Published in final edited form as:

Curr Opin Neurobiol. 2005 October ; 15(5): 585–592. doi:10.1016/j.conb.2005.08.002.

Spike train metrics

Jonathan D Victor

Department of Neurology and Neuroscience, Weill Medical College of Cornell University, 1300 York Avenue, New York, NY 10021 USA

Abstract

Quantifying similarity and dissimilarity of spike trains is an important requisite for understanding neural codes. Spike metrics constitute a class of approaches to this problem. In contrast to most signal-processing methods, spike metrics operate on time series of all-or-none events, and are, thus, particularly appropriate for extracellularly recorded neural signals. The spike metric approach can be extended to multineuronal recordings, mitigating the ‘curse of dimensionality’ typically associated with analyses of multivariate data. Spike metrics have been usefully applied to the analysis of neural coding in a variety of systems, including vision, audition, olfaction, taste and electric sense.

Theoretical background

Introduction

Variability is a prominent feature of neural activity and its sources and functional implications are the focus of much investigation. Variability places limits on the reliability of signals, but can also provide a rich language for neuronal populations and their interactions. To analyze variability, one must first quantify the extent to which two patterns of neural activity are dissimilar, that is, one needs a ‘metric’ for comparing patterns. Metrics that are specifically applicable to spike trains (see glossary) are the focus of the present review.

Although laboratory measurements never correspond to a mathematical ideal [1], the choice of an appropriate mathematical framework is a prerequisite for rigorous data analysis. This choice is particularly crucial to the study of neural coding, because neural coding is fundamentally an abstraction: the relationship between stimuli, actions, and/or behavioral states, and the activity of one or more neurons.

A voltage record, the starting point for most neurophysiologic analyses of neural activity, is typically acquired as a set of closely spaced digital measurements. It seems natural to conceptualize a voltage record as a continuous function of time. However, we might also choose to conceptualize a neurophysiologic voltage record as a sequence of stereotyped events. This viewpoint of a voltage record as a realization of a ‘point process’ (see glossary) is especially appropriate for extracellularly recorded neural signals, given the all-or-none nature of the action potential and its effects on postsynaptic neurons [2–4]. The point process viewpoint (discrete events in continuous time) is intermediate between the vector space viewpoint (continuous time, continuous events) and a third alternative, symbol sequences (discrete time, discrete events). The symbol sequence viewpoint is often used in application of information theory and nonlinear dynamics to neuronal data [5–8]. The choice of viewpoint can influence the results of data analysis, both quantitatively and qualitatively. These considerations are fundamental to the rationale for the spike metric approach.

As detailed in Box 1, the continuous-record viewpoint leads naturally to metrics that have a particular mathematical structure (Euclidean distances [see glossary] within a vector space). Given the prominence of nonlinearities at many stages of neural processing and the kinds of ‘spaces’ (see glossary) that neural activity ultimately must represent, this restriction might not be desirable. Metrics based on the point process viewpoint do not have these constraints.

This review describes several kinds of metrics that can be applied to point processes in general, but that are motivated by neurophysiologic considerations specific to series of action potentials (‘spike trains’). I consider metrics applicable to single-unit and multineuronal activity and recent examples of their use. Computational considerations are summarized in Box 2.

A simple example

To describe the nature of metrics and metric spaces, we use a simple example, consisting of a set of cities $A, B, \dots$ linked by commercial airlines. We begin with a table of airfares that provides a list of costs for travel between various pairs of cities. We assume that fares between any two cities are independent of the direction of travel, that is, the costs are ‘symmetric’. The table of pairwise costs does not constitute a metric, because the cost $c(A,B)$ of a flight from A to B might be greater than the sum of the cost of a flight from A to X , and from X to B . However, we can turn the table of costs into a metric by defining the metric distance between two cities as the least total cost of transportation between them, making use of however many stopovers are necessary. More formally, we define the metric distance $d(A,B)$ between two points A and B as

$$d(A,B) = \min \left\{ \sum_{j=0}^{n-1} c(X_j, X_{j+1}) \right\} \quad (1)$$

where $\{X_0, X_1, \dots, X_n\}$ is an itinerary from A to B , with $X_0 = A$ and $X_n = B$. This definition guarantees that the metric distances between three points A, B , and C satisfy the triangle inequality $d(A,C) \leq d(A,B) + d(B,C)$, because the cheapest itinerary from A to C cannot be more expensive than an itinerary that is constrained to stop at B . It is easy to construct examples of metrics that violate the rules of Euclidean geometry.

Cost-based metrics for spike trains

With some fine print, Equation 1 provides a way to turn any set of (symmetric) costs into a metric. In the present context, each spike train $A, B, \dots$, is a point in the metric space (i.e. a ‘city’). The sequence of steps corresponding to each term on the right side of Equation 1 constitutes a path of elementary steps that transforms a spike train A into a spike train B . In spike train metrics, the costs $c(X_j, X_{j+1})$ of the elementary steps are determined by a small number of rules [9,10] that are intended to capture basic biological features of how spike trains influence postsynaptic neurons.

Motivation for the point process view

A spike train can be considered as a continuous function of time (e.g. the extracellularly recorded voltage record) or as abstracted into a ‘point process’, a sequence of stereotyped events. Here, we briefly discuss some implications of these viewpoints.

Vector spaces and point processes

Continuous functions of time form a vector space, in which specific algebraic operations are defined: addition of vectors; multiplication of vectors by scalars; and scalar product

(also known as an ‘inner product’ or a ‘dot product’), which produces a scalar from two vectors. These operations form the underpinning of signal-processing methods such as filtering, averaging, spectral estimates and signal detection. For point processes, these operations are not directly applicable. They are replaced by other procedures, such as the analysis of inter-event intervals, which are not directly applicable to continuous signals. Procedures specifically applicable to spike trains (but not continuous signals) are recognized as crucial to the descriptive [48–50], statistical [51] and information-theoretic [52,53] analysis of neural data.

Distances and metrics

In vector spaces of continuous signals, the scalar product defines a distance. The distance has the familiar properties of Euclidean geometry. For example, a line is the set of points that lie along the shortest path between two points, and a perpendicular is the unique shortest path from a point to a line that does not run through the point. Distances, known as ‘metrics’, can also be defined in a space of event sequences. However, these distances do not typically have Euclidean properties [54•]. There is no guarantee that a continuous ‘line’ can be constructed as the shortest distance between two event sequences, and ‘perpendiculars’ might not be unique. The present use of the term ‘metric’ is consistent with the topological definition of this term [55], and emphasizes that it endows a set of event sequences (the spike trains) with the properties of a topological ‘metric space’. Importantly, perceptual spaces are not typically Euclidean. For example, although human color vision is a three-parameter space determined by the three cone absorption spectra, perpendiculars (as defined by perceptual distances) need not be unique [56]. In olfaction, the situation is far more complex. Olfactory perceptual space might not even have a well-defined dimension, and mixing of odorants need not lead to intermediate percepts. Thus, the Euclidean geometry implied by vector spaces might be too confining to support a correspondence between neural activity and sensory perception [57]. These strictures of vector spaces are not present in the more general metric-space formalism.

Spike time metrics—The ‘spike time’ family of metrics is motivated by a caricature of a neuron as a coincidence detector [11–14]. For these metrics, deleting or inserting a spike has a cost of 1. This rule sets an overall scale for the metrics and ensures that any spike train can be transformed to any other spike train by a path: the path that successively deletes all spikes from train *A*, and then successively inserts all spikes into train *B*. The second rule, which confers sensitivity to spike timing, is that the cost of moving a single spike in time is proportional to the amount of time that it is moved. That is, if two spike trains *A* and *B* are identical except for a single spike that occurs at t_a in *A*, and t_b in *B*, then $c(A,B) = q|t_a - t_b|$, where q (in units of 1/sec) is a parameter that determines the relative sensitivity of the metric to spike count and spike timing.

Computational considerations

Computation of the metrics that correspond to Euclidean distances [25,26] is straightforward: spike trains are convolved with a smoothing kernel, and the scalar product of the resulting time series is calculated. The computational burden is proportional to the duration of the spike trains.

The cost-based metrics $D^{spike}[q]$ and $D^{interval}[q]$ are similar to ‘alignments’, ‘edit-length distances’ and ‘string-matching procedures’ used for comparison of genetic sequences [58] and electroencephalogram (EEG) analysis [59], in which there is a cost associated with insertions, deletions and shifts. This kinship provides an efficient dynamic programming algorithm to calculate these metrics [9,10]. In particular, efficient calculation of the minimal cost sequence is enabled by the following observations concerning the elementary steps in

a minimal-cost sequence. First, if the last spikes in the two trains are neither inserted nor deleted, then they must be connected by a shift. Second, the paths taken by two spikes cannot cross. Third, no spike needs to be shifted more than once. The computational burden of the algorithm is proportional to the product of the number of spikes in the trains to be compared.

For the multineuronal metric $D^{spike}[q,k]$, a clever algorithm in which the two spike trains are treated asymmetrically has been devised [60], which has a computational burden proportional to N^{L+1} , where N is the number of spikes in a typical train, and L is the number of labels (neurons).

Other cost-based metrics can be envisioned [9] (e.g. enabling both single-spike shifts as in D^{spike} and interval changes as in $D^{interval}$). However, no efficient computational algorithm for such metrics is available at present.

These rules, along with equation 1, suffice to provide a metric distance between all spike trains, denoted $D^{spike}[q]$. Intuitively, the distance between two spike trains is the minimum total cost of a set of elementary steps that transforms one spike train into the other (Figure 1). When q is very small, then the times of individual spikes have little influence on the calculated distance between spike trains. For $q = 0$, then spike timing is irrelevant, in that spikes can be shifted in time ‘for free’. Thus, for $q = 0$, a minimal path between spike trains A and B consists of deleting or inserting enough spikes into A so that the total count matches that of B , and then shifting the spikes in time so that they match. That is, $D^{spike}[0](A,B) = |n(A) - n(B)|$, where $n(X)$ denotes the number of spikes in the train X . $D^{spike}[0]$ can, thus, be thought of as formalizing a ‘spike count’ code: spike trains are considered different only if they contain a different number of spikes.

As q increases, the metric $D^{spike}[q]$ becomes increasingly sensitive to spike timing. To see this, consider spike trains A and B that each contain only one spike, at t_a in A , and t_b in B . There are two paths to consider in Equation 1: first, deleting the spike in A and reinserting it into B ; second, shifting the spike from t_a to t_b . The first path has a cost of 2; the second path has a cost of $q|t_a - t_b|$. The distance $D^{spike}[q](A,B)$ is the minimum of these two numbers. The break-even point is for $|t_a - t_b| = 2/q$. That is, in the metric $D^{spike}[q]$, two spikes are only considered as comparable if they occur within an interval of $2/q$ sec. Depending on whether the spike times in two trains A and B are similar, $D^{spike}[q](A,B)$ can range from $|n(A) - n(B)|$ (spike times in A and B match, no shifting of spike times needed) to $n(A) + n(B)$ (no spike times in A and B are within $2/q$ sec; spikes must be deleted from A and then reinserted into B). Thus, q explicitly represents the relative importance of spike times and spike counts: a change in the time of a spike by $1/q$ sec influences the total cost of a path as much as deleting the spike altogether. For neurons that act as a coincidence detector with integration time (or temporal resolution) $1/q$, spike trains will have similar postsynaptic effects if they are similar in the sense quantified by $D^{spike}[q]$. Often, temporal resolution is not known in advance — so, q is retained as a parameter, with the goal of using the dependence on q to analyze coding (see below and Box 3)

Spike interval metrics—For the ‘spike interval’ metrics [9,10], denoted $D^{interval}[q]$, the heuristic is that the postsynaptic effects of a spike might depend strongly on the recent activity at that synapse [4,15,16]. Correspondingly, the temporal dependence of $D^{interval}[q]$ is based on the intervals between spikes, rather than their absolute times. For $D^{interval}[q]$, the cost of insertion or deletion of an interspike interval is 1. The second rule is that shortening or extending an interspike interval by an amount t has a cost qt . Note that changing the length of an interval differs from changing the time of a spike, in that when an interval length is changed, the time of all later spikes are also changed. For this reason, $D^{interval}[q]$ and $D^{spike}[q]$ have fundamentally different topological characteristics [9].

Multineuronal cost-based metrics—Multineuronal recordings that enable the study of patterns of activity across neurons are increasingly available [17–19] and methods to analyze such data effectively are receiving increasing attention [20,21]. A multineuronal recording might be considered to be a sequence of labeled events. In this view, cost-based metrics are readily extended to the multineuronal context. The simplest approach [22] is to add a rule that sets the cost of changing the label associated with an event. This results in a two-parameter family of metrics, $D^{spike}[q,k]$, where k is a parameter that sets the cost of changing a label. $k = 0$ corresponds to the notion that the neuron of origin of a spike is irrelevant (because there is no cost associated with changing this label). $k = 2$ corresponds to a ‘labeled line’ interpretation, because changing the neuron of origin of a spike has the same cost as deleting a spike associated with one neuron, and inserting a spike associated with another neuron.

Spike metrics and information theory

Information theory (IT) [61] (see Cover and Thomas [62] for a general review) forms a natural framework for the analysis of neural coding [7]. However, IT was originally developed for the understanding of man-made communication systems, and its application to neural systems — especially the empiric estimation of mutual information from laboratory data — is not straightforward. One barrier is that the experimentalist must estimate stimulus–response probabilities from limited data, leading to estimates that are biased, imprecise, or both. The estimation problem is only exacerbated in experiments that explore a stimulus space comprehensively or involve recording from several neurons simultaneously — because both increase the number of stimulus–response probabilities that need to be estimated. For these reasons, naive application of IT to datasets of the size typically obtained in the laboratory can yield misleading estimates of information [63,64].

Use of spike metrics can mitigate this problem, in that estimates of IT quantities based on spike metrics [9] make implicit use of the hypothesis that similar responses have similar stimulus-dependent probabilities. Such estimates are typically downwardly biased, because the notion of similarity postulated by a spike metric is likely to deviate from the true ‘neural code’. However, the combined use of metric-spaced measures or related techniques, along with nonparametric approaches with complementary bias characteristics, supports the validity of the estimates [65]. Additionally, the dependence of an estimate of information on the parameters of the spike metric describes which features of spike trains are crucial for carrying information [9,10,22,32–36,37•].

Other metrics—Within the cost-based framework, one can also construct metrics sensitive to motifs of spikes [23] (by having a rule associated with the cost of moving a set of spikes), and metrics that combine the rules of $D^{spike}[q]$ and $D^{interval}[q]$. Metrics on spike trains can also be obtained by binning [24] or convolving them with a smoothing kernel (see glossary) [25, 26], and then using standard vector-space distances between the derived temporal functions. These latter approaches necessarily lead to Euclidean distances.

Applications

Spike metrics have been applied to data obtained from a variety of neural systems, to quantify variability *per se*, [27,28,29•,30,31], to characterize neural coding (at the level of temporal coding within single neurons [9,10,32–36] or neuronal pairs [22,37•]), and to evaluate models [38,39•]

Electrosensory

Kreiman and co-workers [27] studied the variability of the P-receptor afferents [40] in the weakly electric fish *Eigen-mannia*, whose discharges are loosely phase-locked to the periodic (200–600 Hz) discharge of its electrosensory organ. They made extensive use of measures of variability based on D^{spike} , because these, but not measures based on spike count or its variance, appeared to capture trial- to trial- to variability in P-receptor activity. They found that the intrinsic variability of spike trains was not likely to degrade information transmission, but enabled improvement in coding by averaging across multiple afferent fibers.

Vision

Blowfly H1—Grewe *et al.* [28] used D^{spike} , along with a Euclidean metric, to examine variability in the responses of the blowfly wide-field, motion-sensitive neuron H1, driven by motion stimuli with various levels of added noise. By determining the ‘equivalent noise’ (the maximum amount of added noise that enabled a criterion level of response classification), they deduced that internal noise, rather than photon noise, limited the performance of the H1 neuron.

Mammalian retina—Chichilnisky and Rieke [29•] used D^{spike} to analyze near-threshold signaling in rod photoreceptors and retinal ganglion cells of the tiger salamander. Although grouping of responses into signal versus no-signal clusters was no more accurate than could be achieved by standard methods such as the Fisher discriminant (see glossary) [41], the observation that optimal grouping was achieved for $D^{spike}[q]$ at $q = 0.1$ indicated a meaningful temporal resolution of 100 ms.

Lateral geniculate nucleus—In response to full-field random flicker, retinal and lateral geniculate neurons often fire in discrete ‘firing events’ consisting of several spikes, at times that are reproducible across trials [30,38]. Reinagel and Reid [30] used D^{spike} to show that other than a possible small overall latency shift, these timing events were conserved not only across trials, but also across animals. Keat *et al.* [38] later used D^{spike} to evaluate the ability of models to predict these firing events.

Visual cortex—Application of spike metrics to single-unit and multi-unit (single units not separated) recordings in primary visual cortex (V1) and early extrastriate cortex (V2 and V3) of the awake macaque [10] revealed that the temporal structure of spike trains contributed significantly to coding of visual information. Because information estimates based on D^{spike} were generally larger than those based on $D^{interval}$, it was concluded that spike timing (relative to stimulus onset), rather than the interval structure, was the more crucial aspect of temporal structure. Moreover, because information estimates for the several stimulus attributes (contrast, spatial frequency, orientation, size, and texture) had distinctive dependences on q , it was concluded that typical neurons multiplexed visual information, representing contrast with high temporal precision (ca. 10–30 ms), and spatial aspects with lower temporal precision (ca. 100 ms).

In V1 of the anesthetized macaque, Reich *et al.* [32] used D^{spike} to show that most of the information about contrast could be extracted from the latency of the first spike, although additional information could be extracted from the temporal structure of the response without regard to latency, and that temporal coding was particularly important at higher contrasts, at which the spike rate response neared saturation. Mechler *et al.* [33] used a variant of D^{spike} appropriate for responses to periodic stimuli to demonstrate that temporal structure played a much larger role in the coding of edges (square-wave gratings) than of smooth variations (sine gratings). Cyclic variants of D^{spike} were later used [34] to characterize the coding of edge-like, line-like, and intermediate one-dimensional features in V1, and showed that many neurons demonstrated feature opponency and/or feature selectivity for compound gratings.

Two studies focused on the role of the activity pattern within a local cluster of cortical neurons. Aronov *et al.* [22] used $D^{spike}[q,k]$ to characterize the coding of spatial phase across pairs of V1 neurons, isolated using tetrodes (four-element microelectrodes) [17]. Dependence of information estimates on q indicated that spike times had an informative precision of ca. 30 ms; dependence of information estimates on k indicated that the neuron of origin, and not just the total activity of the local population, contributed to coding of spatial phase. The geometry of the stimulus set (the circle of spatial phase) corresponded to the response similarities as determined by $D^{spike}[q,k]$. Greater fidelity of the representation was achieved for nonzero values of k , thus demonstrating that within a local cluster, the neuron of origin of a spike, in addition to its timing, carries information. This study also introduced a technique (designated ‘temporal profiles’) to identify the time course of temporal features that are crucial to stimulus representation. Samonds, Bonds and co-workers [37•,42] examined signaling of orientation in cat primary visual cortex with D^{spike} and $D^{interval}$. Although large angular differences were adequately represented by firing rate, type analysis [43] of responses of neuronal pairs suggested that cooperative signaling was present for small angular differences [42]. Metric-space analysis [37•] demonstrated that orientation differences of 10 deg or less were signaled by the temporal fine structure (2–10 ms) of spike times and spike intervals.

Audition

Middlebrooks and co-workers [44] showed that responses of single neurons in cat ectosylvian gyrus, when analyzed with a neural-network approach, represent the azimuth of a sound stimulus in a panoramic (360 deg) fashion. A reanalysis [9] of these data using D^{spike} and $D^{interval}$ showed that this representation relied on spike timing, at a resolution of ca. 4 ms. A comparable conclusion was also reached in a neural-network analysis of A2 neurons [45] using surrogate-data methods.

Machens *et al.* [36] showed that 100 ms samples of natural songs from up to eight different conspecific grasshoppers could be distinguished by spike trains of single auditory neurons via $D^{spike}[q]$ at $q = 100$, corresponding to an informative precision of 10 ms.

Chemical senses

Laurent and co-workers [31] identified a population of neurons in the olfactory system of the locust, the β -lobe neurons, that are crucial to reading out the output of the mushroom body, a structure involved in odor learning. Desynchronization of projection neurons (PNs), two synapses upstream of these β -lobe cells, could be achieved by blocking fast GABA inhibition with picro-toxin. This leads to loss of behavioral discrimination of similar odors, although coarse odor discrimination remains intact [46]. D^{spike} was used [31] to show that information about odors contained in the spike trains of β -lobe cells is lost when the PNs are desynchronized, even though no such loss is observed within spike trains of individual PNs. These elegant experiments demonstrate the functional relevance of neuronal synchronization.

Di Lorenzo and Victor [35] applied D^{spike} and $D^{interval}$ to the analysis of gustatory coding in single-neuron responses in the nucleus of the solitary tract of the rat. In 10 of 19 neurons, the temporal structure of the initial 2 s of the response contributed to the coding of the four basic taste qualities. The informative precision of spike timing was much lower than in the other studies reviewed here, typically about 300 ms, and response dynamics contributed the most to coding in neurons that were the most variable, in terms of their overall firing rate. As in the visual cortex studies described above [10] analysis of surrogate datasets obtained by shuffling spikes across trials served to demonstrate that the informative aspects of spike timing went beyond that of a Poisson sampling of a firing rate envelope.

Motor

In preliminary work [47] D^{spike} was used to identify aspects of single-neuron activity in parietal cortex of the macaque that were correlated with arm approach and grasp style. Of note, this re-analysis made use of a public database and analysis toolkit.

Conclusions

Spike metrics are applicable to data that can be viewed as discrete events in continuous time. Cost-based spike metrics are a general strategy for formalizing biologically motivated notions of distance, and, thus, constitute a principled approach for the analysis of variability of single- and multineuronal extracellular recordings. They have been fruitfully applied in a variety of neural systems, to characterize and consequently help to understand neuronal variability and coding. Analyses in several sensory systems have shown that spike count might suffice for signaling gross sensory differences, but spike timing is important for signaling subtle differences [31,32,37••].

By parameterizing biologically motivated notions of distances between time series of events, metric-space methods complement non-parametric approaches to estimate information-theoretic quantities from limited data. Spike metrics are related to edit-length distances (see glossary) used for genetic analysis, and present related computational challenges.

Glossary

Note to readers, definitions are intended to explain the concepts, not to provide mathematical rigor ; Edit-length distance, a distance defined on a space of sequences, based on minimizing the cost of transforming one sequence into another by a sequence of ‘edits’ (e.g. insertions, deletions, and substitutions). Edit-length distances are typically not Euclidean.; Euclidean distance, a distance that has the familiar properties of Euclidean geometry, including the existence of unique parallel lines and perpendiculars; Fisher discriminant, a classical approach [41] to divide multivariate data into two categories, based on a weighted sum of the data values. Point process: a process that produces a time series of identical events. The mathematical specification of a point process determines the probability of every such sequence. A ‘realization’ of a point process is an individual time series, chosen according to this probability rule ; Smoothing kernel, a mathematical function describing how nearby values of a time series are to be weighted when these values are averaged together to smooth a signal. Typical smoothing kernels have rectangular (‘boxcar’), Gaussian, or triangular shapes; Space, a generic mathematical term for a collection of objects, usually with one or more properties and/or operations that apply to these objects. In a ‘metric space’, there is a well-defined distance between any two objects. In a ‘vector space’, the operations of vector addition and multiplication by scalars (e.g. the real numbers) are defined on the objects (vectors). Typical vector spaces also have a scalar product (also known as an ‘inner product’ or a ‘dot product’) that produces a scalar from two vectors. The existence of a scalar product implies the existence of a Euclidean distance; Spike train, a sequence of action potentials originating from a single neuron..

Acknowledgements

This work is supported in part by National Eye Institute Grant 1RO1 EY9314 to J Victor and a National Institute of Mental Health Grant 1RO1 MH68012 to D Gardner. The author thanks K Purpura for helpful comments on the manuscript.

References and recommended reading

Papers of particular interest, published within the annual period of review, have been highlighted as:

- of special interest
- of outstanding interest

1. Slepian D. On bandwidth. *Proc IEEE* 1976;64:292–300.
2. Segundo, JP.; Perkel, DH. The nerve cell as an analyzer of spike trains. In: Brazier, MAB., editor. *The Interneuron*. University of California Press; 1969. p. 349–390.
3. Abbott LF. Integrating with action potentials. *Neuron* 2000;26:3–4. [PubMed: 10798385]
4. Sen K, Jorge-Rivera JC, Marder E, Abbott LF. Decoding synapses. *J Neurosci* 1996;16:6307–6318. [PubMed: 8815910]
5. Kennel MB, Mees AI. Context-tree modeling of observed symbolic dynamics. *Phys Rev E Stat Nonlin Soft Matter Phys* 2002;66:056209
6. Strong SP, Koberle R, Ruyter van Steveninck RR, Bialek W. Entropy and information in neural spike trains. *Phys Rev Lett* 1998;80:197–200.
7. Rieke, F.; Warland, D.; Ruyter van Steveninck, RR.; Bialek, W. *Spikes: Exploring the Neural Code*. MIT Press; 1997.
8. Kennel MB, Shlens J, Abarbanel HD, Chichilnisky EJ. Estimating entropy rates with Bayesian confidence intervals. *Neural Comput* 2005;17:1531–1576. [PubMed: 15901407]
9. Victor JD, Purpura KP. Metric-space analysis of spike trains: theory, algorithms and application. *Network* 1997;8:127–164.
10. Victor JD, Purpura KP. Nature and precision of temporal coding in visual cortex: a metric- space analysis. *J Neurophysiol* 1996;76:1310–1326. [PubMed: 8871238]
11. Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 1993;13:334–350. [PubMed: 8423479]
12. Egger V, Feldmeyer D, Sakmann B. Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nat Neurosci* 1999;2:1098–1105. [PubMed: 10570487]
13. Kuba H, Yamada R, Fukui I, Ohmori H. Tonotopic specialization of auditory coincidence detection in nucleus laminaris of the chick. *J Neurosci* 2005;25:1924–1934. [PubMed: 15728832]
14. Abeles M. Role of the cortical neuron: integrator or coincidence detector? *Isr J Med Sci* 1982;18:83–92. [PubMed: 6279540]
15. Dan Y, Poo MM. Spike timing-dependent plasticity of neural circuits. *Neuron* 2004;44:23–30. [PubMed: 15450157]
16. Markram H, Lubke J, Frotscher M, Sakmann B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 1997;275:213–215. [PubMed: 8985014]
17. Gray CM, Maldonado PE, Wilson M, McNaughton B. Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J Neurosci Methods* 1995;63:43–54. [PubMed: 8788047]
18. Kralik JD, Dimitrov DF, Krupa DJ, Katz DB, Cohen D, Nicolelis MA. Techniques for long-term multisite neuronal ensemble recordings in behaving animals. *Methods* 2001;25:121–150. [PubMed: 11812202]
19. Nicolelis MA, Dimitrov D, Carmena JM, Crist R, Lehew G, Kralik JD, Wise SP. Chronic, multisite, multielectrode recordings in macaque monkeys. *Proc Natl Acad Sci USA* 2003;100:11041–11046. [PubMed: 12960378]
20. Panzeri S, Schultz SR. A unified approach to the study of temporal, correlational, and rate coding. *Neural Comput* 2001;13:1311–1349. [PubMed: 11387048]
21. Barbieri R, Frank LM, Nguyen DP, Quirk MC, Solo V, Wilson MA, Brown EN. Dynamic analyses of information encoding in neural ensembles. *Neural Comput* 2004;16:277–307. [PubMed: 15006097]

22. Aronov D, Reich DS, Mechler F, Victor JD. Neural coding of spatial phase in V1 of the macaque monkey. *J Neurophysiol* 2003;89:3304–3327. [PubMed: 12612048]
23. Abeles M, Prut Y. Spatio-temporal firing patterns in the frontal cortex of behaving monkeys. *J Physiol (Paris)* 1996;90:249–250. [PubMed: 9116677]
24. Lim D, Capranica RR. Measurement of temporal regularity of spike train responses in auditory nerve fibers of the green treefrog. *J Neurosci Methods* 1994;52:203–213. [PubMed: 7967723]
25. van Rossum MC. A novel spike distance. *Neural Comput* 2001;13:751–763. [PubMed: 11255567]
26. Richmond BJ, Optican LM. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. II. Quantification of response waveform. *J Neurophysiol* 1987;57:147–161. [PubMed: 3559669]
27. Kreiman G, Krahe R, Metzner W, Koch C, Gabbiani F. Robustness and variability of neuronal coding by amplitude-sensitive afferents in the weakly electric fish *eigenmannia*. *J Neurophysiol* 2000;84:189–204. [PubMed: 10899196]
28. Grewe J, Kretzberg J, Warzecha AK, Egelhaaf M. Impact of photon noise on the reliability of a motion-sensitive neuron in the fly's visual system. *J Neurosci* 2003;23:10776–10783. [PubMed: 14645469]
29. Chichilnisky EJ, Rieke F. Detection sensitivity and temporal resolution of visual signals near absolute threshold in the salamander retina. *J Neurosci* 2005;25:318–330.330 [PubMed: 15647475] The authors used D^{spike} to decode ganglion cell responses, and to deduce that spike times, at a resolution of 100 ms, contributed to the detection of near-threshold signals
30. Reinagel P, Reid RC. Precise firing events are conserved across neurons. *J Neurosci* 2002;22:6837–6841. [PubMed: 12177180]
31. MacLeod K, Backer A, Laurent G. Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature* 1998;395:693–698. [PubMed: 9790189]
32. Reich D, Mechler F, Victor J. Temporal coding of contrast in primary visual cortex: when, what, and why? *J Neurophysiol* 2001;85:1039–1050. [PubMed: 11247974]
33. Mechler F, Victor JD, Purpura KP, Shapley R. Robust temporal coding of contrast by V1 neurons for transient but not for steady-state stimuli. *J Neurosci* 1998;18:6583–6598. [PubMed: 9698345]
34. Mechler F, Reich DS, Victor JD. Detection and discrimination of relative spatial phase by V1 neurons. *J Neurosci* 2002;22:6129–6157. [PubMed: 12122074]
35. Di Lorenzo PM, Victor JD. Taste response variability and temporal coding in the nucleus of the solitary tract of the rat. *J Neurophysiol* 2003;90:1418–1431. [PubMed: 12966173]
36. Machens C, Prinz P, Stemmler M, Ronacher B, Herz A. Discrimination of behaviorally relevant signals by auditory receptor neurons. *Neurocomputing* 2001;38:263–268.
37. Samonds JM, Bonds AB. From another angle: differences in cortical coding between fine and coarse discrimination of orientation. *J Neurophysiol* 2004;91:1193–1202.1202 [PubMed: 14614106] The authors used metric-space analysis to show that neurons in primary visual cortex use firing rate to signal large orientation differences, whereas the fine temporal structure of their responses signal orientation differences of 10 deg or less
38. Keat J, Reinagel P, Reid RC, Meister M. Predicting every spike: a model for the responses of visual neurons. *Neuron* 2001;30:803–817. [PubMed: 11430813]
39. Tiesinga PHE. Chaos-induced modulation of reliability boosts output firing rate in downstream cortical areas. *Phys Rev E Stat Nonlin Soft Matter Phys* 2004;69 031912. The author used D^{spike} to analyze the activity of model neural networks, demonstrating that gamma-frequency input modulations can induce large increases in the reliability of spike timing with only modest changes in firing rate
40. Nelson ME. Multiscale spike train variability in primary electrosensory afferents. *J Physiol (Paris)* 2002;96:507–516. [PubMed: 14692498]
41. Fisher RA. The use of multiple measurements in taxonomic problems. *Ann Eugen* 1936;7:179–188.
42. Samonds JM, Allison JD, Brown HA, Bonds AB. Cooperation between area 17 neuron pairs enhances fine discrimination of orientation. *J Neurosci* 2003;23:2416–2425. [PubMed: 12657701]
43. Johnson DH, Gruner CM, Baggerly K, Seshagiri C. Information-theoretic analysis of neural coding. *J Comput Neurosci* 2001;10:47–69. [PubMed: 11316339]

44. Middlebrooks JC, Clock AE, Xu L, Green DM. A panoramic code for sound location by cortical neurons. *Science* 1994;264:842–844. [PubMed: 8171339]
45. Furukawa S, Middlebrooks JC. Cortical representation of auditory space: information-bearing features of spike patterns. *J Neurophysiol* 2002;87:1749–1762. [PubMed: 11929896]
46. Stopfer M, Bhagavan S, Smith BH, Laurent G. Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* 1997;390:70–74. [PubMed: 9363891]
47. Vaknin R, Goldberg DH, Victor JD, Gardner EP, Debowy DJ, Babu KS, Gardner D. Metric space analysis of responses in parietal cortex during prehension [abstract]. *Soc Neurosci Abstracts* 2005;31984.20
48. Perkel DH, Gerstein GL, Moore GP. Neuronal spike trains and stochastic point processes. I. The single spike train. *Biophys J* 1967;7:391–418. [PubMed: 4292791]
49. Perkel DH, Gerstein GL, Moore GP. Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys J* 1967;7:419–440. [PubMed: 4292792]
50. Aertsen AM, Gerstein GL, Habib MK, Palm G. Dynamics of neuronal firing correlation: modulation of “effective connectivity”. *J Neurophysiol* 1989;61:900–917. [PubMed: 2723733]
51. Brown EN, Barbieri R, Ventura V, Kass RE, Frank LM. The time-rescaling theorem and its application to neural spike train data analysis. *Neural Comput* 2002;14:325–346. [PubMed: 11802915]
52. McFadden JA. The entropy of a point process. *Journal of the Society for Industrial and Applied Mathematics* 1965;13:988–994.
53. Stein RB. The information capacity of nerve cells using a frequency code. *Biophys J* 1967;7:797–826. [PubMed: 19210999]
54. Aronov D, Victor JD. Non-Euclidean properties of spike train metric spaces. *Phys Rev E Stat Nonlin Soft Matter Phys* 2004;69 061905. This study proved that the geometry associated with spike metrics is intrinsically non-Euclidean
55. Gaal SA. *Point Set Topology*. Academic Press. 1964
56. Wuerger SM, Maloney LT, Krauskopf J. Proximity judgments in color space: tests of a Euclidean color geometry. *Vision Res* 1995;35:827–835. [PubMed: 7740773]
57. Hopfield JJ. Pattern recognition computation using action potential timing for stimulus representation. *Nature* 1995;376:33–36. [PubMed: 7596429]
58. Sellers P. On the theory and computation of evolutionary distances. *SIAM J Appl Math* 1974;26:787–793.
59. Wu L, Gotman J. Segmentation and classification of EEG during epileptic seizures. *Electroencephalogr Clin Neurophysiol* 1998;106:344–356. [PubMed: 9741763]
60. Aronov D. Fast algorithm for the metric-space analysis of simultaneous responses of multiple single neurons. *J Neurosci Methods* 2003;124:175–179. [PubMed: 12706847]
61. Shannon, CE.; Weaver, W. *The Mathematical Theory of Communication*. University of Illinois Press; 1949.
62. Cover, TM.; Thomas, JA. *Elements of Information Theory*. Schilling, DL., editor. Wiley; 1991.
63. Victor JD. Binless strategies for estimation of information from neural data. *Phys Rev E Stat Nonlin Soft Matter Phys* 2002;66051903
64. Eguia MC, Rabinovich MI, Abarbanel HD. Information transmission and recovery in neural communications channels. *Phys Rev E Stat Phys Plasmas Fluids Relat Interdiscip Topics* 2000;62:7111–7122. [PubMed: 11102068]
65. Nelken I, Chechik G, Mscic-Flogel TD, King AJ, Schnupp JWH. Encoding stimulus information by spike numbers and mean response time in primary auditory cortex. *J Comput Neurosci* 2005;19:199–221. [PubMed: 16133819]

Figure 1.

A diagram of a sequence of elementary steps that transforms spike train A into spike train B. Each rectangle represents one spike, and the line that they rest on denotes time. Each elementary step is one of three types: deletion of a spike (deleted spike shown in red), insertion of a spike (inserted spike shown in green), or shifting a spike in time (blue arrows).