Non-Euclidean properties of spike train metric spaces

Dmitriy Aronov[1,2,*] and Jonathan D. Victor[2]

[1]*Department of Biological Sciences, Columbia University, 1002 Fairchild, Mail Code 2436, 1212 Amsterdam Avenue,
New York, New York 10027, USA*
[2]*Department of Neurology and Neuroscience, Weill Medical College of Cornell University, 1300 York Avenue,
New York, New York 10021, USA*

Quantifying the dissimilarity (or distance) between two sequences is essential to the study of action potential (spike) trains in neuroscience and genetic sequences in molecular biology. In neuroscience, traditional methods for sequence comparisons rely on techniques appropriate for multivariate data, which typically assume that the space of sequences is intrinsically Euclidean. More recently, metrics that do not make this assumption have been introduced for comparison of neural activity patterns. These metrics have a formal resemblance to those used in the comparison of genetic sequences. Yet the relationship between such metrics and the traditional Euclidean distances has remained unclear. We show, both analytically and computationally, that the geometries associated with metric spaces of event sequences are intrinsically non-Euclidean. Our results demonstrate that metric spaces enrich the study of neural activity patterns, since accounting for perceptual spaces requires a non-Euclidean geometry.

PACS number(s): 87.10.+e, 87.19.La, 02.40.Dr, 87.19.Bb

I. INTRODUCTION

Marked point processes are constructs widely used for representing sequences that contain multiple types of discrete events [1]. These constructs have at least two important applications in biology. In molecular biology, genetic sequences can be viewed as marked point processes on a discrete domain, in which each nucleotide or amino acid is viewed as an event of a particular type [2,3]. In neuroscience, neuronal population activity can be viewed as a marked point process on a continuous domain (time), with each action potential (spike) labeled according to the neuron that fired it [4,5]. In both settings, it is important to quantify the similarity (or dissimilarity) of two sequences. In molecular biology, dissimilarity can be considered the evolutionary distance between two sequences. In neuroscience, dissimilarity helps characterize neural variability and coding. A measure of dissimilarity between two sequences can be thought of as the distance between them in some abstract topological space, in which each sequence corresponds to a particular point. Provided that the notion of dissimilarity satisfies the triangle inequality (which it does, in the above applications), this topological space is a metric space, whose intrinsic properties are determined by the pairwise distances between all sequences.

The starting point for analysis of such sequences is typically an embedding into a high-dimensional vector space [6]. For neural activity, this embedding consists of discretizing the sequence into bins and representing the number of events of each type in each bin as a separate Euclidean coordinate. (For genetic sequences, there is no need for discretization.) These methods present both practical and theoretical prob-

lems [4,7]. The practical problem is that a very high-dimensional space is required, and thus experimental data typically represent only a very sparse sample. This makes it difficult to estimate quantities such as information and entropy in a rigorous fashion. A more fundamental issue is that the vector space representation forces the sequence space to have an intrinsically Euclidean geometry. The geometry of perceptual spaces, however, is typically not Euclidean [8,9]. Since, presumably, perceptual similarities are based on patterns of neural activity [8,10], it would be problematic if the intrinsic geometry of spike train sequences were required to be Euclidean.

As a solution to these problems, Victor and co-workers [4,5,11] proposed a metric space formalism for measuring dissimilarities of action potential sequences (spike trains). This formalism is closely related to the quantification of evolutionary distances defined by Sellers [2] for genetic sequences, which takes into account the local features of a sequence rather than the long-range structure [12]. For neural activity, these notions of dissimilarity take into account the number of spikes in a spike train, the timing of individual spikes, and, for multineuronal responses [5], the neuron of origin of each spike. Further, the model parametrizes the relative importance of these three aspects, allowing a continuum of different neural coding schemes to be analyzed. Distances provided by spike train metrics have been used for measuring the variability [13,14] and stimulus-dependent clustering [4,5,11,15–18] of neural responses in a number of sensory systems, as well as for investigating the coding of visual stimulus features by reconstructing neural response spaces [4,5]. The multineuronal metrics have also been used to investigate the mechanisms of joint coding by populations of neurons [5].

Although spike train metrics solve the above-mentioned practical issues, their mathematical relationship with distances based on a traditional vector space embedding of

---

*Corresponding author. Email address: da2006@columbia.edu

spike trains has remained unclear. *A priori*, metric-space distances between spike trains might be equivalent to Euclidean distances in some vector space, even if not the "obvious" one yielded by the natural embedding. For instance, a recently proposed alternative that is in some ways similar to the Victor and Purpura metric [19] is demonstrably Euclidean. If the metric-based distances were also equivalent to Euclidean distances, the metric space approaches would not provide for a richer set of geometries that appear to be necessary for the representation of perceptual spaces. Although empirical analyses demonstrate that metric spaces are superior to vector-based methods in tasks such as stimulus-dependent clustering [20], the lack of a mathematical relationship between the two methods has made the choice between them problematic.

Here, we present analytical and computational analyses demonstrating that spike train metrics and Euclidean distances are not equivalent under any transformation that preserves the metrics' essential properties. The theoretical importance of this result is that metric-space methods do more than merely address practical problems in spike train analysis. Rather, our result shows that the notion of dissimilarity underlying the metric-space analysis of spike trains and genetic sequence analysis is *fundamentally* richer and more general than that of ordinary Euclidean distances and is better suited for the study of sensory perception.

II. METHODS

Numerical calculations were performed with MATLAB and C programming languages.

III. RESULTS

Our goal is to establish the mathematical relationship between spike train metrics and Euclidean distances. Although sequence metrics are not restricted to neuroscience applications, we adopt the neuroscience terminology and refer to events as "spikes" and sequences of events as "spike trains." If a spike train contains discharges of multiple neurons, each spike is marked according to its neuron of origin, and we refer to the resulting sequence as a "labeled spike train." We first discuss the space of spike trains and some topological issues associated with it.

A. Spike train metrics

In general, a sequence of marked events can be considered a single point in an abstract topological space. If the geometric properties of this space represent certain meaningful aspects of the sequence, the space will provide a useful way to analyze experimental data. For instance, the dissimilarity of any two sequences of events, quantified in some way, can be represented as the distance between corresponding points in the abstract space. With this choice, clusters of points correspond to sequences that are similar, and separate clusters correspond to sequences that are readily discriminable.

The distances are defined by a function d that operates on pairs of sequences, which we denote $A, B, \ldots$. For the result-
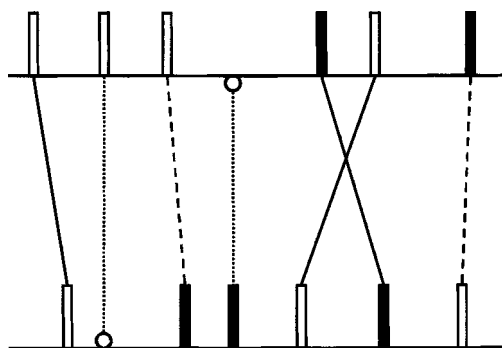

FIG. 1. A minimal transformation path between two spike trains that contain simultaneously recorded discharges of two neurons. Each spike is shaded according to its neuron of origin. Solid lines: shifts linking identically labeled spikes. Dashed lines: shifts accompanied by label reassignments linking differently labeled spikes. Dotted lines with circles: insertions and deletions of spikes.

ing space to be considered a metric space, d must satisfy (i) nondegeneracy: $d(A, A) = 0$, (ii) non-negativity: $d(A, B) \geq 0$, (iii) symmetry: $d(A, B) = d(B, A)$, and (iv) the triangle inequality: $d(A, C) \leq d(A, B) + d(B, C)$.

In the spike train metrics of Victor and Purpura [11], the distance between two spike trains is defined as the minimal "cost" of transforming one spike train into the other via a sequence of elementary steps. Two of the elementary steps consist of deleting or inserting a spike, and are assigned a cost of 1. The third elementary step consists of shifting a spike in time by an amount Δt, for a cost of $q|\Delta t|$. If spike trains contain discharges of multiple neurons, there is a fourth elementary step that consists of changing the label of a spike (i.e., the identity of the neuron that fired it) for a cost of k.

The two metric parameters, q and k, have a simple interpretation for neural coding. The parameter q is measured in $\sec^{-1}$ and quantifies the temporal precision relevant to spike timing. If $q = 0$, there is no cost for shifting a spike in time. In such a case, the metric ignores the timing of individual spikes and degenerates into merely the difference in the number of spikes. If $q > 0$, spikes in two trains can be potentially matched if they occur within $2/q$ of each other (if they are separated by a greater amount, then the cost of shifting one spike to coincide in time with the other is greater than the cost of deleting the two of them). The parameter k is dimensionless and quantifies the importance of distinguishing individual neurons. If $k = 0$, there is no cost for changing the label of a spike, and spikes originating from different neurons are not distinguished from each other. If $k \geq 2$, spikes from different neurons are never considered similar, since deleting two spikes for a cost of 2 is not more expensive than matching their labels. For $0 < k < 2$, differently labeled spikes can be potentially matched if they occur within $(2-k)/q$ of each other. A sample transformation of one labeled spike train into another is illustrated in Fig. 1.

We wish to determine whether the distances defined above are in any way equivalent to a Euclidean distance. Thus, we ask whether any set of spike trains in a metric
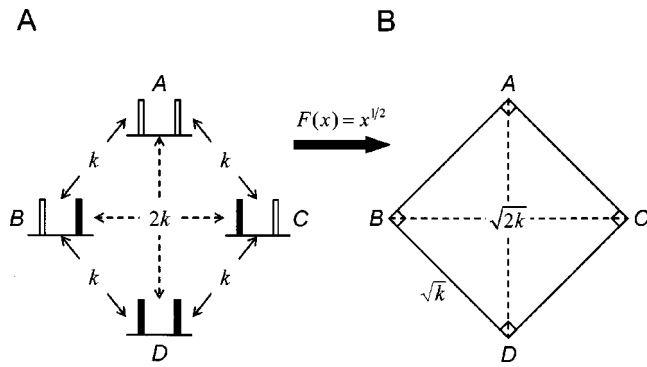
FIG. 2. Sample set of spike trains that cannot be embedded into a Euclidean space in a way that preserves all pairwise metric space distances. (a) The four spike trains contain simultaneously recorded discharges of two neurons, and each spike is shaded according to its neuron of origin. Spikes in each spike train are separated by at least $2/q$ to prevent shifting of spikes in a minimal transformation path. The pairwise metric distances (arrows) are inconsistent with Euclidean geometry. (b) The four spike trains can be mapped to vertices of a square in a Euclidean space after all pairwise distances are transformed with $F(x)=x^{1/2}$.

space can be embedded into a set of points in a Euclidean space of some dimension in a way that preserves all pairwise distances. It is not hard to see that embeddings preserving *exact* distances are generally impossible. For instance, a set of four labeled spike trains that demonstrates this is illustrated in Fig. 2(a). A and D are at a distance $2k$ from each other, while B and C are both at a distance k from A and from D. In Euclidean geometry, both B and C would have to be midpoints of the line segment connecting A and D, and thus be the same. Therefore, the pairwise distances between these spike trains are inconsistent with the uniqueness of a midpoint, which is one of the properties of Euclidean geometry.

This example shows that there is no embedding ϕ in which the metric-space distance $d(A,B)$ is equal to the vector-space distance $|\phi(A)-\phi(B)|$. Thus, a more interesting question is whether there is an embedding preserving the distances up to some transformation that preserves certain crucial properties of the metric (in a sense to be made precise below). That is, we ask whether there is a transformation $F(x)$ and a vector-space embedding ϕ, such that $F[d(A,B)]=|\phi(A)-\phi(B)|$. For instance, a transformation of distances in Fig. 2(a) with $F(x)=x^{1/2}$ maps the four points into vertices of a square with a side of length $\sqrt{k}$ [Fig. 2(b)]. At least for this set of spike trains, the transformation $F(x)=x^{1/2}$ evidently provides a meaningful representation of the original metric. However, it is not clear whether this approach will work more generally. Note that the embedding problem we address is distinct from the question that is often asked about metrics, namely whether *approximate* embeddings are possible in a finite dimensional space [21].

We now state our formal requirements for a transformation F that preserves the geometrical properties of a spike train metric. First, F must preserve the order of distances. That is,

$$F(b) > F(a) \quad \text{if } b > a. \tag{1}$$

Second, we require that the transformation F be scale-invariant. Specifically, two geometrically similar figures in the original metric must remain similar after the transformation. That is,

$$\frac{F(ca)}{F(cb)} = \frac{F(a)}{F(b)} \quad \text{for } c > 0. \tag{2}$$

This property is also necessary for a sensible transformation function. A function that is not scale-invariant may transform one shape in the original metric into multiple distinct shapes that depend on the unit of length used to specify the pairwise distances between the points. However, since the unit of length already has a meaning in spike train metrics (the cost to insert or delete a single spike), we do not want to introduce a second scale via the transformation F. Scale invariance guarantees that a figure in the original metric is transformed into the same figure regardless of the unit of length used for F.

One can show that a function F satisfying Eqs. (1) and (2) must be in the form $F(x)=\alpha x^p$, where $\alpha>0$ and $0<p\leq 1$ in the following way. Equation (2) implies that the quantity $F(ca)/F(a)$ (for $c>0$) is independent of a. We denote this quantity by $G(c)$. We then consider the values $c=c_1$ and $a=c_2 b$ with Eq. (2). By rearranging the equation and dividing both sides by $F(a)$, we can then obtain $G(c_1 c_2)=G(c_1)G(c_2)$. With $H(c)=\ln G(\exp c)$, this property is equivalent to $H(a+b)=H(a)+H(b)$. This implies that $H(c)$ is a linear function $H(c)=pc+\beta$. Consequently, $G(c)$, and hence $F(x)$, is a power-law function, with exponent p. In order to satisfy Eq. (1), the exponent p must be greater than 0. Further, in order to preserve the triangle inequality necessary for distances to form a metric space, the inequality $F(a+b)\leq F(a)+F(b)$ or $(a+b)^p\leq a^p+b^p$ must hold for all non-negative distances a and b. This implies that p must be not greater than 1.

Had we allowed more general forms for F, there would likely be trivial solutions to the embedding problem. For example, transforming all nonzero distances to values very close to 1 would turn every shape into a nearly regular simplex, and thus allow an embedding in a Euclidean space (but one that destroys the original shape). Conversely, a scale-invariant transformation can provide a Euclidean embedding in at least some cases of the spike metric, such as the one in Fig. 2. Moreover, the power-law transformation $F(x)=x^{1/2}$ allows for a Euclidean embedding of points in simple Riemannian manifolds, such as the points on the circumference of a circle (with the distance between points defined as the arc length). Thus, our choice of the scale-invariance constraint (a) precludes trivial solutions to the embedding problem, but also (b) is sufficiently rich so that at least some non-Euclidean metrics can be transformed to a Euclidean one.

Therefore, we now ask whether there exists a transformation satisfying Eqs. (1) and (2) (i.e., a power-law transformation) that will allow for a distance-preserving embedding of *any* set of spike trains. Our first main result is that no such

function exists if there are at least three labels available for spikes. The proof consists of two components: (i) a rule for constructing sets of spikes trains whose pairwise distances form particularly simple and symmetric configurations, and (ii) a demonstration, taking advantage of these properties, that no transformation function can embed each of these sets into a Euclidean space. In view of the above considerations, this shows that the spike metrics are intrinsically non-Euclidean, in a sense that is much stronger than the well-known fact that generic Riemannian metrics need not be Euclidean. Our second result is strong computational evidence demonstrating that such an embedding function does not exist for sets of spike trains with fewer than three labels as well.

B. General criteria for nonembeddability

At most N dimensions are required to embed a set of $N+1$ points whose pairwise distances are consistent with a Euclidean geometry. If all pairwise distances are known, these points can be thought of as vertices of a simplex (i.e., triangle, tetrahedron, etc.) in N dimensions with known edge lengths. The content, or volume generalized to any number of dimensions, of this simplex can then be determined using the Cayley-Menger determinant [22],

$$V_N^2 = \frac{(-1)^{N+1}}{2^N(N!)^2}\det\begin{bmatrix} 0 & 1 & 1 & \cdots & 1 \\ 1 & 0 & d_{12}{}^2 & \cdots & d_{1,N+1}{}^2 \\ 1 & d_{21}{}^2 & 0 & \cdots & d_{2,N+1}{}^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & d_{N+1,1}{}^2 & d_{N+1,2}{}^2 & \cdots & 0 \end{bmatrix}$$

(3)

where d_{ij} is the distance between points i and j. In two dimensions ($N=2$), this equality reduces to the famous Heron's formula for the area of a triangle. If the right side of Eq. (3) is negative, V_N is not a real number, and the distances cannot correspond to edges of a simplex. Thus, the sign of the Cayley-Menger determinant of the pairwise distances can be used to identify nonembeddable sets of points.

C. Two-group arrangements of spike trains

We seek to find configurations of spike trains that are likely to be inconsistent with Euclidean geometry. One such configuration was shown in Fig. 2(a). There, the four spike trains can be thought of as two groups of points—one containing A and D and the other containing B and C. In this *two-group arrangement*, the distance between points from the same group ($2k$) is larger than the distance between any two points from different groups (k). In general, a set of four points arranged symmetrically in two groups can be embedded into a Euclidean space as a square if the ratio of these two distances is $\sqrt{2}$ [Fig. 2(b)]. If this ratio is smaller than $\sqrt{2}$, the four points can also be embedded, but the figure is no longer coplanar (and becomes a tetrahedron). For a ratio greater than $\sqrt{2}$, the points are not embeddable. Since the ratio of distances in Fig. 2(a) is $2k/k=2$, a transformation

function $F(x)=x^p$ with $p \leq 1/2$ is required to embed this set of spike trains.

Here, we develop a method to construct similar two-group arrangements with more than two spike trains per group. Because such arrangements are very symmetric, they can be easily studied using Cayley-Menger determinants. We use this method to show that such arrangements are indeed increasingly "harder" to embed into a Euclidean space as the number of spike trains per group increases. Specifically, we will construct a sequence of arrangements that require successively smaller transforming powers p. Since the limit of the sequence of powers is 0, we conclude that no transformation function [which requires $p>0$ to satisfy Eqs. (1) and (2)] will suffice for an embedding of any given set of spike trains.

Note that a limiting argument involving a sequence of increasingly larger sets of points is necessary to demonstrate nonembeddability rigorously. For any *finite* set of N spike trains, there is a smallest and a largest pairwise distance. Consequently, for a transforming power p sufficiently close to 0, the pth power of all of the pairwise distances will all be within a factor $1+\varepsilon$ of each other (for arbitrarily small ε). For sufficiently small ε, the transformed pairwise distances can therefore be recovered from an $(N-1)$-dimensional approximately regular simplex. Thus, any given set of metric-space distances can be embedded into a Euclidean space, but the number of dimensions necessary to embed them grows without a bound with the number of spike trains. This is fundamentally different from binning techniques, which force points into a Euclidean space of a fixed number of dimensions independently of how many spike trains are being analyzed. As we will show, no fixed number of dimensions, however, suffices to embed *any* set of metric-space distances.

Thus, to demonstrate nonembeddability, we will construct a sequence of sets of spike trains for which any fixed choice of $p>0$ will eventually fail. Each of these sets is a two-group arrangement analogous to that in Fig. 2. The two-group arrangements have an increasing number N of spike trains in each group. In each arrangement, the distance between any two points from the *same* group is a constant, ρ_N, and the distance between any two points from *different* groups is a second constant, ζ_N. As will be shown later, to guarantee nonembeddability, the ratio of distances ρ_N/ζ_N in these sets must be a constant greater than 1 for all N. The set in Fig. 2 is an example of such an arrangement with $N=2$ and $\rho_2/\zeta_2=2k/k>1$. To construct more constraining two-group arrangements, we concatenate spike trains with known pairwise distances to construct arrangements with more and longer spike trains.

Suppose $2M$ spike trains of equal length, numbered 0 through $2M-1$, form a two-group arrangement with one group containing spike trains $0,2,4,\ldots,2M-2$ and the other containing $1,3,5,\ldots,2M-1$. Suppose the distance between any two spike trains from the same group is v and the distance between any two spike trains from different groups is w. We use these spike trains as "tiles" and concatenate them in time to create larger trains. If individual tiles are separated in time by an interval of at least $2/q$, no spike is shifted from one tile to another in a minimal transformation path used to
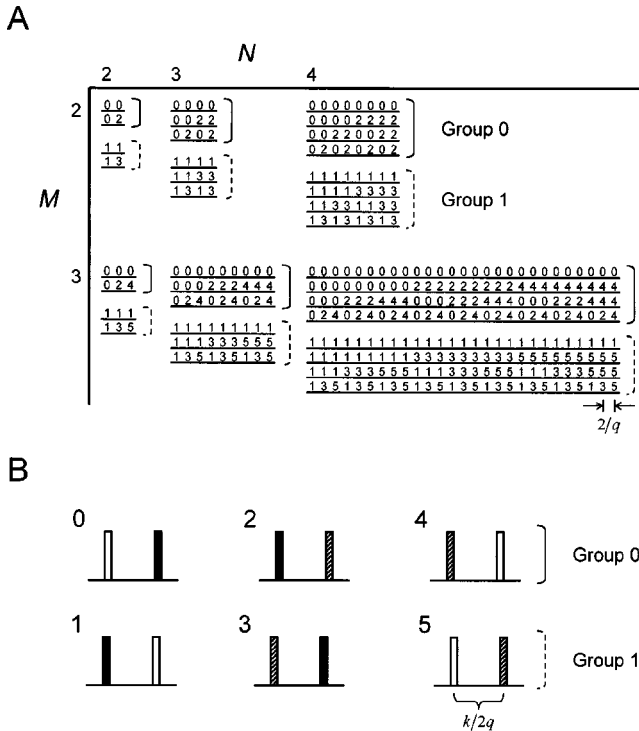
A

B

FIG. 3. Construction of spike trains that are not embeddable into Euclidean spaces. (a) Examples of concatenating short spike trains (tiles) to construct longer spike trains that form two-group arrangements. N is the number of spike trains per group and M is the number of tiles used to construct spike trains in each group. Digits indicate tile numbers. (b) Six tiles that are used in the text with $M=3$ to construct spike trains that are not embeddable into Euclidean spaces. Each spike is shaded according to its neuron of origin.

calculate the distance between two spike trains. Let $T(a_1, a_2, \ldots, a_n)$ denote a spike train formed by concatenating tiles $a_1, a_2, \ldots, a_n$. Then the distance between two spike trains is the sum of the distances between individual tiles,

$$d(T(a_1, a_2, \ldots, a_n), T(b_1, b_2, \ldots, b_n)) = \sum_{i=1}^{n} d(T(a_i), T(b_i)).$$

(4)

Now, consider constructing sets of spike trains by concatenating tiles in a way shown in Fig 3(a). The first row of each array always contains the 0 tile or the 1 tile. The remaining rows of each array constitute a table whose columns list all $(N-1)$-digit numbers that are composed of even digits in a base $2M$ system for one group and all numbers composed of odd digits for the other group. To construct $2N$ spike trains, M^{N-1} tiles are concatenated by reading across each array. Thus, the distance between two spike trains drawn from different groups is $\zeta_N = M^{N-1}w$, since spike trains from two groups contain no matching tiles. Two spike trains in the same group, however, are constructed such that exactly

$1/M$ of their tiles match. Therefore, the distance between them is $\rho_N = (M-1)M^{N-2}v$. The ratio of distances ρ_N/ζ_N is constant,

$$\frac{\rho_N}{\zeta_N} = \frac{(M-1)v}{Mw}.$$

(5)

If $v/w > M/(M-1)$, the ratio ρ_N/ζ_N is greater than 1, which, as we will show below, precludes embedding. For three labels ($M=3$), this condition is $v/w > 3/2$.

As we now show, the six short spike trains in Fig. 3(b) satisfy this requirement. In Fig. 3(b), each tile consists of two spikes with a time separation of $\Delta t = k/2q$ (recall that q is the cost per unit time to translate a spike in time, and k is the cost to change the label of a spike). We consider first spike trains from within the same group ($\{0, 2, 4,\}$ or $\{1, 3, 5\}$). If $0 < k \le 4/3$, then the minimal transformation path consists of either changing the labels of two spikes, for a cost of $2k$, or shifting two spikes in time and then changing the label of one of them, also for a cost of $2q\Delta t + k = 2k$. If $4/3 < k < 2$, then the minimal path consists of matching the two identically labeled spikes by shifting one of them and deleting the other two spikes, for a cost of $q\Delta t + 2 = k/2 + 2$. Thus, $v = 2k$ for $0 < k \le 4/3$ and $v = k/2 + 2$ for $4/3 < k < 2$. Now consider distances between a spike train in one group ($\{0, 2, 4\}$) and a spike train in the other group ($\{1, 3, 5\}$). The transformation path between two tiles from different groups depends on the pair of tiles in question. The transformation paths between tiles 0 and 1, 2 and 3, and 4 and 5 consist of swapping the positions of both spikes, for a cost of $2q\Delta t = k$. The transformation paths for all other pairs of tiles from different groups consist of changing the label of one spike, also for a cost of k. Thus, $w = k$, and the ratio v/w is

$$v/w = \begin{cases} 2, & 0 < k \le \dfrac{4}{3} \\ \dfrac{1}{2} + \dfrac{2}{k}, & \dfrac{4}{3} < k < 2 \end{cases}$$

(6)

and for all $0 < k < 2$, the ratio v/w is greater than $3/2$. It follows from Eq. (5) that concatenation of the tiles as shown in Fig. 3(a) leads to two-group arrangements (for any even number $2N$ of spike trains) for which the ratio ρ_N/ζ_N for all of these sets is a constant greater than 1,

$$\rho_N/\zeta_N = \begin{cases} \dfrac{4}{3}, & 0 < k \le \dfrac{4}{3} \\ \dfrac{1}{3} + \dfrac{4}{3k}, & \dfrac{4}{3} < k < 2. \end{cases}$$

(7)

D. Proof of nonembeddability

We next use Cayley-Menger determinants to prove that no transformation F permits a distance-preserving embedding of these configurations in a Euclidean space. We consider the matrix M_N of pairwise distances between the $2N$ spike trains. The matrix M_N contains pairwise distances between $2N$ points arranged in two groups—one containing spike trains $\{0, 2, 4, \ldots, 2N-2\}$ and the second containing spike trains

$\{1, 3, 5, \ldots, 2N-1\}$. The distance between any two points from the same group is ρ_N, while the distance between any two points from different groups is ζ_N. Further, the ratio of these two distances, ρ_N/ζ_N, is independent of N and is equal to a constant $\rho/\zeta > 1$. Thus, $\rho_N = \alpha_N \rho$ and $\zeta_N = \alpha_N \zeta$, where α_N is a series of positive coefficients. It follows that the Cayley-Menger determinant [Eq. (3)] for $2N$ points is given by

$$V_{2N-1}^2 = \frac{(-1)^{2N}}{2^{2N-1}[(2N-1)!]^2}(\alpha_N)^{4N-4}\det A_N, \qquad (8)$$

, where A_N is a $(2N+1)$-by-$(2N+1)$ matrix,

$$A_N = \begin{bmatrix} 0 & 1 & 1 & 1 & 1 & \ldots & 1 & 1 \\ 1 & 0 & s & r & s & \ldots & r & s \\ 1 & s & 0 & s & r & \ldots & s & r \\ 1 & r & s & 0 & s & \ldots & r & s \\ 1 & s & r & s & 0 & \ldots & s & r \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & r & s & r & s & \ldots & 0 & s \\ 1 & s & r & s & r & \ldots & s & 0 \end{bmatrix}. \qquad (9)$$

Here, $r = \rho^2$ and $s = \zeta^2$. [To relate Eq. (3) to Eq. (9), α_N^2 has been factored out of every entry in the determinant and then factored into the first row and the first column.] The coefficient of the determinant in Eq. (8) is always a positive number. Thus, the arrangement of $2N$ points is not embeddable into a Euclidean space if $\det A_N < 0$. As shown in the Appendix, the determinant of A_N evaluates to

$$\det A_N = 2Nr^{2N-2}[Ns - (N-1)r]. \qquad (10)$$

Combined with Eq. (10), the embedding condition $\det A_N \geq 0$ becomes (in the absence of a power-law transformation)

$$r/s = (\rho/\zeta)^2 \leq N/(N-1). \qquad (11)$$

If a power-law transformation $F(x) = x^p$ is applied to all distances, the embedding condition becomes

$$(\rho/\zeta)^{2p} \leq N/(N-1). \qquad (12)$$

Since our construction of two-group configurations is valid for arbitrarily large N, and $N/(N-1)$ approaches 1 as N increases without bound, this inequality cannot be satisfied for sufficiently high N regardless of the value of p. Thus, any transformation function F that satisfies Eqs. (1) and (2) fails to allow Euclidean embeddings of two-group arrangements that contain sufficiently many points.

E. Computational evidence that spike trains with fewer than three labels cannot be embedded

The proof provided above requires that at least three labels be available for spikes, and thus it leaves open the question of whether a similar result holds for the metrics when there are only one or two spike labels. It also raises the
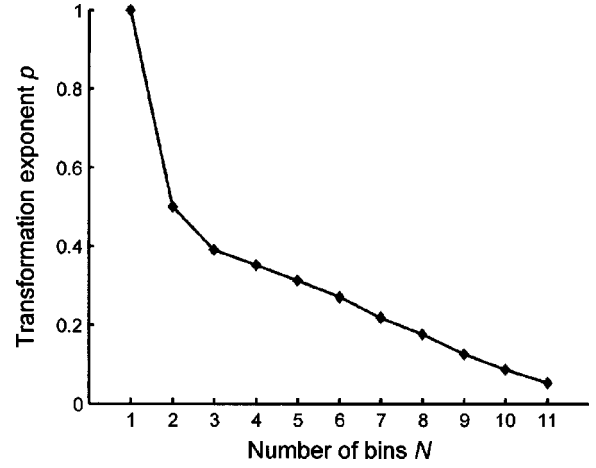

FIG. 4. Computational evidence that metric spaces are non-Euclidean. For each number of bins N, a set of spike trains S_N is constructed according to the scheme presented in the text. Exponent p is the largest value between 0 and 1, such that the transformation function $F(x) = x^p$ embeds S_N into a Euclidean space.

possibility that embedding fails only for very special configurations. Here we provide computational evidence suggesting that the failure of embedding is generic, in that (a) it applies to spike trains with fewer than three labels, and (b) spike trains do not need to be "hand-crafted." In particular, we find configurations of spike trains with only one label for which meaningful embeddings into a Euclidean space are not possible.

Our approach is as follows. We consider the set S_N of 2^N distinct spike trains whose spikes occur only at integer times from 0 to $N-1$ sec. We then calculate all pairwise metric distances between these spike trains for values of the parameter q in the range of $0-2$ sec^{-1}, sampled at intervals of 0.1 sec^{-1}. (We do not consider values of $q > 2$ sec^{-1} since in this range, spikes that are 1 sec apart are never shifted to coincide with each other in a minimal transformation path.) We then determine the largest p between 0 and 1, such that the function $F(x) = x^p$ applied to all pairwise distances embeds the set S_N into a Euclidean space. The largest p appears to become arbitrarily small as N becomes large (Fig. 4). Thus, it seems that the behavior that we have rigorously demonstrated for two-group configurations of spike trains with three labels also applies to generic spike trains with only one label. While we have not proven that this is the case, this computational result suggests that the conclusion about nonembeddability is true regardless of the number of labels.

IV. DISCUSSION

We have shown that spike train metric spaces are not equivalent to Euclidean spaces under any transformation that preserves their basic metrical properties. This is a rather surprising result, in view of the three limiting cases below. The first limiting case is that of $q = 0$ and a single label. In this case, the metric merely compares the number of spikes in

each spike train, and the distance between two spike trains is the difference of their spike counts. This is trivially equivalent to a Euclidean distance, since each train can be mapped into a one-dimensional vector space with the coordinate corresponding to the number of spikes it contains.

The metric also becomes equivalent to a Euclidean distance when q is infinite, but in a less trivial way. When q is infinite, the distance between two spike trains is proportional to the number of spikes that do not occur at exactly the same times in both trains (since each of these spikes must be deleted). For spike trains in discrete time, this distance is identical to the Hamming distance for binary sequences, which counts the number of positions at which two sequences differ. These distances may be recovered from the Euclidean distance in a vector space as follows. Spike trains in which there are N possible positions for spikes are mapped into an N-dimensional space with binary coordinates, indicating whether or not a spike is present. The spike metric distance is equal to the number of locations at which two spike trains differ. By the Pythagorean rule, the Euclidean distance is equal to the square root of the number of locations at which two spike trains differ. Thus, the spike metric distance is equivalent to the Euclidean one under the transformation $F(x)=x^{1/2}$.

It might appear from the above limiting cases that the metrics are Euclidean because their distances are not influenced by the shifting of spikes. However, shifting alone does not make metrics non-Euclidean. This can be seen from the fact that spike train metrics are *locally* Euclidean, even for $0<q<\infty$. Consider a ball of radius 1/2 around a spike train A, consisting of spikes at times $a_1,a_2,\ldots,a_N$. (By definition, this is the set of all spike trains whose distance from A is less than 1/2.) Since deleting or inserting a spike from A corresponds to a distance of 1, the spikes within the ball must each contain exactly N spikes, each of which occur at a time close to the time of one of the spikes in A. That is, each spike train within the ball contains spikes at times $a_1+\delta_1$, $a_2+\delta_2,\ldots,a_N+\delta_N$, where $\Sigma q|\delta_i|<1/2$. The transformation path between any two spike trains within the ball consists only of shifting spikes to match them in time. The distances are again equivalent to a Euclidean distance under the transformation $F(x)=x^{1/2}$ if each spike train is mapped into an N-dimensional space with the ith coordinate corresponding to δ_i.

Thus, deletions and insertions of spikes *combined* with shifting are responsible for the non-Euclidean properties of spike train metrics. This is particularly interesting, since these elementary transformations model a mechanism that detects, with a particular temporal precision, the coincidence of spikes from different spike trains. In such a neural decoding mechanism, only spikes located within a certain period of time from each other are considered coincident. A spike shifted continuously in time can change discretely from being considered coincident with one spike to being coincident with another. Such nonlinearities in coincidence detection systems are thought to be functionally relevant for neural processes such as the coding of retinal disparity in the visual system [23], interaural disparity in the auditory system [24], and, in general, the processing of inputs by a cortical neuron [25–28]. In all these cases, the nervous system must associ-

ate inputs arriving from different regions, such as different eyes, ears, or cortical layers. Coincidence detection is a way for such association to occur on adequately short time scales—a task that is hard to accomplish with most encoding schemes. Our results suggest that the presence of coincidence detectors produces non-Euclidean geometries of neural response spaces, which cannot be described by traditional vector-space models.

Interestingly, a non-Euclidean representation in itself may be crucial for the coding of sensory stimuli in the nervous system. As pointed out by Hopfield [8], there is an intrinsic problem with recognizing stimuli that are represented by multidimensional Euclidean vectors. In order for such recognition to be independent of stimulus intensity, sensory neurons must perform Euclidean rescaling, which is a highly unnatural task for neural networks. In addition, even with rescaling, Euclidean representation is undesirable in that it loses sensitivity to minor vector components, even when these components are crucial for the recognition of a particular stimulus. Finally, Euclidean representation creates a problem for a system that attempts to split the recognition problem into subparts, because each part of an input vector must be scaled by the same factor to preserve stimulus identity. Hopfield [8] demonstrated that these problems can be solved by using action potential timing for encoding sensory stimuli and using coincidence detection for stimulus recognition. The present results demonstrate that a metric based on coincidence detection indeed guarantees a non-Euclidean geometry of the neural response space. Interestingly, the problems pointed out in [8] are independent of the sensory stimulus modality. This supports the studies that use spike train metrics for analyzing various senses, including vision [4,5,11,14,17], audition [16], olfaction [15], gustation [18], and electroreception [13].

The usage of metric spaces, rather than vector spaces, is computationally (and perhaps conceptually) burdensome. However, our results demonstrate that spike train metrics are fundamentally different from the traditional methods involving various forms of vector space embedding. This greater generality is both practically and theoretically relevant to the study of the neural basis of perception.

APPENDIX: THE DETERMINANT OF A_N

We calculate the determinant of the matrix A_N [Eq. (9)] from the product of its eigenvalues, including multiplicities. Let an eigenvalue be λ and the corresponding eigenvector be $\boldsymbol{x}=(x_0,x_1,\ldots,x_{2N})$. Then

$$A_N \mathbf{x} = \begin{bmatrix} x_1 + x_2 + \cdots + x_{2N} \\ x_0 + r(x_1 + x_3 + \cdots + x_{2N-1}) + s(x_2 + x_4 + \cdots + x_{2N}) - rx_1 \\ x_0 + r(x_2 + x_4 + \cdots + x_{2N}) + s(x_1 + x_3 + \cdots + x_{2N-1}) - rx_2 \\ \vdots \\ x_0 + r(x_1 + x_3 + \cdots + x_{2N-1}) + s(x_2 + x_4 + \cdots + x_{2N}) - rx_{2N-1} \\ x_0 + r(x_2 + x_4 + \cdots + x_{2N}) + s(x_1 + x_3 + \cdots + x_{2N-1}) - rx_{2N} \end{bmatrix} = \begin{bmatrix} \lambda x_0 \\ \lambda x_1 \\ \lambda x_2 \\ \vdots \\ \lambda x_{2N-1} \\ \lambda x_{2N} \end{bmatrix}. \tag{A1}$$

Identifying the corresponding entries of the two vectors in Eq. (A1) yields equations for the eigenvalues and eigenvectors,

$$\lambda x_0 = \sum_i x_i,$$

$$(r+\lambda)x_j = x_0 + r \sum_{\text{odd } i} x_i + s \sum_{\text{even } i>0} x_i \text{ for odd } j>0,$$

$$(r+\lambda)x_j = x_0 + r \sum_{\text{even } i>0} x_i + s \sum_{\text{odd } i} x_i \text{ for even } j>0.$$

$$\tag{A2}$$

From the last two lines of Eq. (A2),

$$(r+\lambda)x_1 = (r+\lambda)x_3 = \cdots = (r+\lambda)x_{2N-1},$$

$$(r+\lambda)x_2 = (r+\lambda)x_4 = \cdots = (r+\lambda)x_{2N}. \tag{A3}$$

One solution to these equations (for $N>1$) is $\lambda=-r$. Note that A_N+rI has only three distinct rows, and thus the multiplicity of this eigenvalue is at least $(2N+1)-3=2N-2$.

If $\lambda \neq -r$, then $x_i=x_{\text{odd}}$ for all odd i and $x_i=x_{\text{even}}$ for all even $i>0$. Therefore, the last two lines of Eq. (A2) can be written as

$$x_j = \frac{1}{r+\lambda}x_0 + \frac{Nr}{r+\lambda}x_{\text{odd}} + \frac{Ns}{r+\lambda}x_{\text{even}} \text{ for odd } j,$$

$$x_j = \frac{1}{r+\lambda}x_0 + \frac{Nr}{r+\lambda}x_{\text{even}} + \frac{Ns}{r+\lambda}x_{\text{odd}} \text{ for even } j. \tag{A4}$$

Substituting these values into the first line of Eq. (A2) leads to

$$\lambda x_0 = \frac{2N}{r+\lambda}x_0 + \frac{Nr}{r+\lambda}Nx_{\text{odd}} + \frac{Nr}{r+\lambda}Nx_{\text{even}}$$
$$+ \frac{Ns}{r+\lambda}Nx_{\text{odd}} + \frac{Ns}{r+\lambda}Nx_{\text{even}}. \tag{A5}$$

Knowing that $\Sigma x_i = Nx_{\text{odd}} + Nx_{\text{even}}$ and using the first line of Eq. (A2) again,

$$\lambda x_0 = \frac{2N}{r+\lambda}x_0 + \frac{\lambda N(r+s)}{r+\lambda}x_0. \tag{A6}$$

One solution to this equation is $x_0=0$. From Eq. (A4) we know that $x_1=x_3=\cdots=x_{2N-1}$ and $x_2=x_4=\cdots=x_{2N}$. Combining this information with the first line of Eq. (A2), we find that the eigenvector corresponding to $x_0=0$ is $(0,1,-1,\ldots,1,-1)$. Multiplication of this vector by A_N shows that the corresponding eigenvalue is $(N-1)r-Ns$. If $x_0 \neq 0$, Eq. (A6) becomes a quadratic equation,

$$\lambda^2 - [(N-1)r + Ns]\lambda - 2N = 0. \tag{A7}$$

Thus, A_N has four distinct eigenvalues: $\lambda_1=-r$, $\lambda_2=(N-1)r-Ns$, and the roots λ_3 and λ_4 of the above quadratic equation. The coefficients in Eq. (A7) imply that $\lambda_3\lambda_4=-2N$.

A_N has a total of $2N+1$ eigenvalues. As noted above, λ_1 has a multiplicity of at least $2N-2$ times. Since λ_2,λ_3, and λ_4 are distinct (and also distinct from λ_1), we conclude that λ_1 has a multiplicity of exactly $2N-2$. The determinant of A_N can now be found by multiplying all eigenvalues together,

$$\det A_N = 2Nr^{2N-2}[Ns - (N-1)r]. \tag{A8}$$

[1] K. Sigman, *Stationary Marked Point Processes: An Intuitive Approach* (Chapman & Hall, New York, 1995).
[2] P. H. Sellers, SIAM (Soc. Ind. Appl. Math.) J. Appl. Math. **26**, 787 (1974).
[3] W. R. Pearson and W. Miller, Methods Enzymol. **210**, 575 (1992).
[4] J. D. Victor and K. P. Purpura, Network **8**, 127 (1997).
[5] D. Aronov *et al.*, J. Neurophysiol. **89**, 3304 (2003).
[6] K. Fukunaga, *Introduction to Statistical Pattern Recognition* (Academic, Boston, 1990).
[7] J. D. Victor, Phys. Rev. E **66**, 051903 (2002).
[8] J. J. Hopfield, Nature (London) **376**, 33 (1995).
[9] S. M. Wuerger, L. T. Maloney, and J. Krauskopf, Vision Res. **35**, 827 (1995).
[10] S. Edelman, Behav. Brain Sci. **21**, 449 (1998).
[11] J. D. Victor and K. P. Purpura, J. Neurophysiol. **76**, 1310 (1996).
[12] M. Li *et al.*, Bioinformatics **17**, 149 (2001).
[13] G. Kreiman *et al.*, J. Neurophysiol. **84**, 189 (2000).
[14] J. Keat *et al.*, Neuron **30**, 803 (2001).

[15] K. MacLeod, A. Bäcker, and G. Laurent, Nature (London) **395**, 693 (1998).

[16] C. K. Machens *et al.*, J. Neurosci. **21**, 3215 (2001).

[17] F. Mechler, D. S. Reich, and J. D. Victor, J. Neurosci. **22**, 6129 (2002).

[18] P. M. Di Lorenzo and J. D. Victor, J. Neurophysiol. (to be published).

[19] M. C. van Rossum, Neural Comput. **13**, 751 (2001).

[20] A. Bäcker, Doctoral dissertation, California Institute of Technology, 2002.

[21] P. Indyk and J. Matousek, in *Handbook of Discrete and Computational Geometry*, 2nd ed. edited by J. E. Goodman and J. O'Rouke (CRC, Boca Raton, FL, 2004).

[22] D. M. L. Y. Sommerville, *An Introduction to the Geometry of n Dimensions* (Dover, New York, 1958).

[23] C. W. Lee and B. A. Olshausen, Neural Comput. **8**, 545 (1996).

[24] T. Hattori and N. Suga, J. Comp. Physiol., A **180**, 271 (1997).

[25] M. Abeles, Isr J. Med. Sci. **18**, 83 (1982).

[26] W. R. Softky and C. Koch, J. Neurosci. **13**, 334 (1993).

[27] M. E. Larkum, J. J. Zhu, and B. Sakmann, Nature (London) **398**, 338 (1999).

[28] R. Azouz and C. M. Gray, Proc. Natl. Acad. Sci. U.S.A. **97**, 8110 (2000).