A probabilistic population code based on neural sampling

Ankani Chattoraj, Shu Chen Wu, Richard D. Lange, Sabyasachi Shivkumar, Ralf M. Haefner

Brain and Cognitive Sciences, University of Rochester

III- 89

Introduction

- Visual processing is often characterized as implementing **probabilistic inference**. [1][2]
- We distinguish between **explicit** and **implicit** inference [3]:
  - **explicit:** neurons compute the posterior over variables in an internal model of the world.
  - **implicit:** probabilistically "decoding" a distribution over other variables that aren't explicitly inferred
- One candidate algorithm to do explicit probabilistic inference, is '**neural sampling**' .
  - neural responses represent **samples from the posterior probability distribution over latent variables** in the brain's internal model of the world.
- A second debate concerns whether neural responses represent **samples of latent variables** (sampling with explicit representation) [11] or **parameters of their distributions** (log probabilities for exponential families). [4][5]
- We propose that V1 spikes represent **binary samples** from a linear model of the image.

Marr's Level 1

Computational

Probabilistic Inference

Linear Image Model equation:

$$p(\mathbf{I}|\tilde{\mathbf{x}}) = \mathcal{N}\left(\sum_i \mathbf{PF}_i \mathbf{x}_i, \mathbf{1}\right)$$

$\mathbf{PF}_i \rightarrow$ Projective Field of i^{th} neuron

1. Linear Gaussian Model
2. Explicit Code
3. Binary Sampling [6][7]
4. Direct Probability

Network:**Generative** (**black**),**Inference**(**red**)

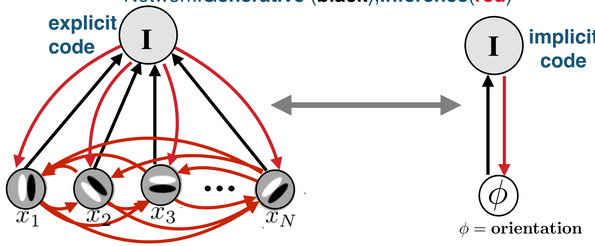

1. Orientation
2. Implicit Code
3. Parametric Code [2][5]
4. Log Probability

Constrast invariant decoding weights [9]

(Observations in support of PPC)

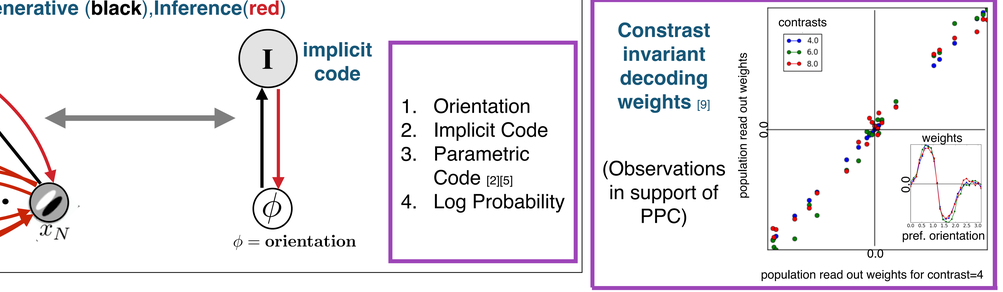

Distribution of **Noise correlation coefficients** matches empirical observation [10][15]

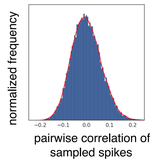

Marr's Level 2

Algorithm/Representation

Binary Sampling

Gibbs Sampling from posterior distribution:

$$p\left(x_k|\tilde{x}_{\neg k}, \mathbf{I}\right) \propto \exp\left(-\frac{1}{\sigma^2}\left|\left|\mathbf{I} - \mathbf{PF}^T\tilde{\mathbf{x}}\right|\right|^2\right) \prod_{i=1}^{N} p_{prior}^{x_i}\left(1 - p_{prior}^{(1-x_i)}\right)$$

Variational Bayesian Inference connection:

$$\mu_k \leftarrow \text{sigmoid}\left(-\frac{1}{\sigma^2}\left(\mathbf{I}^T\mathbf{PF}_k + \frac{1}{2}\mathbf{R}_{kk} + \mu_{\neg k}^T\mathbf{R}_{\neg k,k}\right) + \log\frac{p_{prior_k}}{1 - p_{prior_k}}\right)$$

$\mathbf{R} = -\mathbf{PF}^T\mathbf{PF}$

Parametric Sampling for infinite samples:

$$p(\phi|\tilde{\mathbf{x}}^{(1,2,...T)})_{T\rightarrow\infty} \propto \left|\frac{\partial \mathbf{f}(\phi)}{\partial \phi}\right| \mathcal{N}(\mathbf{f}(\phi), \mathbf{PF}^T\tilde{\mathbf{x}}, (\sigma^2)\mathbb{I})p(\phi)$$

Binary Samples interpreted as spikes [6][7]

x_1 **001111....1..**

x_2 **100001....0..**

x_3 **111000....0..**

$\vdots$ $\vdots$

x_N **000100....1..**

Orientation tuning curves scale with the contrast of grating

Close match with VB (solid line) for 50 neurons

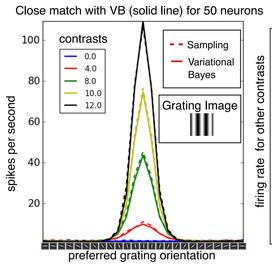

248 neuron model with PF as in [8]

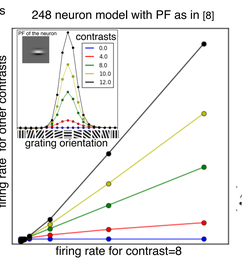

Implicit decoder gives a distribution over orientations given set of binary sample vectors, so sampling in binary hidden state of an explicit model is like sampling in parameter space of an implicit model: **Parametric Sampling**

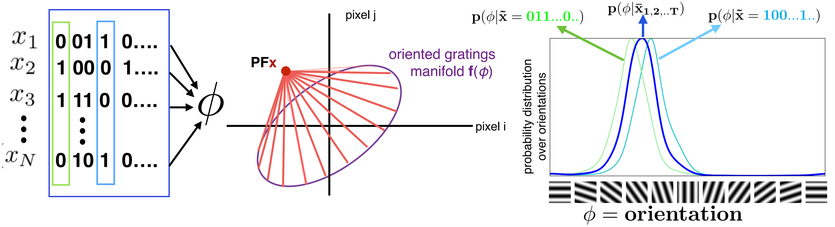

Marr's Level 3

Implementation

Leaky Integration and Fire Network

Connecting **Sampling** and **LIF**:

$$p(\text{spike now}|\mathbf{I}, \text{other recent updates})$$

$\propto$ Firing Rate (feedforward) Input Current (recurrent input from other neurons)

I Vs P in LIF neuron for fixed input:

$$p = \frac{1}{\Delta t} = \left[\tau \ln\left(\frac{I_{input\ current}}{I_{input\ current} + (V_{rest} - V_{threshold})}\right)\right]^{-1}$$

Input Current Vs Probability in LIF neuron for steady input

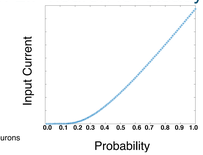

LIF equation(unitless):

$$\tau\frac{dV}{dt} = V - V_{rest} + I_{input\ current}\ ;\quad I_{input\ current} = I_{recurrent} + I_{feedforward}$$

Simulation details:

LIF $\tau = 10ms$ $V_{threshold} = -55mV$

Sampling time $= 5ms$ $V_{rest} = -70mV$

Raster Plot of spikes generated by network of 248 LIF [8] neurons

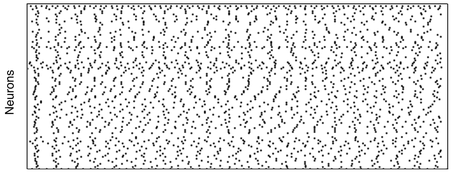

Agreement with sampling: BRIAN simulation [13]

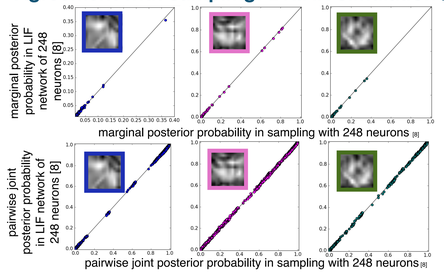

Spike Statistics of network of LIF [8] neurons

More regular but hierarchy and external/internal noise will increase irregularity

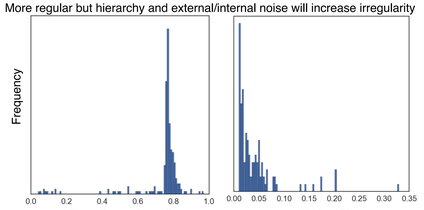

Conclusions

- The spike rate in sampling code is proportional to the marginal posterior probability over the variable represented by the neuron
- Spike rates in our model can be interpreted as parameters of an exponential family (linear PPC)
- Firing rates are proportional to probability over Gabor features (explicit) but log probability over orientation (implicit)
- Neural responses in this model show contrast-invariant tuning.
- LIF network can implement sampling in binary linear-Gaussian image model

References

[1] Fiser, József, et al. "Statistically optimal perception and learning: from behavior to neural representations." Trends in cognitive sciences 14.3 (2010)
[2] Pouget, Alexandre, et al. "Probabilistic brains: knowns and unknowns." Nature neuroscience 16.9 (2013)
[3] Pitkow, Xaq, and Dora E. Angelaki. "Inference in the brain: Statistics flowing in redundant population codes." Neuron 94.5 (2017)
[4] Hoyer, Patrik O., and Aapo Hyvärinen. "Interpreting neural response variability as Monte Carlo sampling of the posterior." Advances in neural information processing systems. 2003.
[5] Ma, Wei Ji, et al. "Bayesian inference with probabilistic population codes." Nature neuroscience 9.11 (2006)
[6] Bornschein, Jörg, Marc Henniges, and Jörg Lücke. "Are V1 simple cells optimized for visual occlusions? A comparative study." PLoS computational biology 9.6 (2013)
[7] Buesing, Lars, et al. "Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons." PLoS computational biology 7.11 (2011)
[8] Orbán, Gergő, et al. "Neural variability and sampling-based probabilistic representations in the visual cortex." Neuron 92.2 (2016)
[9] Beck, Jeffrey M., et al. "Probabilistic population codes for Bayesian decision making." Neuron 60.6 (2008)
[10] Ecker, Alexander S., et al. "Decorrelated neuronal firing in cortical microcircuits." science 327.5965 (2010)
[11] Haefner, Ralf M., Pietro Berkes, and József Fiser. "Perceptual decision-making as probabilistic inference by neural sampling." Neuron 90.3 (2016)
[12] Petrovici, Mihai A., et al. "Stochastic inference with spiking neurons in the high-conductance state." Physical Review E 94.4 (2016)
[13] Goodman, Dan FM, and Romain Brette. "The brian simulator." Frontiers in neuroscience 3 (2009): 26.
[14] David Marr, Vision. A Computational Investigation into the Human Representation and Processing of Visual Information, , 1982
[15] Ecker, Alexander S., et al. "State dependence of noise correlations in macaque primary visual cortex." Neuron 82.1 (2014): 235-248.