Semi-parametric Bayesian entropy estimation for binary spike trains

Evan Archer, Il Memming Park, & Jonathan Pillow

Abstract: The set of possible neural response patterns is frequently so large that its distribution cannot be reliably estimated from limited data. Consequently, information and entropy estimation for neural data presents a significant challenge which has been met by a diverse literature. Most entropy estimators in this literature, however, are general purpose in that they are designed to work on as broad a class of data-generating distributions as possible. For neural applications all such general-purpose estimators have a critical weakness: they ignore by design much of our strong prior knowledge about the structures of spike trains. Neural response patterns, however, are not arbitrary: we can apply our prior knowledge about the basic statistical structure of spike trains to entropy estimation. Here, we augment the nonparametric Bayesian entropy estimation method [4] by incorporating a simple, parametric model of spike trains. Intuitively, we wish to incorporate our prior knowledge that spikes are rare events, and we assign lower prior probability to words with more spikes. Mathematically, we model a spike word as a vector of independent Bernoulli random variables, each with a probability p of firing. Under this model, for typical values of p, very sparse vectors are much more likely than those with many spikes. Alone, this simple model does not provide a good method for entropy estimation, as it cannot flexibly account for data drawn outside the model class. However, by "centering" a Dirichlet process on this parametric model, we obtain a semi-parametric model that can model arbitrary discrete distributions. We derive a computationally efficient entropy estimator under the model, and for real data, we show that this model outperforms conventional estimators.

Additional Information: We focus on the problem of estimating the entropy of the joint response distribution of N simultaneously-recorded spiking neurons. As a prior, we propose an independent, binomial spiking model, $G|p := \mathrm{Binom}(p, N)$. The distribution $G|p$ encodes our prior knowledge on the 2^N possible spiking patterns given just the marginal firing probability p (per bin). Taken alone, $G|p$ is an impoverished model of spike trains, but our goal is not to estimate $G|p$ but to use it to inform our distribution over the full $(2^N - 1)$-dimensional simplex by using it as the *base measure* of a Dirichlet distribution. Under our final model, $\boldsymbol{\pi} \sim \mathrm{Dir}(\alpha G)$, while spike train observations x are multinomial given $\boldsymbol{\pi}$ (See Fig. 1). Due to the conjugacy of Dirichlet and multinomial, the posterior distribution given observations is $\boldsymbol{\pi}|\text{data} \sim \mathrm{Dir}(\alpha_1 + n_1, \ldots, \alpha_{\mathcal{A}} + n_{\mathcal{A}})$, where n_i is the number of observations for the i-th spiking pattern, and $\mathcal{A} = 2^N$. For $\boldsymbol{\pi} \sim \mathrm{Dir}(\alpha_1, \alpha_2, \ldots, \alpha_{\mathcal{A}})$, such that $\sum_{i=1}^{\mathcal{A}} \alpha_i = A$, and letting $\vec{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_{\mathcal{A}})$, the expected entropy is given by [2],

$$\mathbb{E}[H(\boldsymbol{\pi})|\vec{\alpha}] = \psi_0(A+1) - \sum_{i=1}^{\mathcal{A}} \frac{\alpha_i}{A} \psi_0(\alpha_i + 1) \tag{1}$$

where ψ_0 is the digamma function. We compute the expected entropy without explicitly representing the 2^N values of α_i by exploiting the redundancy of binomial probability. The probability depends only on the total number of spikes. With careful bookkeeping, we can efficiently evaluate (1).

Prior design: Following [1,4], we construct an approximately flat prior on entropy using the approximation,

$$P(\alpha) = \frac{\mathrm{d}}{\mathrm{d}\alpha} \mathbb{E}[H(\boldsymbol{\pi}|\alpha)]. \tag{2}$$

Bayes least squares estimator: Given the prior distributions $P(\alpha)$ and $P(p)$, the Bayesian entropy estimate is given by:

$$\hat{H}_{BD} = \mathbb{E}[H|\mathbf{x}] = \iint \mathbb{E}\left[H|\alpha, p\right] P(\alpha, p|\mathbf{x}) P(\alpha, p) \mathrm{d}\alpha \, \mathrm{d}p. \tag{3}$$

The posterior mean (3) describes the final form of our estimator: the Binomial-Dirichlet (BD) entropy estimator.

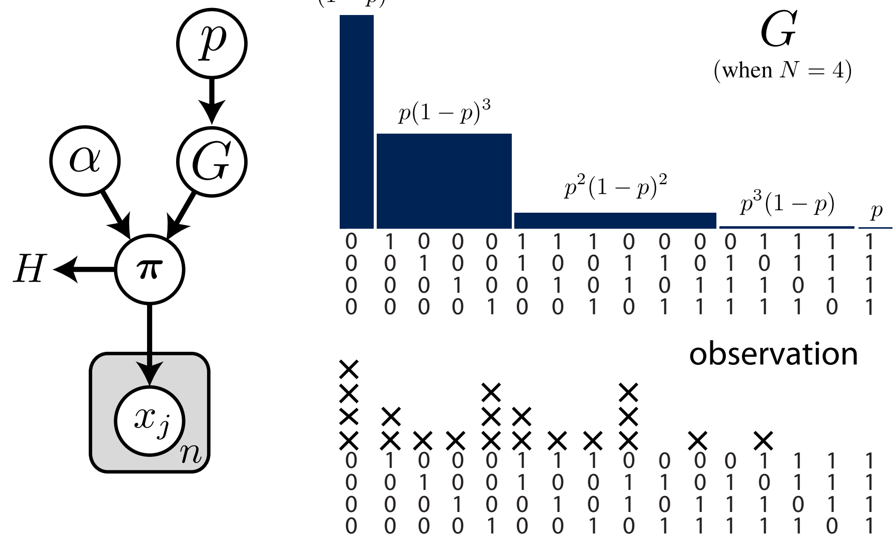

Figure 1: Graphical model illustrating the ingredients for Bayesian entropy estimation. Arrows indicate conditional dependencies between variables. Note that for large N, it is not practical to represent the distribution with a histogram as shown, since there are 2^N possible patterns.

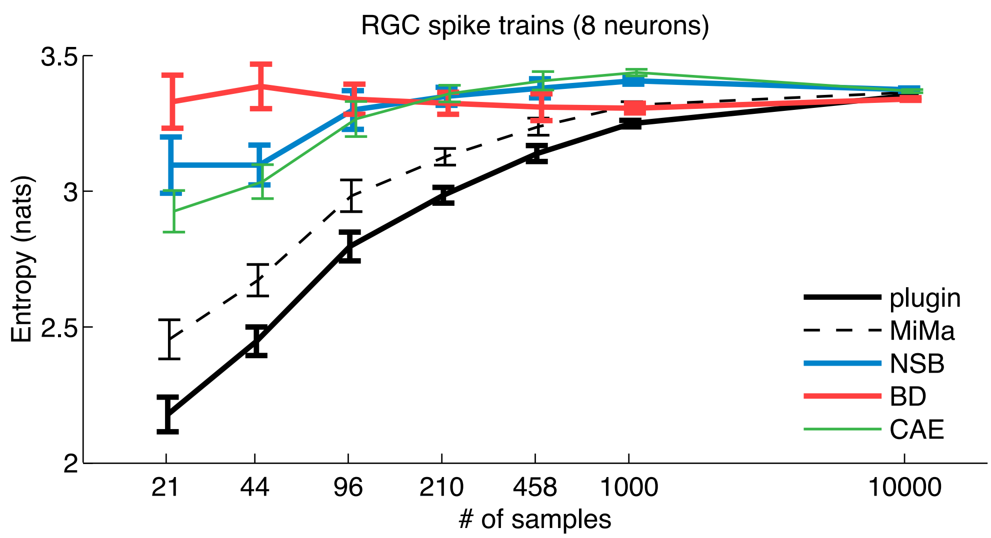

Figure 2: Convergence of entropy estimators with sample size for 8 simultaneously recorded retinal ganglion cells. In the legend, we abbreviate "MiMa" is for "Miller-Maddow" (see, [5]), "CAE" for Coverage Adjusted Estimator [3], and "NSB" for Nemenman-Shafee-Bialek [1]. Lines are means computed across 16 random subsamples of the dataset. Error bars are standard errors.

Significance: We design an efficiently-computable entropy estimator specifically for neural spike trains. Our estimator incorporates domain knowledge about the statistical structure of spike trains.
Relevance: Entropy quantifies information encoded in spike trains. Our estimator allows practical entropy estimation with less data.
Originality: To the best of our knowledge, this is the first semi-parametric Bayesian entropy estimator.

References: [1] Nemenman, Shafee & Bialek 2001. [2] Wolpert, Wolf 1995., [3] Vu, Yu, & Kass 2007. [4] Archer, Park, Pillow 2012. [5] Paninski 2003.