

Variational Information Maximization for Neural Coding

Felix Agakov¹ and David Barber²

¹ University of Edinburgh, 5 Forrest Hill, EH1 2QL Edinburgh, UK
felixa@inf.ed.ac.uk, www.anc.ed.ac.uk

² IDIAP, Rue du Simplon 4, CH-1920 Martigny Switzerland,
www.idiap.ch

Abstract. Mutual Information (MI) is a long studied measure of coding efficiency, and many attempts to apply it to population coding have been made. However, this is a computationally intractable task, and most previous studies redefine the criterion in forms of approximations. Recently we described properties of a simple lower bound on MI [2]. Here we describe the bound optimization procedure for learning of population codes in a simple point neural model. We compare our approach with other techniques maximizing approximations of MI, focusing on a comparison with the Fisher Information criterion.

1 Introduction

The problem of encoding real-valued stimuli \mathbf{x} by a population of neural spikes \mathbf{y} may be addressed in many different ways. The goal is to adapt the parameters of any mapping $p(\mathbf{y}|\mathbf{x})$ to make a desirable population code for a given set of patterns $\{\mathbf{x}\}$. There are many possible desiderata. One could be that *any* reconstruction based on the population should be accurate. This is typically handled by appealing to the Fisher Information which, with care, can be used to bound mean square reconstruction error. Another approach is to bound the probability of a correct reconstruction. Here we consider maximizing the amount of information which the spiking patterns \mathbf{y} contain about the stimuli \mathbf{x} (e.g. [7], [5]). The fundamental information theoretic measure in this context is the mutual information

$$I(\mathbf{x}, \mathbf{y}) \equiv H(\mathbf{x}) - H(\mathbf{x}|\mathbf{y}), \quad (1)$$

which indicates the decrease of uncertainty in \mathbf{x} due to the knowledge of \mathbf{y} . Here $H(\mathbf{x}) \equiv -\langle \log p(\mathbf{x}) \rangle_{p(\mathbf{x})}$ and $H(\mathbf{x}|\mathbf{y}) \equiv -\langle \log p(\mathbf{x}|\mathbf{y}) \rangle_{p(\mathbf{x}, \mathbf{y})}$ are marginal and conditional entropies respectively, and the angled brackets represent averages over all variables contained within the brackets.

The principled information theoretic approach to learning neural codes maximizes (1) with respect to parameters of the encoder $p(\mathbf{y}|\mathbf{x})$. However, it is easy to see that in large-scale systems exact evaluation of $I(\mathbf{x}, \mathbf{y})$ is in general computationally intractable. The key difficulty lies in the computation of the conditional entropy $H(\mathbf{x}|\mathbf{y})$, which is tractable only in a few special cases. Standard

techniques often assume that $p(\mathbf{x}, \mathbf{y})$ is jointly Gaussian, the output spaces are very low-D, or the channels are invertible [4]. Other methods suggest alternative objective functions (e.g. approximations based on the *Fisher Information* [5]), which, however, do not retain proper bounds on $I(\mathbf{x}, \mathbf{y})$. Here we analyze the relation between a simple variational lower bound on the mutual information [2] and standard approaches to approximate information maximization, focusing specifically on a comparison with the Fisher Information criterion.

1.1 Variational Lower Bound on Mutual Information

A simple lower bound on the mutual information $I(\mathbf{x}, \mathbf{y})$ follows from non-negativity of the Kullback-Leibler divergence $KL(p(\mathbf{x}|\mathbf{y})||q(\mathbf{x}|\mathbf{y}))$ between the exact posterior $p(\mathbf{x}|\mathbf{y})$ and its variational approximation $q(\mathbf{x}|\mathbf{y})$, leading to

$$I(\mathbf{x}, \mathbf{y}) \geq \tilde{I}(\mathbf{x}, \mathbf{y}) \stackrel{\text{def}}{=} H(\mathbf{x}) + \langle \log q(\mathbf{x}|\mathbf{y}) \rangle_{p(\mathbf{x}, \mathbf{y})}. \quad (2)$$

Here $q(\mathbf{x}|\mathbf{y})$ is an arbitrary distribution saturating the bound for $q(\mathbf{x}|\mathbf{y}) \equiv p(\mathbf{x}|\mathbf{y})$. The objective (2) explicitly includes¹ both the encoder $p(\mathbf{y}|\mathbf{x})$ (distribution of neural spikes for a given stimulus) and decoder $q(\mathbf{x}|\mathbf{y})$ (reconstruction of the stimulus from a population of neural firings). The flexibility of the choice of the decoder $q(\mathbf{x}|\mathbf{y})$ makes (2) particularly computationally convenient.

2 Variational Learning of Population Codes

To learn optimal stochastic representations of the continuous training patterns $\mathbf{x}_1, \dots, \mathbf{x}_M$ according to (2), we need to choose a continuous density function for the decoder $q(\mathbf{x}|\mathbf{y})$. Computationally, it is convenient to assume that the decoder is given by the isotropic Gaussian $q(\mathbf{x}|\mathbf{y}) \sim \mathcal{N}(\mathbf{U}\mathbf{y}, \sigma^2\mathbf{I})$, where $\mathbf{U} \in \mathbb{R}^{|\mathbf{x}| \times |\mathbf{y}|}$. For simplicity, we limit the discussion to this case only (though other, e.g. correlated or nonlinear cases may also be considered). Then for the empirical distribution $p(\mathbf{x}) = \sum_{m=1}^M \delta(\mathbf{x} - \mathbf{x}_m)/M$ we may express the bound (2) as a function of the encoder $p(\mathbf{y}|\mathbf{x})$ alone

$$\tilde{I}(\mathbf{x}, \mathbf{y}) \propto \text{tr} \{ \langle \mathbf{y}\mathbf{y}^T \rangle \langle \mathbf{y}\mathbf{y}^T \rangle^{-1} \langle \mathbf{y}\mathbf{x}^T \rangle \} + \text{const}. \quad (3)$$

Note that the objective (3) is a proper bound for any choice of the stochastic mapping $p(\mathbf{y}|\mathbf{x})$. We may therefore² use it for optimizing a variety of channels with continuous source vectors.

¹ The bound (2) corresponds to the criteria optimized by Blahut-Arimoto algorithms (e.g. [6]); however, we optimize it for both encoder and decoder subject to enforced tractability constraints.

² From (3) it is clear that if $\langle \mathbf{y}\mathbf{y}^T \rangle$ is near-singular, the varying part of the objective $\tilde{I}(\mathbf{x}, \mathbf{y})$ may be infinitely large. However, if the mapping $\mathbf{x} \mapsto \mathbf{y}$ is probabilistic and the number of training stimuli M exceeds the dimensionality of the neural codes $|\mathbf{y}|$, the optimized criterion is typically positive and finite.

2.1 Sigmoidal Activations

Here we consider the case of high-dimensional continuous patterns $\mathbf{x} \in \mathbb{R}^{|\mathbf{x}|}$ represented by stochastic firings of the post-synaptic neurons $y \in \{-1, +1\}^{|\mathbf{y}|}$. For conditionally independent activations, we obtain

$$p(\mathbf{y}|\mathbf{x}) = \prod_{i=1, \dots, |\mathbf{y}|} p(y_i|\mathbf{x}) \stackrel{\text{def}}{=} \prod_{i=1, \dots, |\mathbf{y}|} \sigma(y_i(\mathbf{w}_i^T \mathbf{x} + b_i)) \quad (4)$$

where $\mathbf{w}_i \in \mathbb{R}^{|\mathbf{x}|}$ is a vector of the synaptic weights for neuron y_i , b_i is its threshold, and $\sigma(a) \stackrel{\text{def}}{=} 1/(1 + e^{-a})$. Optimization of (3) for $\mathbf{W} \stackrel{\text{def}}{=} \{\mathbf{w}_1, \dots, \mathbf{w}_{|\mathbf{y}|}\} \in \mathbb{R}^{|\mathbf{x}| \times |\mathbf{y}|}$ readily gives

$$\Delta \mathbf{W} \propto \sum_{m=1, \dots, M} \text{cov}(y|\mathbf{x}_m) \left(\tilde{\mathbf{D}} \boldsymbol{\lambda}_{x_m} + \boldsymbol{\Sigma}_{yy}^{-1} \boldsymbol{\Sigma}_{yx} (\mathbf{x}_m - \boldsymbol{\Sigma}_{xy} \boldsymbol{\Sigma}_{yy}^{-1} \boldsymbol{\lambda}_{x_m}) \right) \mathbf{x}_m^T, \quad (5)$$

where $\boldsymbol{\Sigma}_{yy} \stackrel{\text{def}}{=} \langle \mathbf{y} \mathbf{y}^T \rangle$, $\boldsymbol{\Sigma}_{yx} \equiv \boldsymbol{\Sigma}_{xy}^T \stackrel{\text{def}}{=} \langle \mathbf{y} \mathbf{x}^T \rangle$ are the second-order moments, $\tilde{\mathbf{D}}$ corresponds to the diagonal of $\boldsymbol{\Sigma}_{yy}^{-1} \boldsymbol{\Sigma}_{yx} (\boldsymbol{\Sigma}_{yy}^{-1} \boldsymbol{\Sigma}_{yx})^T$, and $\lambda_i(\mathbf{x}) \stackrel{\text{def}}{=} \langle y_i \rangle_{p(y_i|\mathbf{x})} = 2\sigma(\mathbf{w}_i^T \mathbf{x} + b_i) - 1$ is the expected conditional firing of y_i . The update for the threshold $\Delta \mathbf{b}$ has the same form as (5) without the post-multiplication of each term by the training stimulus \mathbf{x}_m^T .

From (5) it is clear that the magnitude of each weight update $\Delta \mathbf{w}_i \in \mathbb{R}^{|\mathbf{x}|}$ decreases with a decrease in the corresponding conditional variance $\text{var}(y_i|\mathbf{x}_m)$. Effectively, this corresponds to a variable learning rate – as training continues and magnitudes of the synaptic weights increase, the firings become more deterministic, and learning slows down. One may also obtain a stochastic rule

$$\Delta \mathbf{W} \propto \tilde{\mathbf{D}} \langle \boldsymbol{\lambda}_x \mathbf{x}^T \rangle + \boldsymbol{\Sigma}_{yy}^{-1} \langle \boldsymbol{\lambda}_x \mathbf{x}^T \rangle \left(\boldsymbol{\Sigma}_{xx} - \langle \mathbf{x} \boldsymbol{\lambda}_x^T \rangle \boldsymbol{\Sigma}_{yy}^{-1} \langle \boldsymbol{\lambda}_x \mathbf{x}^T \rangle \right) \quad (6)$$

where $\boldsymbol{\Sigma}_{xx} \stackrel{\text{def}}{=} \langle \mathbf{x} \mathbf{x}^T \rangle$. Clearly, (6) is decomposable as a combination of the stochastic Hebbian and anti-Hebbian terms, with the weighting coefficients determined by the second-order moments of the firings and input stimuli. Additionally, from (6) one may see that the “*as-if Gaussian*” approximations [7] are suboptimal under the variational lower bound (2) – see [1] for details.

3 Fisher Information and Mutual Information

Let $\hat{\mathbf{x}} \in \mathbb{R}^{|\mathbf{x}|}$ be a statistical estimator of the input stimulus \mathbf{x} obtained from the stochastic neural firings \mathbf{y} . It is easy to see that $\mathbf{x} \rightarrow \mathbf{y} \mapsto \hat{\mathbf{x}}$ forms a Markov chain with $p(\hat{\mathbf{x}}|\mathbf{y}) \sim \delta(\hat{\mathbf{x}} - \hat{\mathbf{x}}(\mathbf{y}))$. If $\hat{\mathbf{x}}$ is *efficient*, its covariance saturates the Cramer-Rao bound (see e.g. [6]), which results in an upper bound on the entropy of the conditional distribution $H(p(\hat{\mathbf{x}}|\mathbf{x}))$. From the data processing inequality, one may obtain a lower bound on the mutual information

$$I(\mathbf{x}, \mathbf{y}) \geq H(\hat{\mathbf{x}}) + \langle \log |\mathbf{F}_{\mathbf{x}}| \rangle_{p(\mathbf{x})} / 2 + \text{const}, \quad (7)$$

where $F_x = \{F_{ij}(x)\} \stackrel{\text{def}}{=} -\langle \partial^2 \log p(y|x) / \partial x_i \partial x_j \rangle_{p(y|x)}$ is the Fisher Information matrix. Despite the fact that the mapping $y \mapsto \hat{x}$ is deterministic, exact computation of the entropy of statistical estimates $H(\hat{x})$ in the objective (7) is in general computationally intractable. It was shown that under certain assumptions $H(\hat{x}) \approx H(x)$ [5], leading to the approximation

$$I(x, y) \gtrsim \tilde{I}_F(x, y) \stackrel{\text{def}}{=} H(x) + \langle \log |F_x| \rangle_{p(x)} / 2 + \text{const}, \quad (8)$$

which is then used as an approximation of $I(x, y)$ independently of the bias of the estimator. Since $H(x)$ is independent of $p(y|x)$, maximization of (8) is equivalent to maximization of (7) where the intractable entropic term is ignored.

For sigmoidal activations (4), the criterion (8) is given by

$$\tilde{I}_F(x, y) \propto \sum_{m=1, \dots, M} \log |W^T \text{cov}(y|x_m) W| + \text{const}. \quad (9)$$

Interestingly, if $|x| = |y|$ then optimization of (9) leads to $\Delta W = 2W^{-T} - \langle \lambda_x x^T \rangle$, which (apart from the coefficient at the inverse weight – *redundancy* term) has the same form as the learning rule of [4] derived for noiseless invertible channels. Notably, the weight update has no Hebbian terms. Moreover, from (9) it is clear that as the variance of the stochastic firings decreases, the objective $\tilde{I}_F(x, y)$ may become infinitely loose. Since directions of low variation swamp the volume of the manifold, neural spikes generated by a fixed stimulus may often be inconsistent. It is also clear that optimization of $\tilde{I}_F(x, y)$ is limited to the cases when $W^T W \in \mathbb{R}^{|x| \times |x|}$ is full-rank, which complicates applicability of the method for a variety of tasks involving relatively low-D encodings of high-D stimuli.

4 Variational Lower Bound vs. Fisher Approximation

Since $\tilde{I}_F(x, y)$ is in general not a proper lower bound on the mutual information, it is difficult to analyze its tightness or compare it with the variational bound (2). To illustrate a relation between the approaches, we may consider a Gaussian decoder $q(x|y) \sim \mathcal{N}_x(\mu_y; \Sigma)$, which transforms the variational bound into

$$\tilde{I}(x, y) = -\frac{1}{2} \langle \text{tr} \{ \Sigma^{-1} (x - \mu_y)(x - \mu_y)^T \} \rangle_{p(x, y)} + \frac{1}{2} \log |\Sigma^{-1}| + \text{const}. \quad (10)$$

Here $\Sigma \in \mathbb{R}^{|x| \times |x|}$ is a function of the conditional $p(y|x)$. Clearly, if the log eigenspectrum of the inverse covariance of the decoder is constrained to satisfy

$$\sum_{i=1, \dots, |x|} \log l_i(\Sigma^{-1}) = \sum_{i=1, \dots, |x|} \langle \log l_i(F_x) \rangle_{p(x)}, \quad (11)$$

where $\{l_i(\Sigma^{-1})\}$ and $\{l_i(F_x)\}$ are eigenvalues of Σ^{-1} and F_x respectively, then the lower bound (10) reduces to the objective (8) amended with the average quadratic reconstruction error

$$\tilde{I}(x, y) = -\frac{1}{2} \underbrace{\langle \text{tr} \{ \Sigma^{-1} (x - \mu_y)(x - \mu_y)^T \} \rangle_{p(x, y)}}_{\text{reconstruction error}} + \frac{1}{2} \underbrace{\langle \log |F_x| \rangle_{p(x)}}_{\text{Fisher criterion}} + \text{const}. \quad (12)$$

Arguably, it is due to the subtraction of the non-negative squared error that (10) remains a general lower bound independently of the parameterization of the model and spectral properties of F_x . Another principal advantage of the variational approach to information maximization is the flexibility in the choice of the decoder [1].

5 Experiments

Variational Information Maximization vs Fisher criterion

In the first set of experiments we were interested to see how the value of the true MI changed as the parameters were updated by maximising the Fisher criterion $\tilde{I}_F(x, y)$ and the variational bound $\tilde{I}(x, y)$. The dimension $|y|$ was set to be small, so that the true $I(x, y)$ could be computed. Fig. 1 illustrates changes in $I(x, y)$ with iterations of the variational and Fisher-based learning rules, where the variational decoder was chosen to be an isotropic linear Gaussian with the optimal weights (3). We found that for $|x| \leq |y|$ (Fig. 1 (*left*)), both approaches tend to increase $I(x, y)$ (though the variational approach typically resulted in higher values of $I(x, y)$ after just a few iterations). For $|x| > |y|$ (Figure 1 (*right*)), optimization of the Fisher criterion was numerically unstable and lead to no visible improvements of $I(x, y)$ over its starting value at initialization.

Variational IM: stochastic representations of the digit data

Here we apply the simple linear isotropic Gaussian decoder to stochastic coding and reconstruction of visual patterns. After numerical optimization with an explicit constraint on the channel noise, we performed reconstruction of 196-dimensional continuous visual stimuli from 7 spiking neurons. The training stimuli consisted of 30 instances of digits 1, 2, and 8 (10 of each class). The source variables were reconstructed from 50 stochastic spikes at the mean of the optimal approximate decoder $q(x|y)$. Note that since $|x| > |y|$, the problem could not be efficiently addressed by optimization of the Fisher Information-based criterion (9). Clearly, the approach of [4] is not applicable either, due to its fundamental assumption of invertible mappings between the spikes and the visual stimuli. Fig. 2 illustrates a subset of the original source signals, samples of the corresponding binary responses, and reconstructions of the source data.

6 Discussion

We described a variational approach to information maximization for the case when continuous source stimuli are represented by stochastic binary responses. We showed that for this case maximization of the lower bound on the mutual information gives rise to a form of Hebbian learning, with additional factors depending on the source and channel noise. Our results indicate that other approximate methods for information maximization [7], [5] may be viewed as approximations of our approach, which, however, do not always preserve a proper




Fig. 1. Changes in the exact mutual information $I(x, y)$ for parameters of the coder $p(y|x)$ obtained by maximizing the variational lower bound and the Fisher information criterion for $M = 20$ training stimuli. *Left:* $|x| = 3, |y| = 5$ *Right:* $|x| = 7, |y| = 5$. In both cases, the global maximum is given by $I^* = \log M \approx 3.0$.




Fig. 2. *Left:* a subset of the original visual stimuli. *Middle:* 20 samples of the corresponding spikes generated by each of the 7 neurons. *Right:* Reconstructions from 50 samples of neural spikes (with soft constraints on the variances of firings).

bound on the mutual information. We do not wish here to discredit generally the use of the Fisher Criterion, since this can be relevant for bounding reconstruction error. However, for the case considered here as a method for maximising information, we believe that our method is more attractive.

References

1. Agakov, F. V. and Barber, D. (2004). Variational Information Maximization and Fisher Information. Technical report, UoE.
2. Barber, D. and Agakov, F. V. (2003). The IM Algorithm: A Variational Approach to Information Maximization. In *NIPS*.
3. Barlow, H. (1989). Unsupervised Learning. *Neural Computation*, 1:295–311.
4. Bell, A. J. and Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7(6):1129–1159.
5. Brunel, N. and Nadal, J.-P. (1998). Mutual Information, Fisher Information and Population Coding. *Neural Computation*, 10:1731–1757.
6. Cover, T. M. and Thomas, J. A. (1991). *Elements of Information Theory*. Wiley.
7. Linsker, R. (1989). An Application of the Principle of Maximum Information to Linear Systems. In *NIPS*.