



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A Novel Spike Distance

Citation for published version:

Van Rossum, MCW 2001, 'A Novel Spike Distance', *Neural Computation*, vol. 13, no. 4, pp. 751-763.
<https://doi.org/10.1162/089976601300014321>

Digital Object Identifier (DOI):

[10.1162/089976601300014321](https://doi.org/10.1162/089976601300014321)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Neural Computation

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



A Novel Spike Distance

M. C. W. van Rossum

Department of Biology, Brandeis University, Waltham, MA 02454, U.S.A.

The discrimination between two spike trains is a fundamental problem for both experimentalists and the nervous system itself. We introduce a measure for the distance between two spike trains. The distance has a time constant as a parameter. Depending on this parameter, the distance interpolates between a coincidence detector and a rate difference counter. The dependence of the distance on noise is studied with an integrate-and-fire model. For an intermediate range of the time constants, the distance depends linearly on the noise. This property can be used to determine the intrinsic noise of a neuron.

1 Introduction ---

In the analysis of experimental neural data (e.g., when studying the effect of a certain manipulation or how reproducible the neuron's response is), one often encounters the question, "How similar are these two spike trains?" A problem is that it is not really known what information to look for in the spike train. In particular cases the precise timing is known to be important, whereas in other cases, only the number of spikes in a certain interval seems of importance (Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1996). The nervous system itself is possibly confronted with the same question of similarity, as when it has to respond to a certain visual cue and has to decide what the appropriate response is ("Was that really a tiger?!").

One approach in addressing this question is to introduce a distance that measures the (dis)similarity of two spike trains. If the distance between two spike trains is small enough, one can assume that the inputs were identical. Alternatively, distance measures can be used in a forced-choice experiment in which the spike train is compared to different templates. In that case, the template with the smallest distance to the trial should be chosen. Measures to compare spike trains have been introduced—for example:

- The total spike count. The total number of spikes in the spike trains is compared. This method is quite effective, but it misses all temporal structure in the spike trains.
- To resolve temporal structures, the spikes can be binned and the number of spikes per bin counted. Next, one measures the number of coincident spikes (see, e.g., Kistler, Gerstner, & van Hemmen, 1997). Alter-

natively, the count is interpreted as a vector in N -dimensional space, where N is the number of bins. The distance to other spike trains can be calculated using an N -dimensional Euclidean distance (Geisler, Albrecht, Salvi, & Saunders, 1991; MacLeod, Backer, & Laurent, 1998). A disadvantage of the binning procedure is that it does not distinguish between a spike shifted so that it just drops out of the bin and a spike shifted many bin durations. And a spike that is shifted but stays within the bin is treated like a spike that was not shifted at all. These might not always be desired effects if the spike timing matters.

- A more flexible distance measure is needed that takes the temporal structure into account but avoids the problems associated with binning. In a set of comprehensive papers Victor and Purpura (1996, 1997) introduced a measure based on a cost function. Two processes differentiate one spike train from another. (1) spikes can be deleted or inserted, and (2) they can be shifted in time. Victor and Purpura attributed to both processes a certain (arbitrary) cost and calculated how much it would cost to transform one spike train into the other one. The cost of insertion or deletion was fixed to one. If no cost was associated with the shifting process, the total cost was just the difference in the total number of spikes. Otherwise the cost function for the shift was taken to be a monotonic increasing function of the spike time difference. Although this method has been applied successfully (MacLeod et al., 1998), the calculation of the full cost function is quite involved. The reason is that it is not always clear where a displaced spike came from, and if the number of spikes in the trains is unequal, it can be difficult to determine which spike was inserted or deleted.

Here we introduce a spike distance closely related to the distance introduced by Victor and Purpura, yet it is easier to calculate and has a physiological interpretation. The distance is used to measure the intrinsic noise of a model neuron.

2 A Novel Measure

With the goal of a simple distance measure, we propose the following: Given a spike train with spike times t_i ,

$$f^{\text{orig}}(t) = \sum_i^M \delta(t - t_i), \quad (2.1)$$

where we will assume that all $t_i > 0$. Replace the delta function associated with each spike with an exponential function, that is, add an exponential tail to all spikes,

$$f(t) = \sum_i^M H(t - t_i) e^{-(t-t_i)/\tau_c}. \quad (2.2)$$

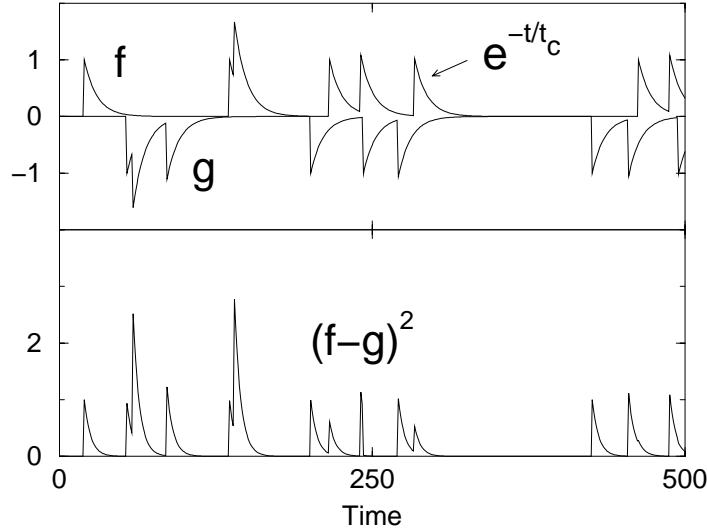


Figure 1: Definition of the distance. (Top) Two spike trains (one flipped) are convolved with an exponential with time constant t_c . (Bottom) The difference squared of the spike trains. The distance is given by the integral of this curve.

Here t_c is the time constant of the exponential function and H is the Heaviside step function ($H(x) = 0$ if $x < 0$ and $H(x) = 1$ if $x \geq 0$). In principle, one could convolve with functions other than the exponential. Our choice for the exponential function was motivated by its causality, simplicity, and possible biological interpretation (see section 4). The distance between two trains f and g we define as (see Figure 1)

$$D^2(f, g)_{t_c} = \frac{1}{t_c} \int_0^\infty [f(t) - g(t)]^2 dt. \quad (2.3)$$

The distance is the Euclidean distance of the two filtered spike trains with t_c as a free parameter.

To get a sense of the distance, consider the two limits of t_c . For t_c much smaller than the interspike interval, the smeared functions f and g contribute to the integral only if the spikes are not more than t_c apart. This is very similar to a coincidence detection. For most spike trains, coincident spikes can be neglected in the limit of zero t_c ; thus, if f contains M and g contains N spikes, one has

$$\lim_{t_c \rightarrow 0} D^2(f, g)_{t_c} = \frac{1}{t_c} \int_0^\infty [f^2(t) + g^2(t)] dt = \frac{M + N}{2}. \quad (2.4)$$

The distance counts the noncoincident spikes.

In the opposite limit, for large t_c , the main contribution to the integral comes from times when the last spike has passed but the exponent has still not decayed. Assuming that f (g) contains M (N) spikes, one can approximate

$$\lim_{t_c \rightarrow \infty} D^2(f, g)_{t_c} = \frac{1}{t_c} \int_0^\infty (Me^{-t/t_c} - Ne^{-t/t_c})^2 dt = \frac{(M - N)^2}{2}. \quad (2.5)$$

In this limit, D measures the difference in total spike count. Note that it is important that the upper limit of the integral is taken to be infinity (or, in practice until the tail of the last spike died out) instead of the time of the last spike.

The distance thus interpolates between the two extremes of coincidence detection and measuring difference in total spike count. The distance measure is very easy to implement numerically. For the best results, the integral between spikes and after the last spike can be done analytically. After filtering, the spike trains this leaves a sum of $N + M$ terms.

Changing integration variables leads to an alternative expression for the distance,

$$D^2(f, g) = \frac{1}{2} \int_{-\infty}^\infty C_{f-g, f-g}(t) e^{-|t|/t_c} dt, \quad (2.6)$$

where $C_{f-g, f-g}(t)$ is the autocorrelation of the difference of the raw spike trains, $f^{\text{orig}}(t) - g^{\text{orig}}(t)$. This shows that the distance can be interpreted as the weighted integral over the autocorrelation, with the weighting depending on t_c . It also shows that the distance is invariant under time reversal, that is, the distance is the same if the exponential tails were attached to the other sides of the spikes.

2.1 Analytical Results. For some simple cases, we can derive analytical results. Let us consider the distance between two almost identical spike trains. First consider the insertion of a single spike at time t_i into train f , other than that, the spike trains are identical, that is,

$$g(t) = f(t) + H(t - t_i)e^{-(t-t_i)/t_c}.$$

The distance is

$$\begin{aligned} D_{\text{insertion}}^2(f, g) &= \frac{1}{t_c} \int_{t_i}^\infty e^{-2(t-t_i)/t_c} \\ &= \frac{1}{2}. \end{aligned} \quad (2.7)$$

The removal of a spike yields the same answer. Note that the result is independent of t_c .

Now suppose a spike is shifted from t_i in spike train f to time $t_i + \delta t$ in g ,

$$g(t) = f(t) - H(t - t_i)e^{-(t-t_i)/t_c} + H(t - t_i - \delta t)e^{-(t-t_i-\delta t)/t_c},$$

which yields a distance,

$$\begin{aligned} D_{\text{displace}}^2(f, g) &= \frac{1}{t_c} \int_{t_i}^{t_i+\delta t} e^{-2(t-t_i)/t_c} + \frac{1}{t_c} \int_{t_i+\delta t}^{\infty} [e^{-(t-t_i)/t_c} - e^{-(t-t_i-\delta t)/t_c}]^2, \\ &= 1 - e^{-|\delta t|/t_c}. \end{aligned} \quad (2.8)$$

which is unity for small t_c and vanishes for large t_c . The ratio of distances due to spike insertion and displacement depends on t_c . For a shift $\delta t = t_c \ln(2)$, insertion and displacing cost the same. Note, however, that displacing a spike never costs more than insertion or removal of two spikes. The reason is that shifting a spike can always be done by removing the spike at time t and reinserting it at time $t + \delta t$.

2.2 Correlation Effects. When one considers two trains in which more spikes are changed, the distance is more complicated and certainly not always equal to the sum of the individual displacement and insertion distances. As an example, consider the case where two spikes, time T apart, are both displaced at time δt . No other spike occurs between the two. For this case one finds

$$D^2(f, g) = 2[1 - e^{-|\delta t|/t_c}] - 2e^{-|T|/t_c} [\cosh(\delta t/t_c) - 1]. \quad (2.9)$$

In Figure 2 the distance is plotted for various amounts of shift. The first term is twice the single spike displacement distance, equation 2.8. The second term is due to the correlation and is always negative. Thus, the total distance of displacing two spikes is less if the spikes are close. This seems a natural phenomenon, as the spike trains will look more similar when two neighboring spikes are shifted than when two spikes far apart are shifted. However, proper description of such effects calls for a more complex distance measure with the introduction of extra timescales (Victor & Purpura, 1997).

One can also study the insertion of two spikes, time T apart. For that case, one finds

$$D^2(f, g) = \frac{2}{2} + e^{-T/t_c}, \quad (2.10)$$

that is, the cost of insertion two spikes close together is larger than the cost of inserting two distant spikes. Equivalently, when more than one spike is inserted, the distance will increase with increasing t_c (see Figure 2).

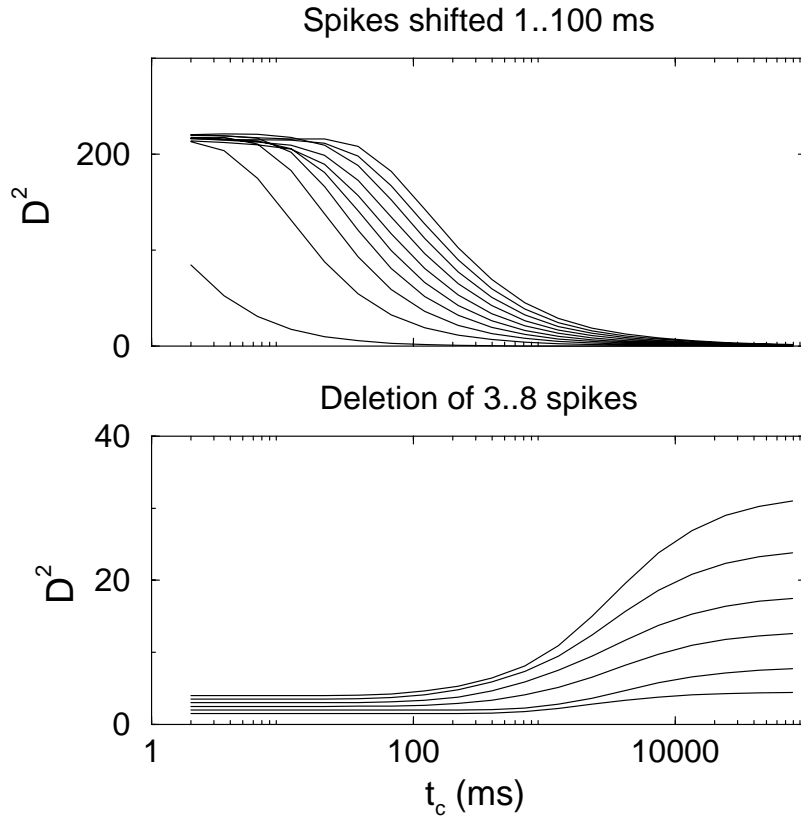


Figure 2: Effect of the deletion and shifting of spikes on the distance versus t_c . (Top) All the spikes in the second spike train are shifted 1, . . . , 100 ms with respect to the first train. At large t_c , this shift does not contribute to the distance. The spike trains were Poisson trains with a duration of 10 s and a rate of 20 Hz. (Bottom) At random positions, 3, . . . , 8 spikes are deleted from the second spike train (lower to upper curve). The distance is largest at large t_c .

2.3 Distance Between Uncorrelated Poisson Trains. Next we calculate the distance between two uncorrelated Poisson spike trains, both with a rate ρ . For small t_c , the distance measures the number of spikes (see equation 2.4). Assuming a total duration of T , on the average ρT spikes will be produced, so for small t_c , $D^2 = \rho T$. On the other hand, for large t_c , according to equation 2.5, the distance approaches $(M - N)^2/2$, where M and N again denote the number of spikes in the trains. The expectation value for $(M - N)^2$ for two Poisson processes is $2\rho T$; hence, for large t_c , one has $D^2 = \rho T$. The average distance is thus identical at small and large t_c . Alternatively, one




Figure 3: Distance between two independent Poisson trains. The distance is on average constant across t_c , but shows large fluctuations at large t_c .

can use equation 2.5 and $C_{f-g,f-g} = C_{f,f} + C_{g,g} - 2C_{f,g} = 2\rho T\delta(t)$, which yields $D^2 = \rho T$ independent of t_c .

Numerical simulations are shown in Figure 3. The distance shows large fluctuations for large t_c ; the difference in the number of spikes fluctuates more strongly than the number of nonoverlapping spikes. The reason is that the difference can be interpreted as a variance, while the number of nonoverlapping spikes behaves more like a mean value. It is well known that the variance is a more variable quantity than the mean.

3 Distance Between Two Noise-Driven Spike Trains

Next we study how the distance varies in response to changes in the input for a spiking neuron. We use a leaky integrate-and-fire neuron with a time constant of 50 ms. The stimulus is a gaussian distributed variable with zero mean and 10 ms temporal correlation. Its variance is adjusted to give an average spike frequency of 20 Hz. To compare spike trains, this stimulus is repeated across trials, and nonrepeating gaussian noise is added.

The standard deviation of the noise, σ_{add} , was up to 10% of the standard deviation of the stimulus. In Figure 4 we plot the distance as t_c is varied. As expected, the distance initially drops as spike time differences due to slightly displaced spikes are smoothed out. However, for large t_c , the distance increases again. This contrasts with the distance between two Poisson spike




Figure 4: Distance between a spike train and the spike train with noise added to the input. The different lines correspond to different amounts of noise: standard deviation 2%, 4%, 6%, 8%, and 10% of stimulus standard deviation (lower to upper curve). The crosses indicate the limits $t_c \rightarrow 0$ and $t_c \rightarrow \infty$. Note the minimum of the distance at intermediate t_c .

trains where the average distance was constant across t_c . For the integrate-and-fire neuron, noise will both shift as well as insert and delete spikes. The membrane time constant introduces temporal correlations, leading to a minimal distance when t_c is roughly of the order of the integration time constant. Why the precise value is larger than the membrane time constant is not clear.

The dependence of the distance on the amount of added noise varies for different values of t_c (see Figure 5). Unfortunately, there is no obvious analytical approach to determine the distance as a function of the noise for the integrate-and-fire model; instead we rely on simulations. We found approximately a power law relation between distance and noise, with the exponent depending on t_c . For small t_c , it holds that $D^2 \propto \sqrt{\sigma_{add}}$, whereas for large t_c , one has $D^2 \propto \sigma_{add}^{2, \dots, 4}$. However, for a large region of intermediate t_c , roughly corresponding to the valley region in Figure 4, one has $D^2 \propto \sigma_{add}$.

3.1 Estimating Intrinsic Noise. The smooth dependence of the distance on the noise can be used to measure the intrinsic noise of a neuron. To this end, we assume that there is an intrinsic noise source in the neuron that is additive to the input (Tuckwell, 1988; Gerstein & Mandelbrot, 1964). By




Figure 5: Same data as Figure 4, but here the distance is plotted versus the noise added to the stimulus. For different t_c there are different power laws between the distance and the noise. The standard deviation of the noise is denoted as a fraction of the stimulus standard deviation. Because of the different y-scaling, the plots were split for different t_c values: (top) $t_c = 0, 1, 10$; (middle) $t_c = 100, 1000, 10,000$ (linear regime); (bottom) $t_c = 10^5, 10^6, \infty$.

mixing additional nonrepeating noise with the input and extrapolating the distance, the intrinsic noise can be estimated (see Figure 6). The standard deviation at the input of the spike generator is $\sigma_{total} = \sqrt{\sigma_{add}^2 + \sigma_{intrinsic}^2}$; therefore, the distance is best considered as a function of σ_{add}^2 . In practice, one plots D^4 versus σ_{add}^2 and fits a straight line. The intrinsic noise is reconstructed from the intercept of the line with the x -axis (see Figure 6B).

This method is borrowed from electrical engineering and psychophysical studies. There one measures, for instance, the detection threshold of a visual stimulus as noise is added to the stimulus (Pelli, 1990; Lu & Doshier, 1999). This detection threshold is commonly proportional to the signal-to-noise ratio of the stimulus, which then allows determination of the intrinsic noise.




Figure 6: Using the distance to estimate the intrinsic noise of a cell. (A) Model of the cell. The intrinsic noise in the cell is assumed to be additive to the input. (B) The distance is plotted versus the variance of noise added to the stimulus. By adding noise, the distance to the original response increases (circles). Note that at zero added noise, the distance is nonzero because the cell will not reproduce the exact same spike train due to the presence of the intrinsic noise. Extrapolation yields an estimate for the intrinsic noise.

To test the approach, we simulated an integrate-and-fire neuron with an intrinsic noise source and tried to estimate the amount of intrinsic noise by extrapolation. Because the linear regression is not a good fit for all t_c (see Figure 5), the quality of the fit was measured with χ^2 , and the result with minimal χ^2 was selected. The method works well for small amounts of intrinsic noise. An intrinsic noise with a standard deviation of 1% (as

compared to the stimulus standard deviation) was estimated as 1.25% ($t_c = 600$ ms, shown in Figure 6B); an actual value of 5% was estimated as 6%, 10% was estimated as 12.6%, and 20% was estimated as 33%. Thus, the method slightly overestimates the intrinsic noise. At higher levels of added or intrinsic noise, the fitting deteriorates as the linear relation between σ^2 and D^4 gets higher-order corrections. Nevertheless, we obtain a reasonable estimate for the intrinsic noise, which is otherwise hard to access.

4 Conclusion

We have introduced a distance measure that computes the dissimilarity between two spike trains. To calculate the distance, filter both spike trains, and calculate the integrated squared difference of the two trains. The simplicity of the distance allows for an analytical treatment of simple cases. Numerical implementation is straightforward and fast.

The distance interpolates between counting noncoincident spikes and counting the squared difference in total spike count. In order to compare spike trains with different rates, total spike count can be used (large t_c). However, for spike trains with similar rates, the difference in total spike number is not useful, and coincidence detection is sensitive to noise. Instead, intermediate values of t_c , somewhat longer than the membrane time constant, are optimal.

The distance uses a convolution with the exponential function. This has an interpretation in physiological terms. Indeed, among other distances (see section 1), our distance does not seem the most unlikely one to be implemented physiologically. For short and medium t_c , the convolution can be interpreted as postsynaptic potentials in a higher-order neuron. For longer t_c , slower second messenger or calcium-induced currents seem more appropriate. It would be interesting to see if such detection schemes are implemented biologically.

As an alternative measure, one could convolve the spikes with a square window. In that case, the situation becomes somewhat similar to binning followed by calculating the Euclidean distance between the number of spikes in the bins. But in standard binning, the bins are fixed on the time axis; therefore, two different spike trains yield identical binning patterns as long as spikes fall in the same bin. However, with the proposed distance (convolving with either square or exponential), this does not happen; the distance is zero only if the two spike trains are fully identical (assuming t_c is finite).

The distance squared is related to the distance measure introduced by Victor and Purpura (1996, 1997). One difference with this work is that their displacement distance was linear in the time difference, although they did suggest the use of an exponential displacement distance (see equation 2.8). Another difference is that the distance introduced here is explicitly embedded in Euclidean space, which makes it less general but easier to analyze than the Victor and Purpura distance.

Interestingly, the distance is related to stimulus reconstruction techniques, where convolving the spike train with the spike-triggered average yields a first-order reconstruction of the stimulus (Rieke et al., 1996). Here the exponential corresponds roughly to the spike-triggered average, and the filtered spike trains correspond to the stimulus (the exponentials are now attached to the other side of the spikes, but that does not change the distance). The distance thus approximately measures the difference in the reconstructed stimuli (see equation 2.3). This might well explain the linearity of the measure for intermediate t_c . It also suggests that the distance measure might be refined by using the actual spike-triggered average instead of an exponential.

A possible application for the distance measure is the insect olfactory system, which uses a timing-based code to distinguish odors (MacLeod et al., 1998). There it was shown, using the Victor and Purpura distance, that information was encoded in the temporal structure of the spike train but not in the mean rate. Another application is the measurement of the intrinsic noise of a neuron or network, which is possible because the distance varies smoothly as noise is added to the input. For sensory systems, stimulation with "natural" noisy stimuli has become quite common. Although natural stimuli are relevant in determining the system's response, the characterization of the noise in the system is less straightforward (Reich, Victor, Knight, Ozaki, & Kaplan, 1997). Measuring the intrinsic noise could be helpful. We stress that the intrinsic noise is a quantity that is otherwise difficult to measure experimentally. For instance, measurements of subthreshold fluctuations in the membrane potential fail to detect noise in the spike generator itself. By adding noise to the stimulus, one can determine this intrinsic noise, which gives an effective description of the neuron's variability.

Acknowledgments

It is a pleasure to thank Alex Bäcker, Bob Cudmore, Kate Macleod, Rob Smith, and Ken Sugino for discussions and Larry Abbott and Paul Tiesinga for helpful comments on the manuscript. I thank one of the referees for guiding me toward equation 2.6 and indicating its application to Poisson trains and pointing out the relation of the distance to stimulus reconstruction. Supported by the Sloan Foundation.

References

- Geisler, W. S., Albrecht, D. G., Salvi, R. J., & Saunders, S. S. (1991). Discrimination performance of single neurons: Rate and temporal information. *Journal of Neurophysiology*, 66, 334–362.
- Gerstein, G. L., & Mandelbrot, B. (1964). Random walk models for the spike activity of a single neuron. *Biophysical Journal*, 4, 41–68.

- Kistler, W. M., Gerstner, W., & van Hemmen, J. L. (1997). Reduction of the Hodgkin-Huxley equations to a single variable threshold model. *Neural Computation*, *9*, 1015–1046.
- Lu, Z.-L., & Doshier, B. A. (1999). Characterizing human perceptual inefficiencies with equivalent internal noise. *Journal of the Optical Society of America A*, *16*, 764–778.
- MacLeod, K., Backer, A., & Laurent, G. (1998). Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature*, *395*, 693–698.
- Pelli, D. G. (1990). The quantum efficiency of vision. In C. Blakemore (Ed.), *Vision: Coding and efficiency* (pp. 9–24). Cambridge: Cambridge University Press.
- Reich, D. S., Victor, J. D., Knight, B. W., Ozaki, T., & Kaplan, E. (1997). Response variability and timing precision of neuronal spike trains in vivo. *Journal of Neurophysiology*, *77*, 2836–2841.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R., & Bialek, W. (1996). *Spikes: Exploring the neural code*. Cambridge, MA: MIT Press.
- Tuckwell, H. C. (1988). *Introduction to theoretical neurobiology*. Cambridge: Cambridge University Press.
- Victor, J. D., & Purpura, K. P. (1996). Nature and precision of temporal coding in visual cortex: A metric-space analysis. *Journal of Neurophysiology*, *76*, 1310–1326.
- Victor, J. D., & Purpura, K. P. (1997). Metric-space analysis of spike trains: Theory, algorithms and application. *Network: Comput. Neural Syst.*, *8*, 127–164.

Received September 24, 1999; accepted July 3, 2000.