Linear Fisher Information in a network of LNP spiking neurons

Beck JM, Ma WJ, Pouget A
Dept of Brain and Cognitive Science. Univ. of Rochester. Rochester, NY 14627

Deriving an analytical expression for the rate of Fisher information transfer across multiple layers of recurrently connected spiking neurons is a critical step toward understanding the neural basis of a wide variety of problems such as perceptual learning, decision making, attention and optimal Bayesian inference. Here, we address this issue by deriving an expression for linear Fisher information loss in a network of Linear-Nonlinear-Poisson (LNP) spiking neurons receiving spatially correlated spiking inputs. By linear Fisher information, we mean the Fisher information available to an observer with knowledge of the first and second order statistics of the output layer of the network. Linear Fisher information is particularly important in computational neuroscience as linear Fisher information is recoverable by delta rule learning and linear Fisher information has been shown to represent the quantity of information available to decoder implemented by an attractor neural network. Our analysis is based on linearization of the networks dynamics, and yield a surprisingly simple expression for the rate of information loss which takes into account both the spatial and temporal covariance structure of the spike trains. This equation allows us to determine precisely where information loss occurs. We show that critical factors are:

1- The feedforward weights: information loss is roughly inversely proportional to feedforward connection strength provided noninvertible feedforward connection matrices (such as the ones between the LGN and orientation selective V1 cells) are tuned so as to avoid filtering out informative collections of neurons.
2- The nonlinear activation function: any activation threshold will lose information but, more importantly, the shape of the activation function is critical. The most informative neurons should respond where the derivative of the square root of the activation function is the highest in order to minimize information loss. This indicates the utility of a gain control mechanism.
3- Large recurrent weights: the information loss increases with the norm of the recurrent weights. This can be balanced by a comparable increase in feedforward connection strength or optimizing output rate
4- Amplitude of the output activity: surprisingly, information loss may grow with firing rate, i.e., large firing rates are not necessarily advantageous.
5-The Poisson step: this is the only stochastic step in LNP neurons and, not surprisingly, it leads to information loss. However, this factor can be minimized if the number of neurons is very large and the information is finite. In other words, this source of information loss is minimized when the neurons are highly correlated. This implies that, once a redundant code has been established, a the number of neurons is subsequent may be fixed with only small consequences to information loss.

This work also suggests an interesting perspective on the origin of noise in the brain. Information loss is equivalent to adding noise. When we say that the nervous system is ‘noisy’, we are effectively saying that it looses information. Interestingly, the main source of noise in the brain is thought to be the Poisson step (or its underlying cause, such as chaotic dynamics). Yet, according to our analysis, this source of information loss (or ‘noise’) is in fact easy to minimize across multiple layers of cortex. Thus, it seems likely that most of the ‘noise’ in the cortex has little to do with the Poisson step but is, in fact, related to accurate perceptual filtering and task relevant stimulus identification. There is a natural Bayesian interpretation of this effect. The brain cannot know with certainty the statistics of its inputs. As a result, it learns approximate distributions which are embedded in the connectivity. The mismatch between the presumed and true statistics could potentially be the main source of noise.

Acknowledgments: This work was supported by NSF, NIDA and the James S. McDonnell Foundation.