

Short communication

Drawing inferences from Fano factor calculations

Uri T. Eden*, Mark A. Kramer

Department of Mathematics and Statistics, Boston University, 111 Cummington St, Boston, MA 02215, USA

ARTICLE INFO

Article history:

Received 23 February 2010

Received in revised form 12 April 2010

Accepted 14 April 2010

Keywords:

Fano factor

Poisson spiking

Spike train analysis

Hypothesis testing

ABSTRACT

An important characterization of neural spiking is the ratio of the variance to the mean of the spike counts in a set of intervals—the Fano factor. For a Poisson process, the theoretical Fano factor is exactly one. For simulated or experimental neural data, the sample Fano factor is never exactly one, but often appears close to one. In this short communication, we characterize the distribution of the Fano factor for a Poisson process, allowing us to compute probability bounds and perform hypothesis tests on the distribution of recorded neural spike counts. We show that for a Poisson process the Fano factor asymptotically follows a gamma distribution with dependence on the number of observations of spike counts, and that convergence to this asymptotic distribution is fast. The analysis provides a simple method to determine how close to 1 the computed Fano factor should be and to formally test whether the observed variability in the spiking is likely to arise in data generated by a Poisson process.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

The number of action potentials (or “spikes”) generated by a neuron in an interval often exhibits high variability. For example, an individual neuron presented multiple times with the same stimulus will typically produce a different number of spikes following each presentation. Remarkably, a consistency often emerges in the count statistics of neural spiking over many stimulus presentations: namely, the ratio between the sample mean and sample variance of the spike counts over fixed intervals remains constant. This ratio of the spike count variance to the mean—the Fano factor (Fano, 1947)—has been found to span a range of values near one (Softky and Koch, 1993; de Ruyter van Steveninck et al., 1997; Kara et al., 2000; Maimon and Assad, 2009; Shadlen and Newsome, 1998; Tolhurst et al., 1983).

The importance of the Fano factor extends beyond a simple characterization of the spiking data. This statistic can also be used to draw inferences about how the observed neural variability compares with the expected variability for a Poisson process, the most common statistical model of neural spiking. A central feature of the Poisson process is that the number of spikes fired in any time interval is random, with a Poisson distribution. For a homogeneous Poisson process, for which the rate of spiking is constant in time, the spike counts in any set of time intervals of fixed size are all independent identically distributed Poisson random variables. For an inhomogeneous Poisson process, for which the rate of spiking can change as a function of time, the spike counts over multiple inter-

vals that have the exact same rate function will also be independent identically distributed Poisson random variables. Since the theoretical variance and mean of a Poisson distribution are equal, the theoretical Fano factor for the spike counts is *exactly* equal to one. Yet, we never expect this precision in any sample from a Poisson process, let alone in noisy neural data or in computational models of neural activity. Instead, we compute the sample Fano factor from these data and expect—if the data are Poisson—a value near one. Typical techniques to perform this estimate include regression analysis to estimate the linear relationship between the spike count variance and mean (Tolhurst et al., 1983; Gur et al., 1997; Gershon et al., 1998) and direct calculation of the Fano factor (spike count variance over spike count mean) from the data (Teich et al., 1997; Kara et al., 2000).

The characterization of neural data exhibiting a Fano factor “close to” one as Poisson (and rejection of this characterization when the Fano factor is “distant” from one) is unsatisfying. Ideally, we would prefer to make these claims with an associated measure of uncertainty. In this short communication, we propose a simple statistical method to determine how close to 1 the computed Fano factor should be and to formally test against the null hypothesis of Poisson spiking based on the observed Fano factor. We describe this test here, show simulation results, and provide simple routines in MATLAB to perform this test.

2. Results

Let N_i for $i = 1, \dots, n$ be a collection of independent identically distributed samples from a Poisson distribution with parameter λ . These samples might arise as counts of the number of spikes fired over repeated trials of an experiment, as binned spike counts over

* Corresponding author. Tel.: +1 617 353 9553; fax: +1 617 353 8100.
E-mail address: tzvi@bu.edu (U.T. Eden).

fixed intervals for a homogeneous Poisson process, or any number of other experimental conditions. The Fano factor for these data is given by the ratio of the sample variance to the sample mean:

$$\hat{F} = \frac{\frac{1}{n-1} \sum_{i=1}^n (N_i - \bar{N})^2}{\bar{N}}, \quad \text{where } \bar{N} = \frac{1}{n} \sum_{i=1}^n N_i.$$

Defined this way, the Fano factor is a statistic, with a sampling distribution that defines the probability of \hat{F} taking on any particular value. By comparing the computed value of the Fano factor for a given dataset to critical values of this sampling distribution, we can perform hypothesis tests and compute significance levels (p -values) for the null hypothesis of Poisson spiking.

The sampling distribution of \hat{F} can depend on two factors, the Poisson parameter, λ , which also defines the expected number of spikes per sample, and the number of samples obtained, n . Previous analyses (Hoel, 1943; Kathirgamatamby, 1953) have shown that as λ and n both tend to infinity, the sampling distribution of $(n-1)\hat{F}$ tends to a Chi-square distribution with $(n-1)$ degrees of freedom, which is equivalent to a gamma distribution with shape parameter $(n-1)/2$ and a scale parameter equal to 2. From there, it follows from the scaling property of the gamma distribution that the sampling distribution of the Fano factor itself tends to a gamma distribution with shape parameter $(n-1)/2$ and scale parameter $2/(n-1)$, $\hat{F} \sim \Gamma((n-1)/2, 2/(n-1))$.

This result provides a simple method for constructing bounds for the Fano factor under the Poisson assumption, and for performing hypothesis tests to reject the Poisson hypothesis and compute significance levels. In particular, if the data were generated by a Poisson process, with 0.95 probability the computed Fano factor would lie in the interval $[\Gamma_{0.025}, \Gamma_{0.975}]$, where Γ_α is the critical value of the $\Gamma((n-1)/2, 2/(n-1))$ distribution with probability mass α to the left of that value. For example, if the Fano factor were computed from $n=50$ trials, we would expect the value to fall in the range $[0.64, 1.43]$ with probability 0.95. Therefore, a Fano factor of 1.4 computed over 50 trials should not be considered exceptionally strong evidence that the increments are more variable than a Poisson process. These bounds can be computed trivially in any statistical computing software. For example, in MATLAB, the 0.95 bounds are generated by the command `gaminv([.025,.975],(n-1)/2,2/(n-1))`.

This result can also be used to compute a hypothesis test for whether the data could have arisen from a Poisson process. In this case, the null hypothesis H_0 , that the data are Poisson, is equivalent to the statement that the theoretical Fano factor is equal to 1. We can compare this to an upper tailed alternative hypothesis that the increments are more variable than the Poisson, a lower tailed alternative that the increments are less variable, or a two-tailed alternative that the variability is either larger or smaller than for a Poisson. The test statistic is just the empirical Fano factor, and the approximate p -value for the test is $\Pr(\Gamma > \hat{F})$ for an upper tailed test, $\Pr(\Gamma < \hat{F})$ for a lower tailed test, and $2 \times \min(\Pr(\Gamma > \hat{F}), \Pr(\Gamma < \hat{F}))$ for a two-tailed test, where Γ is a random variable with $\Gamma((n-1)/2, 2/(n-1))$ distribution and \hat{F} is the computed value of the Fano factor. Once again, these quantities are simple to compute in standard statistical and computing software. For example, to compute the p -value for the upper tailed test that the increments are more variable than a Poisson in MATLAB, the command is `1-gamcdf(F, (n-1)/2,2/(n-1))`. The power of any of these tests depends on the specific form of the true alternate hypothesis, that is, the specific history dependent structure of the neuron. Previous analyses (Kathirgamatamby, 1953) have examined the power of these tests for a variety of alternatives.

All these methods rely on the gamma approximation to the sampling distribution of the Fano factor. How large do λ and n need to be for the gamma approximation to be accurate? Early research (Sukhatme, 1938) used random number tables to examine deviations from the gamma distribution in low rate, small sample instances. However, it is now possible through simulation to generate highly accurate empirical estimates of the true distribution of the Fano factor for any λ and n pair.

For values of λ ranging from 0.1 to 100 expected spikes per sample (in steps of 0.1) and values of n ranging from 2 to 100 samples, we simulated 100,000 sets of n samples from a Poisson distribution with parameter λ , and computed the resulting Fano factor from each set. This provided an empirical estimate of the distribution of the Fano factor for every λ and n . The top panel of Fig. 1 shows a scaled histogram of the resulting empirical distribution for the case where $\lambda=10$ and $n=50$. In comparison, the probability density for the $\Gamma(49/2, 2/49)$ distribution is shown with the solid black line. 95% bounds for the value of the Fano factor under the Poisson condition are shown as the transition from white to gray bars and back for the empirical distribution and as vertical black lines for the approximate gamma distribution. From the empirical distribution, 95% of the computed Fano factors ranged between 0.63 and 1.4. This suggests that with as many as 50 samples from a neuron firing as a Poisson process with physiological firing rates, it is not surprising to get sample Fano factor values notably larger or smaller than 1.

From Fig. 1, it is clear that the empirical and approximate gamma distribution deviate only slightly from one another. In particular, the empirical distribution has a slightly smaller variance, and a slightly larger skewness than the gamma approximation. The bounds on the Fano factor based on the empirical and approximate distributions are also in close agreement, with the lower bounds nearly identical, and the upper bound for the gamma approximation providing a slightly conservative estimate.

The bottom panels of Fig. 1 show raster plot realizations of two distinct point processes. The data in the left panel was generated from 50 trials of a Poisson process with rate $\lambda=10$ spikes/trial. The Fano factor of the number of spikes/trial for this data was 1.32, a value far from 1 but still within the 95% bounds for the empirical distribution of such Poisson processes. The data in the right panel was generated not as a Poisson process, but as a renewal process with $\Gamma(0.5, 12)$ interspike intervals. This process has an expected Fano factor larger than 1, suggesting more variability in the number of spikes per trial. The computed Fano factor for this data is 1.63, outside the 95% bounds for the empirical distribution for a Poisson. Visually, it is very difficult to distinguish which dataset is compatible or incompatible with a Poisson process. However, by comparing the computed Fano factors to the computed bounds, or by performing an explicit test, it is possible to differentiate these processes and express degrees of confidence.

Fig. 2 shows 0.025 and 0.975 quantiles of the empirical distributions of Fano factors for the simulated range of values for λ and n , as well as the 0.025 and 0.975 critical values for the gamma approximation. The empirical upper and lower bounds as a function of number of samples are shown as the lines above and below the value 1, respectively, with different line styles indicating different expected numbers of spikes. Clearly, as the number of samples increases, the bounds about the expected value of 1 for the Fano factor narrows. However, even at relatively large values of n , there is still a great deal of variability in the observed values of the Fano factor. In comparison, the 0.025 and 0.975 critical values of the $\Gamma((n-1)/2, 2/(n-1))$ distribution are shown with thick dashed lines.

For large values of n and λ , the empirical and approximate gamma bounds are in close agreement. For small values of n and λ , the approximate gamma bounds are overly conservative. This is most noticeable in the upper bound when $n < 10$. Previous analy-




Fig. 1. Upper: The empirical (white and gray bars) and approximate (gamma, black curve) distributions of the Fano factor are in close agreement. The empirical distribution was generated for Poisson data with $\lambda = 10$ and $n = 50$. The 95% bounds for the Fano factor occur at the transition from the white to gray bars in the empirical distribution and between the vertical black lines in the approximate gamma distribution. Lower: Examples of data generated from a Poisson process (left) and from a renewal process with theoretical Fano factor greater than 1 (right).

ses (Hoel, 1943) have shown that the ratio of the variance of the true distribution of a scaled version of the Fano factor to its gamma approximation is less than 1. This suggests that p -values computed using the gamma approximation will be slightly larger than using the empirical distribution, and the chance of incorrectly rejecting a null hypothesis that the data is Poisson will not increase. However, the power of such a test may decrease for very small sample sizes. In this case, a test based on the empirical distribution, which can now be easily calculated, would be preferable. However, in most realistic neural data analyses such that $n \geq 20$ and $\lambda \geq 1$, tests based on the gamma approximation and empirical distribution will be nearly identical.

3. Discussion

The Fano factor provides one method to describe the variability in spiking activity and to relate it to the variability associated with a Poisson process. Multiple other procedures, both exact (Amarasingham et al., 2006) and approximate (Jackson and Redish, 2007; Nawrot et al., 2008), exist to determine whether the observed variability of spike counts agrees with that expected for a Poisson process. The measure proposed here differs from these measures in the following ways. First, the statistical test acts directly on the Fano factor (and does not require the computation of any additional measures). Second, the analytic procedure allows a simple computation (or tabulation) of probability bounds and p -values for specific hypothesis tests.

Another natural approach for characterizing the variability of the Fano factor is the bootstrap. An implementation of this technique could involve resampling with replacement the observed N_i to create a large number of surrogate data sets, each with the same number of data points as the original data set. The Fano factor is then computed for each surrogate data set to create an empirical distribution of Fano factors. If the 95% confidence bounds based on this distribution do not include 1, then we reject the null hypothesis of Poisson spiking. We note that generating the surrogate data for the bootstrap procedure may be computationally expensive. Additionally, further statistical analysis is required to determine the number of original data points required for the bootstrap distribution to be accurate.

It is important to note that variability of spike counts is only one of a multitude of features that can be used to determine whether spike train data is Poisson in nature (Baddeley et al., 1997). Examination of these other features can be particularly important in some cases. For example, a non-Poisson renewal process can produce of Fano factor near or exactly equal to 1. In this case, any procedure based on the Fano factor alone would not be able to reject a




Fig. 2. Comparison of the empirical and approximate (dashed) bounds of the Fano factor for different values of λ and n . The approximate bounds are overly conservative for small values of n .

Poisson model, but tests based on other statistics could. For example, the increments of a Poisson processes should be independent, and testing methods based on the autocorrelation of the observed increments can be used to identify non-Poisson dependence structure. Additionally, the spectrum of a Poisson process should be flat, and testing methods related to spectral estimators can be used to identify frequency-dependent variation. As with any statistical characterization, examining the data through multiple measures provides a more complete view of the structure present therein.

One attractive feature of the Fano factor is that by varying the length of the time intervals over which the spike counts are collected, we can investigate the variability of the spiking over any time scale (Teich et al., 1997). Typically, the choice of an interval length is made based on the time scale over which one believes the variability of the process to differ most from Poisson variability. However, the analysis presented here suggests that by changing the interval size, one can also adjust the quality of the gamma approximation and the power of the test based on the Fano factor. For example, for a collection of 10 intervals from a homogeneous Poisson process averaging 20 spikes per interval, we can choose to break each interval in half, providing 20 intervals averaging 10 spikes per interval, where the gamma approximation is very good. It is important to note, however, that no matter how the data is divided, the product of $\lambda \times n$ is preserved. Hoel's (1943) work suggests that the ratio of the second moment of the sampling distribution of the Fano factor to the gamma approximation only depends on this product, so up to second order, choice of interval length has no effect.

Overall, the Fano factor is one of the most ubiquitously computed descriptive statistics of spike train variability. The result we describe provides a simple and powerful approach to express the statistical confidence for inferences based on the Fano factor.

Acknowledgements

This work was partially supported by grant IIS-0643995 from the National Science Foundation to U.T.E. M.A.K. holds a Career

Award at the Scientific Interface from the Burroughs Wellcome Fund.

References

- Amarasingham A, Chen TL, Geman S, Harrison MT, Sheinberg DL. Spike count reliability and the poisson hypothesis. *J Neurosci* 2006;26:801–9.
- Baddeley R, Abbott LF, Booth MCA, Sengpiel F, Freeman T, Wakeman EA, et al. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc R S Lond B* 1997;264(1389):1775–83.
- Fano U. Ionization yield of radiations. II. The fluctuations of the number of ions. *Phys Rev* 1947;72(1):26–9.
- Gershon ED, Wiener MC, Latham PE, Richmond BJ. Coding strategies in monkey V1 and inferior temporal cortices. *J Neurophys* 1998;79(3):1135–44.
- Gur M, Beylin A, Snodderly DM. Response variability of neurons in primary visual cortex (V1) of alert monkeys. *J Neurosci* 1997;17(8):2914–20.
- Hoel PG. On indices of dispersion. *Ann Math Stat* 1943;14(2):155–62.
- Jackson J, Redish AD. Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. *Hippocampus* 2007;17(12):1209–29.
- Kara P, Reinagel P, Reid RC. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* 2000;27:635–46.
- Kathirgamatamby N. Note on the Poisson index of dispersion. *Biometrika* 1953;40:225–8.
- Maimon G, Assad JA. Beyond poisson: increased spike-time regularity across primate parietal cortex. *Neuron* 2009;62:426–40.
- Nawrot MP, Boucsein C, Rodriguez Molina V, Riehle A, Aertsen A, Rotter S. Measurement of variability dynamics in cortical spike trains. *J Neurosci Methods* 2008;169(2):374–90.
- de Ruyter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W. Reproducibility and variability in neural spike trains. *Science* 1997;275:1805–8.
- Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 1993;13(1):334–50.
- Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 1998;18:3870–96.
- Sukhatme PV. On the distribution of χ^2 in samples of the Poisson series. *J R S Soc Suppl* 1938;5:75–9.
- Teich MC, Heneghan C, Lowen SB, Ozaki T, Kaplan E. Fractal character of the neural spike train in the visual system of the cat. *J Opt Soc Am* 1997;14(3):529–46.
- Tolhurst DJ, Movshon JA, Dean AF. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 1983;23:775–85.