-dimensional space endowed with Euclidean or hyperbolic geometry with modest curvature. Remarkably, the inferred geometry depends on timescale, and is most distinct for the timescales that are important for encoding contrast, orientation, and spatial correlations.

Introduction

The broad goal of understanding the relationship of neural activity patterns to cognition and behavior entails the intermediate step of characterizing the space in which neural activity lives. This is a challenging problem, even for small populations of neurons, because of the dimensional explosion and data limitations: it is simply not feasible to obtain sufficient experimental data to exhaustively sample the space of all possible activity patterns [28]. For this reason, indirect methods are required. A promising strategy, applicable in a wide variety of contexts, is the persistent homology approach to topological data analysis (TDA) [15, 20, 21, 30, 7]. Here, an abstract notion of distance is used to compare samples of neural activity, nodes in a neural network, or points in a sensory domain. Based on these pairwise distances, a sequence of network graphs is constructed, with successive graphs built by connecting the nodes at increasingly greater distances. These graphs are then associated with higher-dimensional topological objects called simplicial complexes, which are characterized via their Betti numbers, describing (intuitively speaking) the number of disconnected components, the number of holes, the number of “bubbles”, etc. The way that the Betti numbers change as a function of the distance criterion then provides a topological characterization of the space.

Here, we use this approach to analyze the activity patterns of clusters of cortical neurons, focusing on the dynamics of neural activity. This requires confronting another issue, but one that

*MCEN Team, BCAM - Basque Center for Applied Mathematics, 48009 Bilbao, Basque Country, Spain and Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

[†]MathNeuro Team, Inria at Université Côte d’Azur, Sophia Antipolis, France

[‡]Feil Family Brain and Mind Research Institute, Weill Cornell Medical College, New York, NY 10065, USA, jdvicto@med.cornell.edu (JV)

[§]Feil Family Brain and Mind Research Institute, Weill Cornell Medical College, New York, NY 10065, USA, kpurpura@med.cornell.edu (KPP)

[¶]MCEN Team, BCAM - Basque Center for Applied Mathematics, 48009 Bilbao, Basque Country, Spain and ikerbasque - The Basque Foundation for Science, Bilbao, Basque Country, Spain, srodrigues@bcamath.org (SR)

the persistent homology approach is also well-suited to handle: even for a single neuron, activity is not a scalar quantity, but rather, a sample of a point process. Thus, it is natural to measure distances between neuronal responses in a way that is sensitive to temporal patterns, such as the spike distance introduced by Victor and Purpura [26]. While it is well-recognized that TDA methods can be applied to any distance, or even dissimilarity measure, defined on the data space, previous applications of the persistent homology approach to neural data have not exploited this generality. Rather, they have been based on distances computed either from correlations [8, 9], effectively dot products in a vector space, or via a more indirect embedding of the data in a vector space, e.g. based on cross-correlograms integrated over different timescales [15]. Whenever the spike train data is embedded in a vector space as an intermediate step to compute a distance between spike trains, the resulting metric structure may reflect some geometrical aspects of the vector space embedding. The spike distances used here, based on a notion of “cost” required to transform a spike train into another one, are fundamentally different from any distance calculated via a vector space embedding of the spike train data. While spike distances do not correspond to Euclidean distances [4], they do satisfy the triangle inequality. As is well known, TDA extends seamlessly to this more general framework.

While it is natural to consider an instance of population activity (e.g. the recent history of firing in each of the sampled neurons) to be a point in a high-dimensional vector space, and straightforward to use this vector space to construct distances, there are theoretical and empirical reasons for applying TDA to more general notions of distance, such as the Victor-Purpura spike distance.

The theoretical consideration is that using a vector space framework to construct distances implicitly assumes that the set of spike trains has a vector space structure relevant to the computation of distances: for example, that between any two spike trains, there is a third spike train that is halfway between them. The empirical consideration is the emerging evidence that in some sensory domains perceptual distances between sensory stimuli imply a non-Euclidean (specifically, hyperbolic) geometry [30], which is at odds with the Euclidean distances that underlie standard vector space methods. Another aspect of the approach advanced here - applying TDA to spike distances - is that it makes use of single responses, rather than average responses collected over an extended period of time or repeated trials. Thus, correlation-based measures and spike distances are in some sense complementary: correlation-based measures capture what is typical about a neuron’s firing patterns, while spike distances focus on information that is available to the organism on a single trial.

We applied this analysis to tetrode recordings obtained under anesthesia in primary visual (V1) and secondary visual (V2) cortex of the macaque monkey. To increase the likelihood that the observed activity patterns would represent the range of physiologic firing patterns, we presented visual stimuli that contained a range of types of statistical structures. Specifically, stimuli were drawn from a 10-parameter space of synthetic visual textures [24], and included textures that varied in first-, second-, third-, and fourth-order local spatial correlations known to be present in natural images [16] and to which V1 and V2 neurons (and human observers [27]) are sensitive [29].

A key aspect of spike distances is that they form a family parametrized by a characteristic timescale, thus allowing temporal structure of neural activity to be probed across a range of scales. Applying TDA in parallel across these scales reveals the timescale at which the neural data has the greatest structure, i.e., where the disparity between the observed activity and various types of surrogates with shuffled or resampled spike times is the greatest. Interestingly, our analysis shows that this timescale is similar to the timescales that are most relevant for neural coding of visual attributes such as contrast, orientation and spatial patterns [25].

Finally, we find that a novel way of passing from distance measures to graphs – linking nodes according to greatest dissimilarity rather than greatest similarity (using a “decreasing filtration” rather than the standard “increasing filtration”) – sharpens the distinction between the experimental data and candidate models.

Results

The starting point of our analysis is 28 datasets of single-unit spike train activity recorded from areas V1 and V2 of the visual cortex of anesthetized macaques during stimulation with visual patterns drawn from a high-dimensional space of visual textures [24]. This stimulus space is sampled along its 10 axes, which correspond to 10 independent and visually salient kinds of spatial correlation. We sample each axis at four points: two levels of positive correlation and two levels of negative correlation (see Materials and Methods for details). 64 examples of each of these 40 texture types (10 axes, 4 sample points per axis), along with 64 different examples of an uncorrelated texture (responses not analyzed), were presented in random order, with each stimulus lasting 320 ms, a sequence lasting approximately 14 minutes. This sequence was then repeated four times at each recording site. Responses to each of the stimulus types for each repeat were then assembled to yield a total of $40 \times 4 = 160$ collections per dataset. Thus, a collection consists of a set of neuronal responses to 64 stimuli that are perceptually similar. Depending on the dataset, the responses consisted of the spike trains of one or more isolated neurons. If more than four neurons were simultaneously recorded, we considered only the four neurons with the highest firing rates.

To study the topology of the space that neural activity occupies, we selected the 80 collections in each dataset with the smallest number of empty responses (0 spikes) (the vast majority of these collections had at least 60 non-empty responses; see Fig. S1 (D)). This choice is motivated by the fact that our analysis maps spike trains with at least one spike to distinct points in the topological space, while all empty spike trains are mapped to the same point. Thus, this selection criterion maximized the sampling of the space occupied by a collection. We then characterized the topology of each of these collections via the persistent homology approach, and summarized the resulting characterizations to identify their consistent features.

The TDA method we used follows very closely the framework proposed by [15] but with the difference that we first evaluate a matrix of distances between neuronal responses. Specifically, we used the multineuron Victor-Purpura distance [25, 3] to define the distances between all pairs of neuronal responses. This choice to employ a distance rather than the correlations between neuronal responses to quantify (dis)similarity is motivated by several considerations. First, our measured neuronal responses are sparse and only lasted 320 ms. Second, correlations are intrinsically Euclidean - in that they correspond to dot products or angles in a vector space - while spike train distances need not be. Use of the Victor-Purpura distance [25, 3] also allows us to build on previous studies of the visual cortex that employ data analytical techniques other than TDA. Finally, our approach focused on information that is available from a single trial - as would be the case under natural viewing conditions - rather than information that can only be obtained by pooling responses across multiple stimulus presentations.

The Victor-Purpura distance is a cost-based distance that has two parameters: a parameter q that controls the timescale used to quantify dissimilarity in spike timing, and a parameter k that determines the relevance of the neuron of origin of each spike. The distance between two neuronal responses is defined as the minimum cost of transforming one into the other via a sequence of basic moves: addition or deletion of a single spike, with a cost of 1, shift of a single spike by a time interval Δt , with a cost $q|\Delta t|$, or change in the neuron of origin, with cost k . Thus, $q = 0$ corresponds to ignoring spike time but retaining spike count, $q > 0$ corresponds to considering temporal structure at a scale of $1/q$. If a spike time in one train is within $1/q$ of a spike time in a second train, they are considered to correspond, as they contribute less than 1 unit to the dissimilarity. Conversely, if spikes are more than $2/q$ apart, they are considered to be unrelated, since shifting them into correspondence would incur a higher cost than deleting the spike from one train and inserting it into the other. Thus, the family of distances range from measures of dissimilarity that ignore spike timing altogether ($q = 0$), measuring distance based only on spike count, to measures that consider spike timing meaningful at an arbitrarily high level of precision ($q \rightarrow \infty$). Similarly, $k = 0$ corresponds to ignoring the neuron of origin, while $k = 1$ corresponds to considering a change in the neuron of origin equivalent to the insertion of a spike. Further

background on the Victor-Purpura distance, as well as efficient dynamic-programming algorithms for calculating it, can be found in [26, 2]. We carried out TDA for a grid of values of these two parameters: $q = 1, 2, 5, 10, 20, 50, 100, 200$ (sec^{-1}) and $k = 0, 1$. The resulting grid covered the range found relevant to neural coding in previous studies [25, 3].

For each choice of parameters of the Victor-Purpura distance we computed topological summaries via Betti curves [15] (Fig. 1) and characterized relevant features of the topological space by computing mean Betti curves (averaged over 80 collections of each dataset). Specifically, the Betti curves are constructed from a sequence of graphs, in which the nodes of all graphs are the individual responses, and the edges of the graph are determined by thresholding the distance matrix. Starting from a set of disconnected nodes, the graph sequence is built by adding one edge at a time, according to the pairwise distances. When the graph edges are weighted by a dissimilarity measure (e.g. a distance) between the nodes, it is standard to rank these weights in increasing order (the *increasing filtration*), like in the *clique topology* method of [15]; here, we also consider the weights in decreasing order (the *decreasing filtration*), a method introduced in [18] by the name *weight rank clique filtration*. As the edges of the graph fill in, the graph becomes enriched with higher-dimensional information, as each n -clique (subgraph of n all-to-all connected nodes) in the graph is regarded as an $(n - 1)$ -dimensional simplex: 3-cliques are viewed as triangles (2-dimensional simplices), 4-cliques are viewed as tetrahedra (3-dimensional simplices), etc. In this way, each graph becomes a *simplicial complex*, called the *clique complex* of the graph, in which arrangements of cliques can enclose “holes” of different dimensions. The number of 1-dimensional tunnels, 2-dimensional voids and 3-dimensional “cavities”, known as Betti numbers and denoted by β_1 , β_2 and β_3 respectively, are computed at each step of the filtration, starting with no edges and progressively adding the edge with the next smallest weight (increasing filtration) or next largest weight (decreasing filtration). The Betti numbers are thus functions of the edge density ρ , the fraction of potential edges that have been filled in. These functions ($\beta_1(\rho)$, $\beta_2(\rho)$ and $\beta_3(\rho)$) are the Betti curves, which we computed over the range from $\rho = 0$ to 0.6. Betti curves are sensitive to the geometry and topology of the metric space of data points [15, 21, 30], and reflect the overall structure of the set of distances, however they are invariant under monotonic transformations of the distance [15]. This means that they are sensitive only to the relative ordering of all distances. To summarize our application of the Betti curves method: each of the 80 collections from the 28 datasets, containing at most 64 neuronal responses to stimuli of the same texture type, is regarded as a set of nodes in the described filtration process, in which the weights are given by the Victor-Purpura distance between pairs of responses. Fig. 1B displays the mean Betti curves, averaged over the 80 collections, of one dataset.

Below, we will use the mean Betti curves over 80 collections for further analysis, but first we investigated whether the Betti curves had a substantial dependence on the two experimental variables that distinguished the collections: whether the recordings were in V1 vs. V2, and whether the spatial structure of the stimulus was low-order (first- and second-order) vs. high-order (third- or fourth-order). Low-order correlation structure is extracted by linear receptive fields; third- and fourth-order correlation structure is primarily extracted in V2 [29]. To compare the topological characterizations of the collections of responses, we summarized the information of the Betti curves by considering the integral of the curves [15], referred to as integrated Betti values. Fig. 2 shows that the integrated Betti values of the collections recorded in areas V1 and V2 of the visual cortex have similar distributions. A remarkable consistency can be observed across all values of the parameters of the Victor-Purpura distance, and for both the increasing and decreasing filtration. A similar behavior of the distributions of integrated Betti values is observed (Fig. 3) if the data is split in two groups according to the visually-salient structure of the textures of the stimuli, to compare responses to first- and second-order spatial correlations to third- and fourth-order spatial correlations (see [27] for details on the definitions). These results, together with the remainder of the analysis, suggest that our topological methods applied to spike train data are sensitive to intrinsic population dynamics, rather than stimulus-driven dynamics. The similarity across recording areas and texture types of the results of TDA justifies our choice of pooling these results in the further

analysis below and, in particular, of considering mean Betti curves over the 80 collection of each dataset.

To assess which aspects of the spike train data might be responsible for the shape of these mean Betti curves (e.g., Fig. 1B), we synthesized surrogate spike train data of four different types by perturbing specific aspects of the original experimental data.

1. *Uniform resampling of spike times* (U). Each sequence of spikes from each neuron from each response is replaced by a sequence of the same number of spikes randomly distributed over the length of the response interval (0 – 320ms). The sequence of labels indicating which unit fired the spikes is preserved, as is the number of spikes of each neuron in the response.
2. *Exchange resampling of spike times between collections* (EB). Spikes are randomly swapped between the 80 selected collections of a dataset, preserving their time of occurrence in the 320ms response and their neuron of origin. Thus, each response in the surrogate dataset has the same number of spikes as the original, and the overall distribution of spike times across the dataset are maintained.
3. *Exchange resampling of spike times within collections* (EW). As in 2, but with the swapping restricted to responses within each collection. Thus, spike counts are unchanged within each response, as is the distribution of spike times within each collection.
4. *Poisson generated spike trains* (P). For each neuron contributing to a dataset, we determine its overall firing rate across the dataset. We then replace each neuron’s spike train in all non-empty responses in the dataset with a sample of a Poisson process with that rate. We set the bin width in the Poisson generator to 0.0001ms. Note that this surrogate dataset does not preserve the number of spikes of each neuron in each response, only the average.

For each experimental dataset, we generated 20 surrogate spike datasets for each of the above four types, and computed Betti curves for these surrogates just as the curves were computed for the experimental data.

An example of this analysis is shown in Fig. 4. As is typical across the datasets, the Betti curves associated with the experimental data were often quite distinct from those of the surrogate sets. These differences were most prominent in the mid-range values of q (5 to 50 sec^{-1}), and also more prominent for the decreasing filtration than for the increasing filtration. In addition, the Poisson surrogates usually show a greater divergence from the real data than the other surrogates. One possible reason for this is that, unlike the other surrogates, the Poisson surrogates do not preserve the number of spikes in each response.

To determine the consistency of these features across the 28 datasets, we used the integrated Betti values of the Betti curves. Thus, for every fixed Betti number (β_1 , β_2 or β_3), we considered the difference between the integrated Betti value of each surrogate (mean value over the 20 computations) and the integrated Betti value of the experimental data, and averaged them over the 80 collections from each dataset (Fig. 5). Consistent with Fig. 5 and with the single dataset of Fig. 4, the behavior of the experimental data departs from the behavior of the surrogates, especially for the mid-range values of q (5 to 50 sec^{-1}). This holds both when the neuron of origin is ignored ($k = 0$) and when it is considered relevant ($k = 1$), and it is seen for both increasing and decreasing filtrations. All surrogates behave in a way that is distinct from the experimental data. As in Fig. 5, three of the surrogates (U, EB, EW) behave similarly to each other, while the Poisson surrogate (P) deviates from the data more extensively.

Fig. 6 assesses the statistical significance of these observations, via the two-sample Kolmogorov-Smirnov (KS) test applied to the distributions of integrated Betti values. Even for the EW surrogate, which perturbs the recorded spike trains the least, the KS test statistic takes on large values, especially for mid-range values of q , rejecting (at the 0.05 significance level) the null hypothesis that the EW surrogate and the experimental data samples come from the same distribution. More

generally, the analysis shows that the surrogates for U, EB, and EW diverge in a similar, time-scale-dependent manner from the experimental data, while the Poisson generated data is always significantly different from both the experimental data and the other surrogates.

In sum, Figs. 2-3 and 5-6 demonstrate that the neural activity elicited by images with a range of statistical properties, and recorded in either V1 or V2, have a common topological structure that is distinct from that of surrogates with matching spike counts and spike timing distributions. These differences are present over a broad range of temporal scales, and are most prominent at a temporal scale of $5 - 50 \text{ sec}^{-1}$, i.e., 20 to 200 ms.

To gain insight into this topological structure, we compared the measured Betti curves with those obtained from randomly-assigned distances or by sampling points in simple geometric spaces [15, 20, 30]. Specifically, we computed the Betti curves associated with: (i) a random symmetric 64×64 matrix with zeros on the diagonal [15], (ii) a sample of 64 random points in a unit (hyper)cube within a d -dimensional Euclidean space [15], and (iii) a sample of 64 random points in the hyperbolic ball model of the d -dimensional hyperbolic space [30]. For Euclidean and hyperbolic spaces, we considered dimensions d from 1 to 15. For the hyperbolic models, we also varied the effective curvature. As in [30], we implemented this by keeping intrinsic (Gaussian) curvature fixed at -1 but changing typical distances between the points (i.e., varying from 1 to 5 the maximum radius R_{\max} of the hyperbolic ball model, see Materials and methods). We therefore refer to models with low R_{\max} as hyperbolic spaces with moderate curvature. For each random or geometric model, we computed the integrated Betti values of 300 sets of samples. We then compared these distributions with the integrated Betti values of the Betti curves selected from our data collections. We restricted the analysis to collections with at most one empty response, in order to have 64 data points (and to exactly match the geometric models), and we randomly chose one such collection for every dataset that had collections meeting this criterion. This yielded integrated Betti values from 18 different datasets.

To determine compatibility with the geometric models, we tallied how many of the experimental integrated Betti values were within 3 standard deviations of the mean of the 300 values of the model for all Betti numbers β_1 - β_3 and for both increasing and decreasing filtrations. This procedure was followed for the grid of values of the k and q parameters of the Victor-Purpura distance.

While none of these models were fully compatible with the data, the Euclidean and low-curvature hyperbolic models ($R_{\max} = 1$) came the closest (Fig. 7). Compatibility is maximal for low dimensions ($d = 3$ to 5) in the range we considered, consistently across the geometric models. In all cases, Betti curves associated with the experimental data were found to be incompatible with the random symmetric matrix model. Interestingly, across the geometric models, compatibility between the experimental data and the models is maximized in a consistent region of the parameter space of the Victor-Purpura distance. These regions approximately correspond to $q = 5$ to 20 sec^{-1} for $k = 0$, and $q = 1$ to 10 sec^{-1} for $k = 1$, while for large values of q the compatibility is very low in all cases.

The data analyzed here were collected from monkeys under propofol anesthesia/sufentanil analgesia and neuromuscular blockade. As a consequence, fixational eye movements could not have contributed to the fluctuations in spiking activity we observed in V1 and V2, but conversely, the natural dynamics of the visual input due to such eye movements is only partially approximated by the transient mode of stimulus presentation that we used. Another caveat is that the anesthesia and opiate analgesia may have produced noise correlations in the neural activity that contributed to the topological structure we extracted through TDA. While it is impossible to rule out any impact of anesthesia and analgesia on our results, there are two reasons that it is probably minor. First, as mentioned above, the timescales associated with the most consistent and distinctive state-space structure observed here corresponded to the timescales that are most important for carrying visual information in neurons in the visual cortex of the awake macaque [25]. Second, available evidence suggests that the effect of drug-induced changes in network state would be more likely to dilute any underlying structure, than to create it. Specifically, Ecker and co-workers [11] compared the variability of the spiking activity in V1 in awake monkeys with that seen in monkeys under sufentanil,

inferring the presence of a sufentanil-related state variable. This state variable fluctuated with a timecourse that varied from 50 to 1650ms. If similar sufentanil-related fluctuations were present during our recording sessions, the impact would have been distributed randomly across the many 320ms spiking response samples.

Discussion

Understanding how the cortex accomplishes the many components of behavior and cognition requires understanding the dynamics of spiking activity across populations of neurons. Thus, characterizing the state space in which these dynamics operate – the space of configurations of neural population activity – is a matter of fundamental importance in systems and cognitive neuroscience. The present manuscript identifies the topological and geometrical features of this space by an analysis of the spiking activity generated in V1 and V2 of the macaque monkey during robust visual stimulation. While no simple model such as a circle, torus, or sphere recapitulates these features, they are most consistent with low-dimensional models of Euclidean geometry or hyperbolic geometry with a modest amount of negative curvature (Fig. 7). Moreover, these characteristics appear to be largely independent of the statistics of the visual stimuli or the cortical region of the recordings (Figs. 2-3).

We also show that topological data analysis (TDA) provides a way to characterize neuronal state spaces in a way that avoids the limitations of vector space methods. TDA begins with the pairwise distances between instances of population activity. While several previous applications of TDA to neural spiking data have assumed a vector space structure or used a vector space embedding to compute these distances [15, 8, 9, 20], there is no fundamental reason to do so. Rather, the primary requirement is that the distance between two samples of neural activity corresponds to their dissimilarity. In [5], TDA methods combined with non-parameteric (dis)similarity measures between spike trains were applied to synthetic spike train data for regimes classification in artificial neural networks. Here, we apply TDA in conjunction with a parametrized family of distances that have been shown to capture the meaningful dissimilarities between spike trains in several contexts [25, 1, 10, 12], and that are not Euclidean [4].

Recognizing that neural activity is the result of an interaction of the inputs to the population and its intrinsic network properties, we sought to identify characteristics of neuronal state space that might apply to a wide range of visual stimuli. To this end, we used stimuli consisting of mathematically-defined visual textures [27, 29] of known relevance to natural vision [22]. Critically, this stimulus set included some textures whose visually-salient structure (first- and second-order spatial correlations) could be extracted by simple center-surround operations in the retina, and other textures whose structure (third- and fourth-order spatial correlations) can only be extracted by cortical computations, primarily in V2 [29]. As we show (Figs. 2-3), the characteristics we identify are independent of this distinction. Moreover, these characteristics are shared by V1, where only first- and second-order spatial correlations are robust drivers of activity, and V2, where many neurons are sensitive to higher-order correlations [29]. Second, we matched the orientation and spatial scale of the texture stimuli to the receptive field properties of the well-isolated single-units in the tetrode recordings. This allowed us to collect robust spiking activity even under propofol anesthesia/sufentanil analgesia [29]. These responses included strongly driven modulations of firing rate driven by the texture type, as well as trial-to-trial variability, thus facilitating wide exploration of the state space. We further promoted the exploration of the state space through the use of 64 distinct exemplars of each texture type. Finally, to ensure sensitivity to the details of neural activity, we based the analysis on individual responses rather than their averages, and used a distance that was sensitive to the timing of individual spikes rather than just overall spike count [20] or firing rate envelope [23].

Our approach recognizes that the timing of individual spikes is potentially important but it is agnostic as to what the relevant timescales are. This viewpoint is implemented by applying TDA to a family of distances controlled by a parameter q (in units of sec^{-1}) that sets the resolution with

which spike timing influences the measure of dissimilarity. Applying TDA to the resulting family of distances for $q = 1$ to $q = 200$ ms allows us to focus on a range of timescales from 1sec down to 5ms.

The family of distances has a second parameter, k , which controls the importance of the neuron of origin in determining dissimilarity. For $k = 0$, the neuron of origin is irrelevant; for $k = 1$, changing the cost of the neuron of origin has unit cost. These extremes proved useful in analyzing multineuronal coding of spatial phase [3], demonstrating maximal information near $k = 1$; here, the dependence of state space geometry on k was relatively small.

Importantly, the distinction between experimental data and the random surrogates (U, EB, EW, P) depends systematically on the temporal resolution q of the distance between spike trains (Figs. 5-6 and S2-S5). The distinction is greatest for intermediate values of $q = 5 - 50 \text{ sec}^{-1}$, corresponding to resolutions of 200 to 20 ms). For low values of $q (< 5)$, the distinction is generally lost, as the Victor-Purpura distance progressively disregards spike timing; for higher values (> 100), it is also generally weaker, indicating that at these timescales (with interspike intervals ≈ 10 ms) the systematic effect on geometry is less pronounced. Although the behavior of Poisson surrogate statistically diverges from the experimental data for the decreasing filtration method and high values of q (Figs. 4, 6 and S2-S3), we observe (Figs. 5-6 and S2-S5) that the distinction is clearly visible across the whole range of parameters and for both filtration methods. One can interpret the dependence between the timescale parameter q of the Victor-Purpura distance and the difference between experimental data and the random surrogates as the temporal scale of distinctive geometries that structure the cortical activity space. It is notable that the timescales which most clearly reveal the geometry of the state space correspond to the timescales that are most informative for carrying visual information about contrast, spatial frequency, orientation, and texture – both as identified by analysis methods based on the Victor-Purpura distance [25], and by unrelated approaches [14]. Because the timescales at which the ongoing activity has the most distinctive geometric structure is similar to the timescales that are most informative about visual features, it is reasonable to hypothesize that this match facilitates transmitting visual information.

One way to interpret the results with the random surrogate data is to consider the degrees of freedom of the spiking activity. The Betti curves of the surrogates typically have higher values than the experimental data for intermediate values of q (Figs. 5 and S2-S3). This may indicate that the activity space for the real spiking activity is more constrained than the synthetic spiking activity. In this range ($q = 5 - 50 \text{ sec}^{-1}$), the random manipulation of the spike data in different ways consistently opens up holes in the cortical activity state space, leading to higher Betti values (Figs. 4-5 and S2-S3). At the neural level this suggests that neural activity in the space produced during visual stimulation is more structured than the surrogates.

The comparison with random and geometric models highlights complementary aspects of the geometry of the space of sampled spike trains endowed with Victor-Purpura distances. Upon defining a notion of compatibility that accounts for the Betti curves of $\beta_1 - \beta_3$ and both increasing and decreasing filtration, we observed that even if none of the considered models consistently fits the data, Euclidean geometry and hyperbolic geometry with moderate curvature ($R_{\max} = 1$) are more compatible with the geometry of the spike train responses (Fig. 7). Compatibility is concentrated in low dimensions of the geometric models ($d = 3-5$). This observation is consistent with other geometrical descriptions of neural population activity [13]. In these descriptions, the instantaneous firing rates of a large number of neurons are found to occupy a low-dimensional manifold within a high-dimensional space. While in the study described here we analyze recordings from at most four neurons, the distances we consider between spike trains are intrinsic and do not depend on the choice of an embedding, and as a consequence, potentially capture aspects of the neural activity manifold. However, recordings of larger numbers of neurons are needed to demonstrate that this is the case. Fig. 7 also shows that the compatibility between the data and the geometric models systematically depends on the timescale q of the Victor-Purpura distance between spike trains, and is maximized for mid-range values $q = 5 - 20 \text{ sec}^{-1}$ when the neuron of origin of each spike is disregarded ($k = 0$), and for low-range values $q = 1 - 5 \text{ sec}^{-1}$ when the origin of each spike is

accounted for ($k = 0$).

In summary, we introduce a new framework for applying TDA to spike train data, and use it to analyze neuronal activity in macaque visual cortex. There are two main methodological innovations. First, in contrast to most previous applications of TDA to neural data, we do not assume that the spike trains have a vector space embedding that induces distances and correlation measures between them; rather, the topological analysis is directly applied to dissimilarity measures (e.g. distances) that result from considering spike trains to be sequences of events – in this case, the Victor-Purpura spike train distances, which are non-Euclidean. A second innovation of the approach is the filtration – the sequence of simplicial complexes derived from the dissimilarities that are used to compute the Betti curves. Typically, an increasing filtration is used for TDA [21, 7]: graphs are progressively filled in for pairs of points at greater and greater dissimilarities; here we show that the decreasing filtration can reveal a clearer picture of the data’s geometry and topology. Betti curves, especially when averaged over several computations, capture the statistical distribution of the persistent topological features (tunnels, voids, etc.) across a filtration, which serves as a topological descriptor of the underlying metric spaces. In contrast to the usual filtration of graphs by increasing weights and their associated clique complexes, the decreasing filtration is obtained by considering, in reverse order, the *independence complex* of each graph (i.e., the clique complex of the complement of the graph). Betti curves of low dimension (β_1 – β_3) of the decreasing filtration therefore carry information on the arrangements of high-order cliques of the original graphs, which is not related to the homology of the increasing filtration. Because the analysis is carried out for a sequence of distances parameterized by timescale, we are able to identify the range in which the geometric structure of state space is most distinctive: the range 5 – 50 sec⁻¹, i.e., 20 to 200ms. As noted above, this timescale corresponds to the temporal precision that is most informative for decoding visual information from spike trains. This matching of the timescale of state space geometry and the timescale of neural coding may be a general feature of brain networks, and we speculate that networks in other domains (e.g., motor planning, learning, decision-making, etc.) will behave similarly.

Materials and Methods

Experiments

All procedures followed the guidelines provided by the US National Institutes of Health and Weill Cornell Medical College Animal Care and Use Committee. Full details concerning the physiological preparation and multi-tetrode single-unit recordings can be found in [19] and [29]. Detailed descriptions of the visual stimuli, their generation, and their display during the experimental sessions are given in [29], which also details how single-unit activity was characterized, and how the time-series of neural firing events, including the procedures utilized for spike sorting, were extracted from the multi-tetrode recordings. These methods are summarized in SI Materials and Methods.

Data analysis and TDA software

For each collection of neuronal responses we determined the distance matrix $D = (D_{ij})$, where D_{ij} is the Victor-Purpura distance with parameters (k, q) between the i th and the j th response in the collection. The parameters k and q ranged over a grid of values: $k = 0, 1$, which sets sensitivity to neuron of origin; $q = 1, 2, 5, 10, 20, 50, 100, 200$ (sec⁻¹), which sets the relevant scale of temporal structure. The multineuron Victor-Purpura distance between spike trains is computed using the “labdist_faster_qkpara_opt” Matlab function implemented by Thomas Kreuz, available at: http://www-users.med.cornell.edu/~jdvicto/labdist_faster_qkpara_opt.html

We applied and extended the clique topology method introduced in [15], which corresponds to computing the Betti curves associated with our increasing filtration. Given a symmetric matrix M , the method transforms it by rank-ordering its entries, computes the persistent homology of the transformed matrix and finally determines the associated Betti curves. Since the input matrix is

transformed by considering only the rank ordering (as opposed to the specific values) of its entries, the output is invariant to monotonic transformation applied entry-wise to M . Given a dissimilarity or distance matrix, the clique topology method orders the entries increasingly. We observe that in [15] the method is also applied to a matrix $C = (C_{ij})$ of correlations by transforming it into a dissimilarity matrix $D = (D_{ij})$ via the application of any function that inverts the ordering between the absolute values $|C_{ij}|$ of the correlations and the entries D_{ij} , e.g., $D_{ij} = -|C_{ij}|$. In this work, we apply the clique topology method by ordering the entries of our distance matrices both increasingly (increasing filtration) and decreasingly (decreasing filtration), the latter case corresponding to the weight rank clique filtration method introduced in [18].

To compute clique topology and persistent homology, we use a faster and equivalent alternative to the original CliqueTop scripts (<https://github.com/nebneuron/cliq-top>; see [15]), namely Ripser [6] (<https://github.com/Ripser/riper>) on a rank-ordered distance matrix. Starting from a distance matrix $D = (D_{ij})$, we added small (symmetric) Gaussian noise and transformed the resulting symmetric matrix by rank-ordering its entries in increasing (resp., decreasing) order for the increasing (resp., decreasing) filtration method. Note that the purpose of the Gaussian noise is to uniquely specify an ordering of the entries in presence of equal values, and its magnitude was chosen small enough to preserve the ordering between entries with different values. Ripser computes the persistent homology of the clique filtration associated with the ordered matrix and outputs the so-called *barcode*, a collection of pairs of real numbers $B = \{(b_i, d_i)\}_{i=1, \dots, m}$ for the chosen dimension of homology, which are $j = 1, 2, 3$ in our setting. The pairs $\{(b_i, d_i)\}$ are the ranges in edge density for which a particular void of dimension j is present. If the size of the input matrix is $n \times n$, the Betti curve is obtained from the barcode B by setting $\beta_j(r/N)$ to the number of (b_i, d_i) in B such that $b_i \leq r < d_i$, where $N = n(n-1)/2$ and r is any integer between 0 and the maximum value r_{\max} such that r_{\max}/N does not exceed a maximum edge density ρ_{\max} . This parameter was set to 0.6 in our analysis for computational efficiency, following [15]. The function $\beta_j(\rho)$ for $\rho \in [0, \rho_{\max}]$ is obtained by linear interpolation.

Boxplots are generated with the Matlab function “boxplot”. The two-sample Kolmogorov-Smirnov test was performed using the Matlab function “kstest2”.

Geometric models

Models of random and geometric spaces were considered for Fig. 7. The random symmetric matrices and the distance matrices of random points in a Euclidean space were generated following [15], while the distance matrices of random points in a hyperbolic space were generated following [30].

1. *Random symmetric matrix model.* The nonzero (off-diagonal) elements of the distance matrix are randomly-chosen positive numbers uniformly distributed in the interval $(0, 1)$, with $D_{ij} = D_{ji}$, unconstrained by the triangle inequality.

2. *Euclidean geometry model.* To generate a random Euclidean distance matrix $D = (D_{ij})$, 64 random points are uniformly sampled in a unit (hyper)cube within the d -dimensional Euclidean space, for a fixed dimension d between 1 and 15. Each entry D_{ij} is set equal to the Euclidean distance between the randomly-chosen i th and j th points.

3. *Hyperbolic geometry model.* Similarly to [30], we generated a random hyperbolic distance matrix $D = (D_{ij})$ by uniformly sampling 64 points in the d -dimensional hyperbolic space (for a fixed dimension d between 1 and 15), using the hyperbolic ball model [17] with curvature $\zeta = 1$. We sampled a standard d -variate Gaussian (using the Matlab function “randn”) and rescaled the radii of the sampled points by selecting radii r within $[0, R_{\max}]$ following the distribution $\rho(r) \sim \sinh((d-1)r)$. We examined R_{\max} values in the range from 1 to 10, and results for $R_{\max} = 1, 2, 5$ are shown in Fig. 7. The distance D_{ij} between two points of radii r_i and r_j with an angle $\Delta\theta$ between them is determined from the hyperbolic law of cosines

$$\cosh \zeta D_{ij} = \cosh(\zeta r_i) \cosh(\zeta r_j) - \sinh(\zeta r_i) \sinh(\zeta r_j) \cos(\Delta\theta)$$

with curvature $\zeta = 1$ in our setting.

Data and code availability

The data related to this article is available at <https://github.com/aguidolin/visual-spike>. A set of Matlab functions and scripts to replicate the analysis is available at <https://github.com/aguidolin/top-spike>. Additional code related to this article may be requested from the authors.

Acknowledgments

JDV and KPP acknowledge Qin Hu, Ferenc Mechler, Eyal Nitzany, Anita Schmid, Dan Thengone and Yunguo Yu for assisting with data collection. They also acknowledge NIH grant numbers EY09314 and EY07977 from the National Eye Institute of the NIH. SR was supported by Ikerbasque (The Basque Foundation for Science); AG and SR were supported by the Basque Government through the BERC 2018-2021 program and by the Ministry of Science, Innovation and Universities: BCAM Severo Ochoa accreditation SEV- 2017- 0718 and through project RTI2018-093860B-C21 funded by (AEI/FEDER, UE) with acronym “MathNEURO”. AG acknowledges the support of the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. MD and SR acknowledge support of Inria through the Associated Team “NeuroTransfSF”.

Additional information

Author contribution

All authors designed the research; KPP and JDV carried out the physiological experiments; AG carried out the computational analyses; SR and JDV provided funding; AG, JDV, KPP, and SR wrote the initial draft; all authors contributed to the final draft.

Competing interest

The authors declare no competing interest.

References

- [1] Alyssa Accomando, Carlos Vargas-Irwin, and James Simmons. Neural spike train similarity algorithm detects differences in temporal patterning of bat echolocation call sequences. In *Proceedings of Meetings on Acoustics 173EAA*, volume 30, page 010002. Acoustical Society of America, 2017.
- [2] Dmitriy Aronov. Fast algorithm for the metric-space analysis of simultaneous responses of multiple single neurons. *Journal of neuroscience methods*, 124(2):175–179, 2003.
- [3] Dmitriy Aronov, Daniel S Reich, Ferenc Mechler, and Jonathan D Victor. Neural coding of spatial phase in V1 of the macaque monkey. *Journal of Neurophysiology*, 89(6):3304–3327, 2003.
- [4] Dmitriy Aronov and Jonathan D Victor. Non-Euclidean properties of spike train metric spaces. *Physical Review E*, 69(6):061905, 2004.
- [5] Jean-Baptiste Bardin, Gard Spreemann, and Kathryn Hess. Topological exploration of artificial neuronal network dynamics. *Network Neuroscience*, 3(3):725–743, 2019.
- [6] Ulrich Bauer. Ripser: efficient computation of Vietoris–Rips persistence barcodes. *Journal of Applied and Computational Topology*, pages 1–33, 2021.

- [7] Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.
- [8] Carina Curto and Vladimir Itskov. Cell groups reveal structure of stimulus space. *PLoS Computational Biology*, 4(10):e1000205, 2008.
- [9] Yuri Dabaghian, Facundo Mémoli, Loren Frank, and Gunnar Carlsson. A topological paradigm for hippocampal spatial map formation using persistent homology. *PLoS Computational Biology*, 2012.
- [10] Patricia M Di Lorenzo and Jonathan D Victor. Taste response variability and temporal coding in the nucleus of the solitary tract of the rat. *Journal of Neurophysiology*, 90(3):1418–1431, 2003.
- [11] Alexander S Ecker, Philipp Berens, R James Cotton, Manivannan Subramanian, George H Denfield, Cathryn R Cadwell, Stelios M Smirnakis, Matthias Bethge, and Andreas S Tolias. State dependence of noise correlations in macaque primary visual cortex. *Neuron*, 82(1):235–248, 2014.
- [12] Luca A Finelli, Seth Haney, Maxim Bazhenov, Mark Stopfer, and Terrence J Sejnowski. Synaptic learning rules and sparse coding in a model sensory system. *PLoS Computational Biology*, 4(4):e1000062, 2008.
- [13] Peiran Gao and Surya Ganguli. On simplicity and complexity in the brave new world of large-scale neuroscience. *Current opinion in neurobiology*, 32:148–155, 2015.
- [14] Timothy J Gawne and Barry J Richmond. How independent are the messages carried by adjacent inferior temporal cortical neurons? *Journal of Neuroscience*, 13(7):2758–2771, 1993.
- [15] Chad Giusti, Eva Pastalkova, Carina Curto, and Vladimir Itskov. Clique topology reveals intrinsic geometric structure in neural correlations. *Proceedings of the National Academy of Sciences*, 112(44):13455–13460, 2015.
- [16] Ann M Hermundstad, John J Briguglio, Mary M Conte, Jonathan D Victor, Vijay Balasubramanian, and Gašper Tkačik. Variance predicts salience in central sensory processing. *eLife*, 3:e03722, 2014.
- [17] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. Hyperbolic geometry of complex networks. *Physical Review E*, 82(3):036106, 2010.
- [18] Giovanni Petri, Martina Scolamiero, Irene Donato, and Francesco Vaccarino. Topological strata of weighted complex networks. *PLoS one*, 8(6):e66506, 2013.
- [19] Anita M Schmid, Keith P Purpura, and Jonathan D Victor. Responses to orientation discontinuities in V1 and V2: physiological dissociations and functional implications. *Journal of Neuroscience*, 34(10):3559–3578, 2014.
- [20] Gurjeet Singh, Facundo Memoli, Tigran Ishkhanov, Guillermo Sapiro, Gunnar Carlsson, and Dario L Ringach. Topological analysis of population activity in visual cortex. *Journal of vision*, 8(8):11–11, 2008.
- [21] Ann E Sizemore, Jennifer E Phillips-Cremens, Robert Ghrist, and Danielle S Bassett. The importance of the whole: topological data analysis for the network neuroscientist. *Network Neuroscience*, 3(3):656–673, 2019.
- [22] Tiberiu Tesileanu, Mary M Conte, John J Briguglio, Ann M Hermundstad, Jonathan D Victor, and Vijay Balasubramanian. Efficient coding of natural scene statistics predicts discrimination thresholds for grayscale textures. *eLife*, 9:e54347, 2020.

- [23] Jonathan D Victor. Spike train metrics. *Current Opinion in Neurobiology*, 15(5):585–592, 2005.
- [24] Jonathan D Victor and Mary M Conte. Local image statistics: maximum-entropy constructions and perceptual salience. *JOSA A*, 29(7):1313–1345, 2012.
- [25] Jonathan D Victor and Keith P Purpura. Nature and precision of temporal coding in visual cortex: a metric-space analysis. *Journal of Neurophysiology*, 76(2):1310–1326, 1996.
- [26] Jonathan D Victor and Keith P Purpura. Metric-space analysis of spike trains: theory, algorithms and application. *Network: Computation in Neural Systems*, 8(2):127–164, 1997.
- [27] Jonathan D Victor, Daniel J Thengone, Syed M Rizvi, and Mary M Conte. A perceptual space of local image statistics. *Vision research*, 117:117–135, 2015.
- [28] Ryan C Williamson, Benjamin R Cowley, Ashok Litwin-Kumar, Brent Doiron, Adam Kohn, Matthew A Smith, and Byron M Yu. Scaling properties of dimensionality reduction for neural populations and network models. *PLoS Computational Biology*, 12(12):e1005141, 2016.
- [29] Yunguo Yu, Anita M Schmid, and Jonathan D Victor. Visual processing of informative multipoint correlations arises primarily in V2. *eLife*, 4:e06604, 2015.
- [30] Yuansheng Zhou, Brian H Smith, and Tatyana O Sharpee. Hyperbolic geometry of the olfactory space. *Science Advances*, 4(8):eaq1458, 2018.

Figures

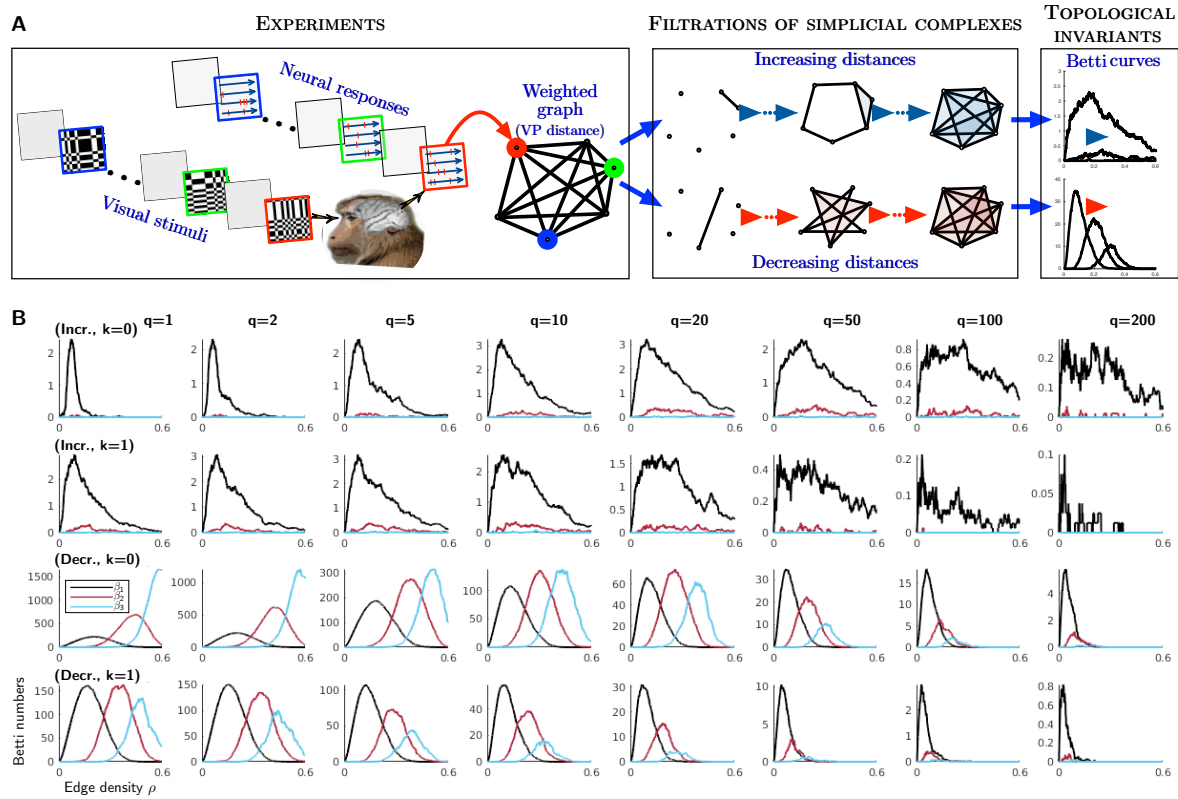


Figure 1: A. Pipeline for extracting topological summaries from the experimental data. Left: For each data set, neuronal responses (single-unit or up to four simultaneously-recorded units) are grouped into collections of 64 responses of duration 320 ms each. Each collection consists of responses to 64 different examples of textures defined by one of 10 texture parameters. The responses in each collection become points (colored balls) in a state space, whose distances, summarized by the weighted graph, are determined by the Victor-Purpura metric. Center: The weighted graph is associated with a filtration of simplicial complexes, using the clique topology method, by successively adding edges of the graph according to their weights, either increasingly or decreasingly. Right: Betti curves are computed from each filtration. Betti curves are a topological descriptor that captures how the number of topological voids of different dimensions (holes, “bubbles”, etc.) depends on the proportion of added edges in the filtration process. **B. Average Betti curves of one dataset.** The average Betti curves for β_1 - β_3 over the 80 collections of one dataset (L7301TT6), containing recordings of spiking activity of 4 neurons in layer 5 of area V2, are displayed for increasing and decreasing filtrations, and for each value of the parameters q (time scale, sec^{-1}) and k (sensitivity to neuron of origin) of the Victor-Purpura distance. The two top (respectively, bottom) rows correspond to the increasing (resp., decreasing) filtration method. The odd (resp., even) rows correspond to $k = 0$ (resp., $k = 1$). In the increasing filtration plots, the Betti curves for β_3 are close to zero.

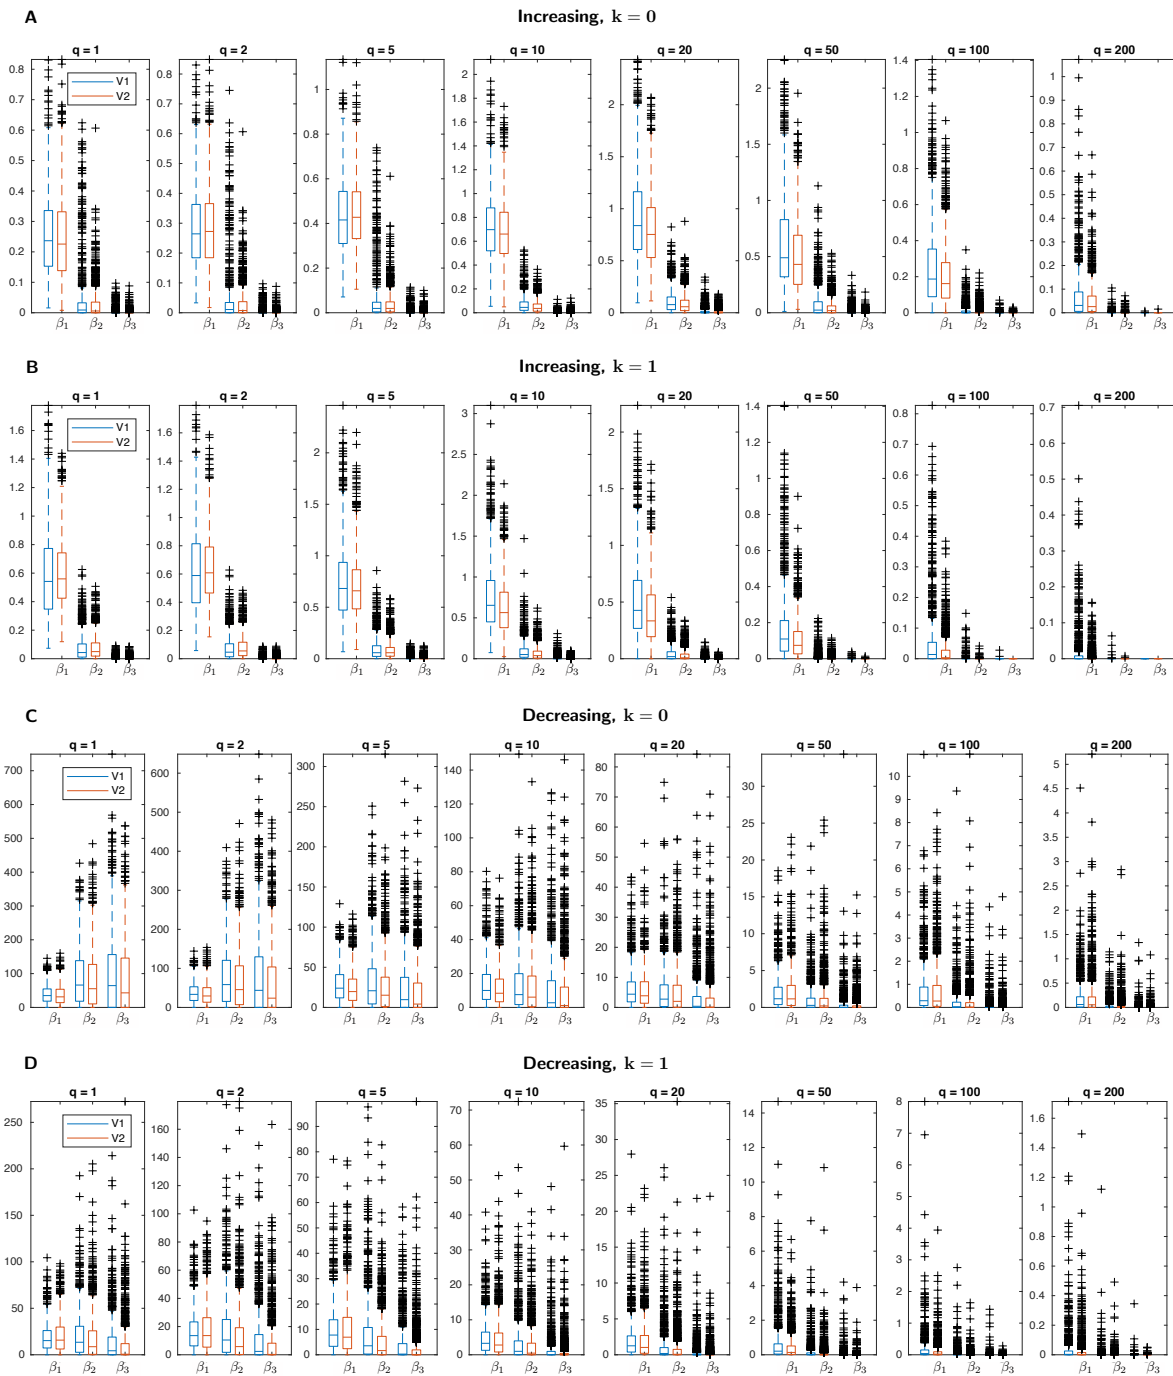


Figure 2: Comparison of integrated Betti values for data recorded in V1 vs. V2. Integrated Betti values (i.e., the values of the integrals of the Betti curves) for β_1 - β_3 of the experimental data are shown divided into two groups, according to whether the data were recorded in area V1 (blue) or V2 (red) of the visual cortex. To generate the distributions of integrated Betti values displayed in the figure, individual (non-averaged over the dataset) Betti curves of each collection of responses are considered. The four panels are for increasing (**A,B**) and decreasing (**C,D**) filtrations, and for $k = 0$ (**A,C**) and $k = 1$ (**B,D**). Each panel shows the distribution of integrated Betti values for all values of the timescale parameter q of the Victor-Purpura distance.

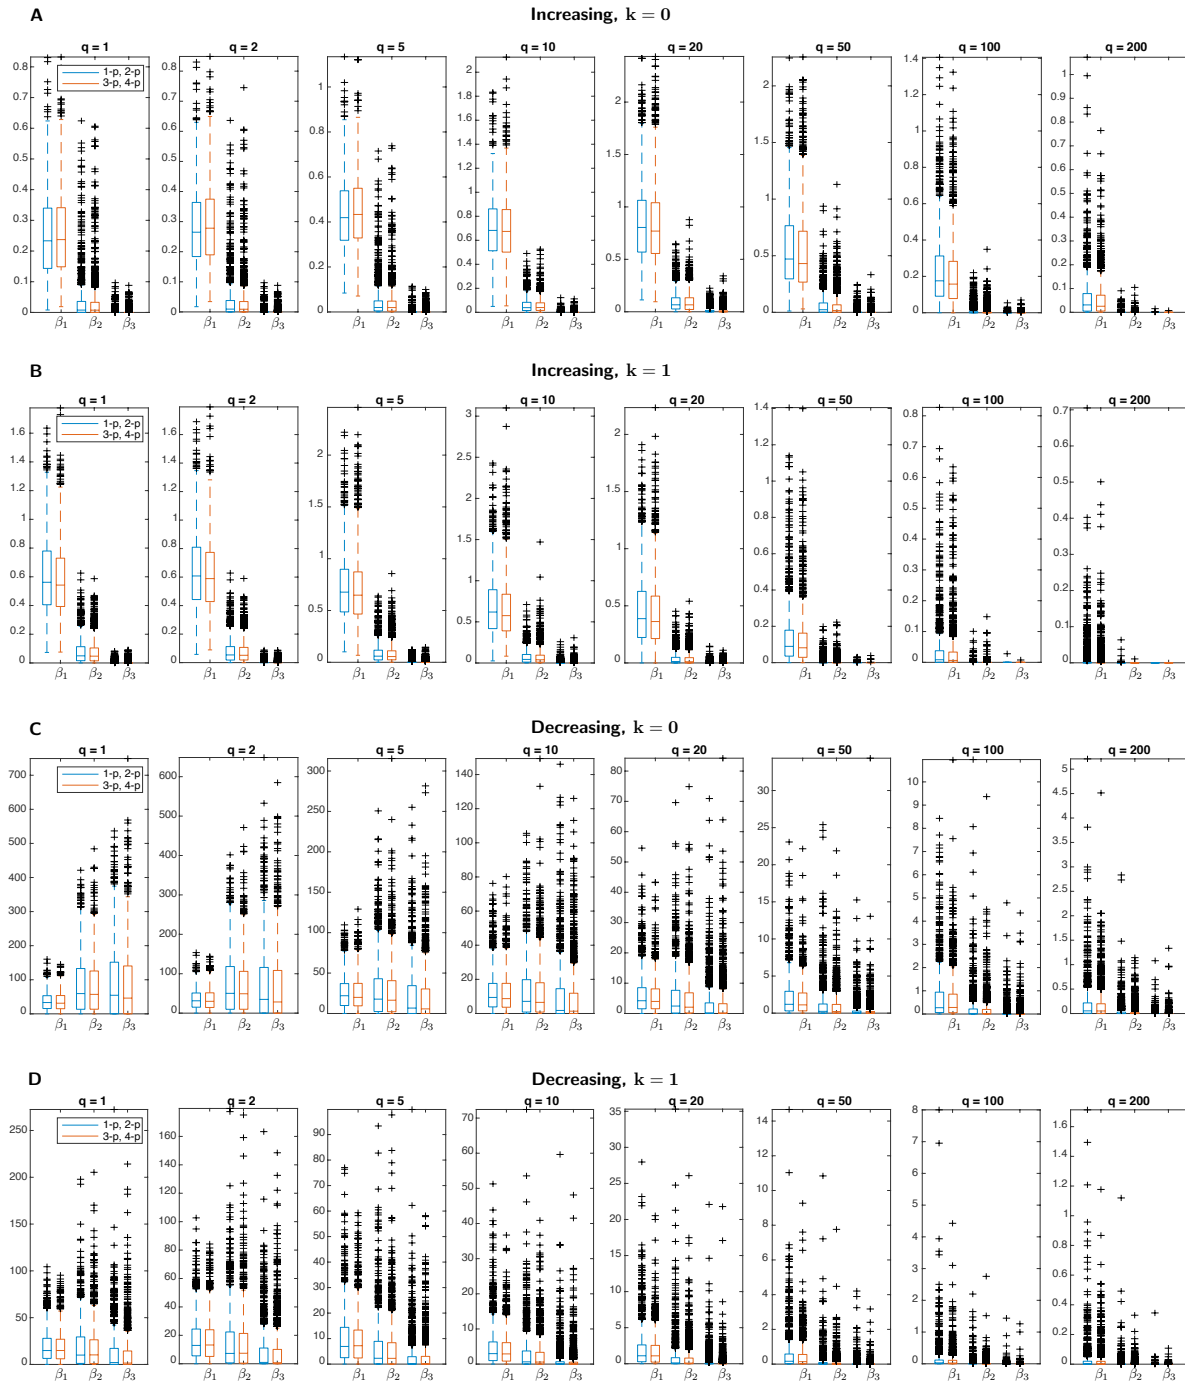


Figure 3: Comparison of integrated Betti values for responses elicited by low- and high-order spatial correlations. Integrated Betti values for β_1 - β_3 of the experimental data are shown divided into two groups, according to whether neuronal responses are driven by visual stimuli with one- and two-point correlations (axes γ , β_- , β_l , β_\setminus , β_l of the stimulus space described in [27]) or three- and four-points correlations (axes θ_\perp , θ_\setminus , θ_\setminus , θ_\setminus and α). The distributions of integrated Betti values, determined as in Fig. 2, are respectively shown in blue and red. High-order (three- and four-points) correlations are extracted primarily in V2 [29]; low-order (one- and two-points) correlations can be extracted by the spatial filtering of retinal processing. To generate the distribution of integrated Betti values, individual (non-averaged over the dataset) Betti curves of each collection of responses are considered. The four panels are for increasing (A,B) and decreasing (C,D) filtrations, and for $k = 0$ (A,C) and $k = 1$ (B,D). Each panel shows the distribution of integrated Betti values for all values of the timescale parameter q of the Victor-Purpura distance.

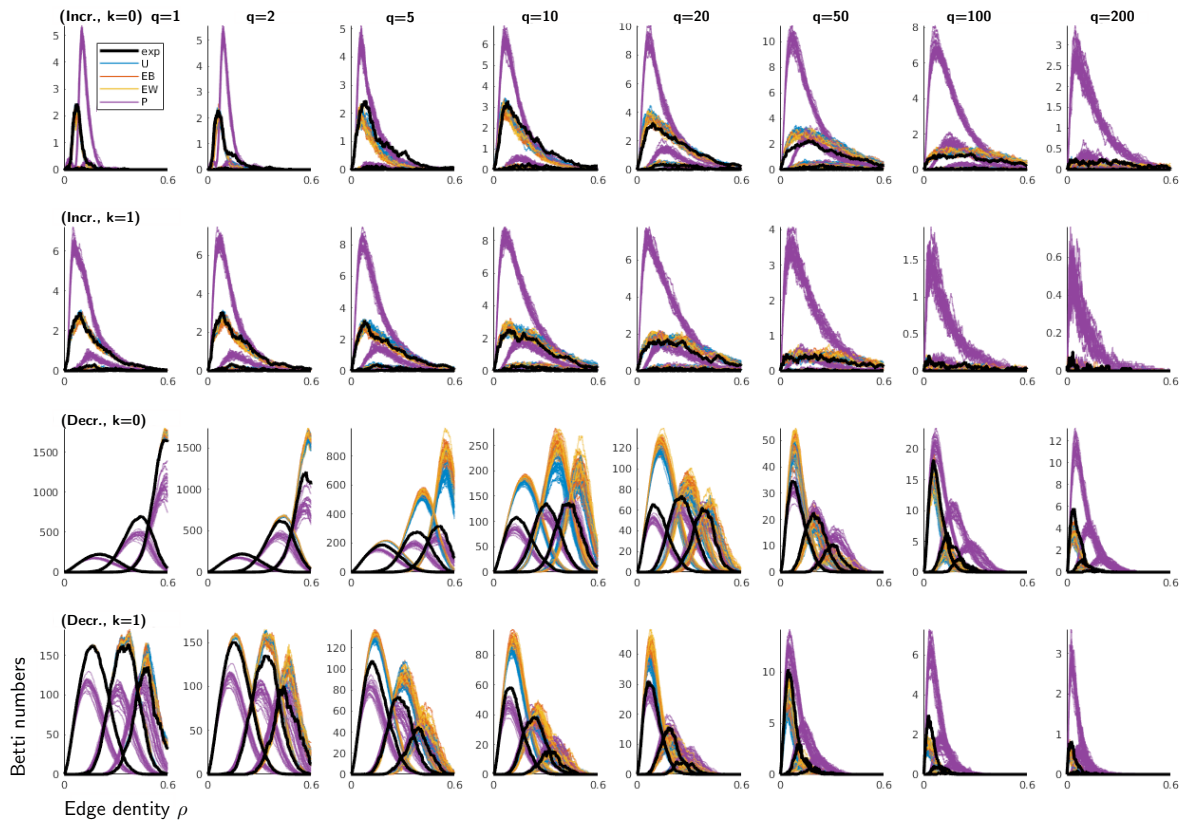


Figure 4: **Average Betti curves of one dataset, and comparison with surrogate data.** The figure superimposes on Fig. 1B the Betti curves of the 20 computations of the four types of surrogate data generated from the single dataset “L7301TT6”. The average Betti curves for β_1 - β_3 of the experimental data (black, “exp”) are highlighted with a thicker line. The average Betti curves for β_1 - β_3 of each computation of the surrogate data are drawn in thin color lines: uniform resampling of spike times (blue, “U”), exchange of spike times between collections (red, “EB”), exchange of spike times within collections (yellow, “EW”), Poisson generated spike data (purple, “P”).

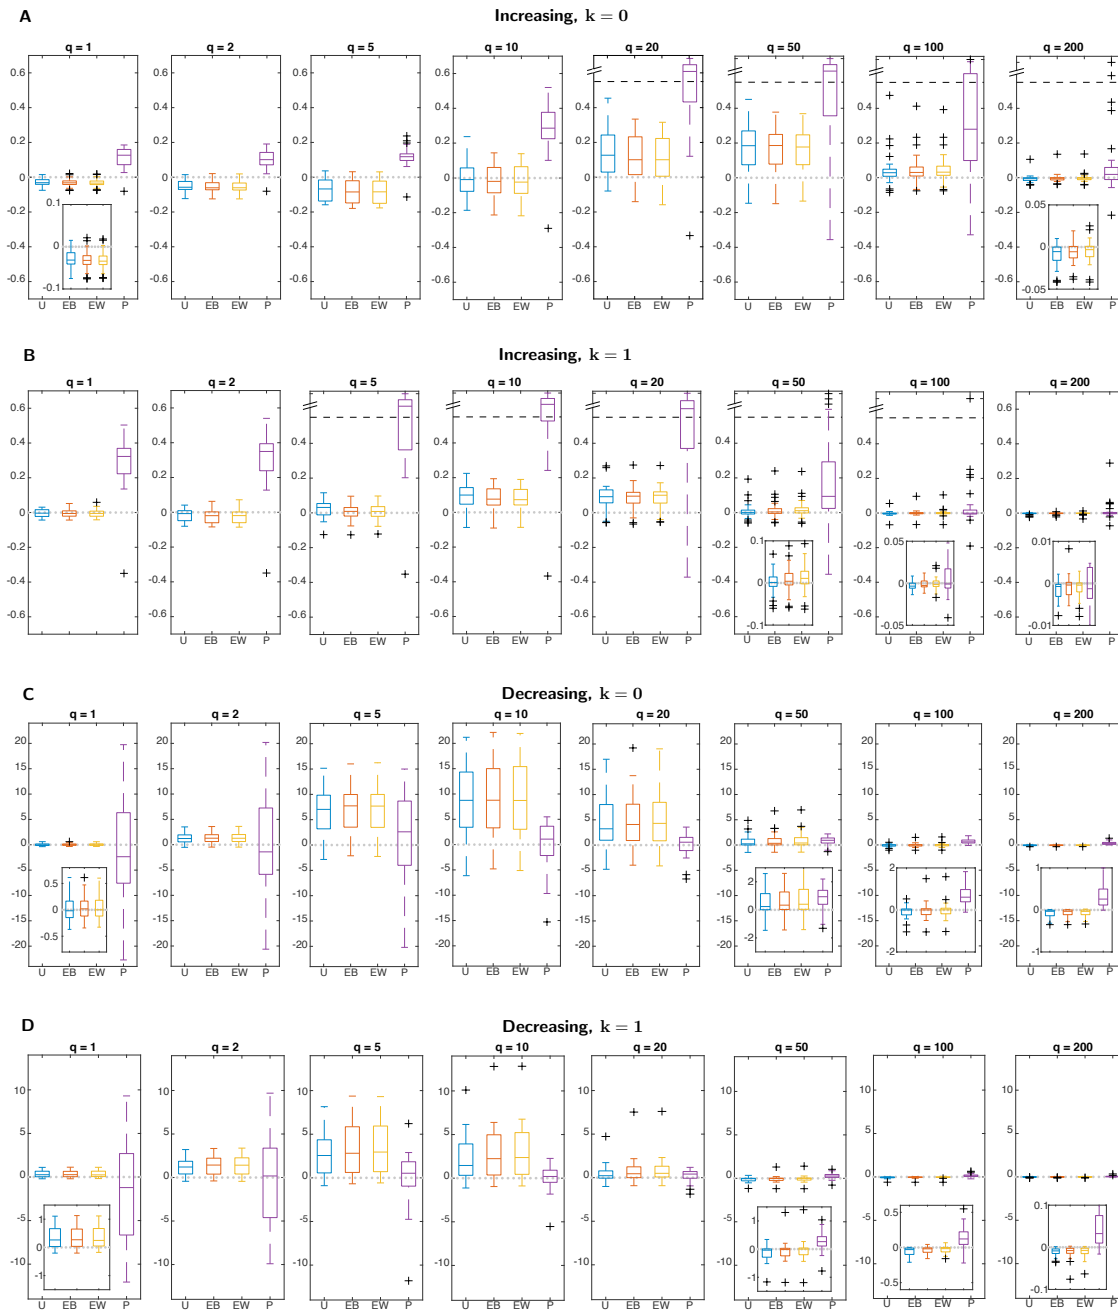


Figure 5: Distribution of the integrated Betti values for β_1 of neural data and surrogates. Distribution across the 28 datasets of the difference between the integrated Betti values for β_1 of each kind of surrogate, averaged over 20 examples, and the integrated Betti value of the experimental data. Each plot shows the four surrogates: uniform resampling of spike times (blue, “U”), exchange of spike times between collections (red, “EB”), exchange of spike times within collections (yellow, “EW”), Poisson generated spike data (purple, “P”). Insets zoom in on some boxplots at a smaller scale. When present, a dashed line bounds an area at the extreme of a plot beyond which data are shown on a compressed ordinate. The four panels are for increasing (**A,B**) and decreasing (**C,D**) filtrations, and for $k = 0$ (**A,C**) and $k = 1$ (**B,D**). For example, in **A**, the means of the integrated Betti values for the Poisson surrogates for all timescales, q are greater than the means of the integrated Betti values for the experimental spiking responses from the 28 datasets. Note that the deviation of the behavior of the surrogates is maximal for the range $q = 5 \text{ sec}^{-1}$ to $q = 50 \text{ sec}^{-1}$.

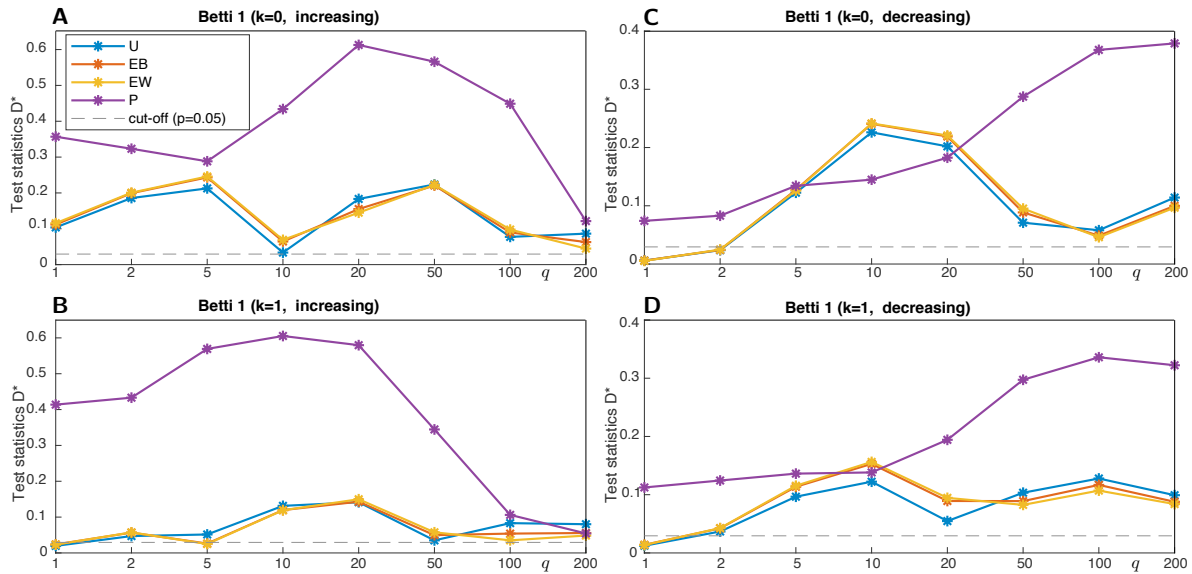


Figure 6: **Two-sample Kolmogorov-Smirnov test for the integrated Betti values for β_1 .** Two-sample Kolmogorov-Smirnov test statistic (ordinate) for comparison of the samples of integrated Betti values for β_1 of the experimental data with each type of surrogate data. The sample for the experimental data consists of the (28×80) integrated Betti values from all 80 collections from all 28 datasets. The samples for each surrogate consist of the $(20 \times 28 \times 80)$ integrated Betti values of the 20 computations. Values above the dashed line indicate $p < 0.05$ for rejection of the null hypothesis that the two samples come from the same distribution. The four panels are for increasing (A,B) and decreasing (C,D) filtrations, and for $k = 0$ (A,C) and $k = 1$ (B,D).

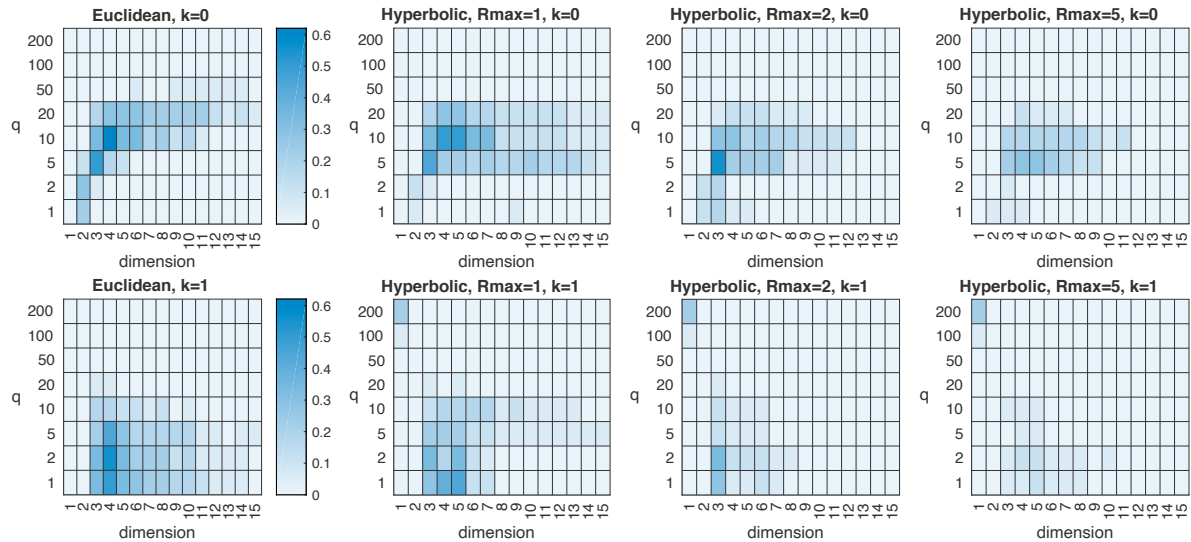


Figure 7: Compatibility of the experimental Betti curves with geometric models. The heatmaps show the fraction of Betti curves of 18 selected collections from different datasets which are compatible with the Euclidean (column 1) and hyperbolic models ($R_{\max} = 1, 2, 5$) (columns 2-4) of dimension $d = 1, \dots, 15$ (abscissas), for all values of the parameters q (ordinates) and k (rows) of the Victor-Purpura distance. The notion of compatibility we introduced requires the experimental integrated Betti values to be within 3 standard deviations of the mean of the 300 values of the model, for all Betti numbers β_1 - β_3 and for both increasing and decreasing filtrations. In the range $q = 5$ to 20 sec^{-1} , the greatest compatibility occurs for dimensions 3-5, and for the Euclidean or hyperbolic models with moderate curvature ($R_{\max} = 1$). As we determined by visual inspection of the Betti curves, the hotspot at dimension 1 and $q = 200$ in the hyperbolic heatmaps for $k = 1$ is an artifact due to the fact that the low dimension constrains the Betti curves of the hyperbolic models to being close to (identically) zero, hence compatible in some cases with the experimental Betti curves at the extreme value $q = 200$.