MCMC with spiking neurons

Lloyd Elliott[1], Chris Eliasmith[2]

[1]Gatsby Computational Neuroscience Unit, University College London,
[2]Center for Theoretical Neurosceince, Univeristy of Waterloo

Abstract

Diverse problems in computational neuroscience are solved by Bayesian models. Human sensorimotor control (Todorov 2004; Körding&Wolpert 2004), and some aspects of perception and decision making have all been shown to be nearly Bayes optimal under appropriately specified loss functions. Bayesian algorithms in machine learning offer ways in which populations of neurons could implement a system that provides empirically observed Bayesian optimal behaviour. In this work, we approximate arbitrary distributions with population codes. These codes are necessarily limited by the dependency structures, saturation and noise inherent in neural signals. We explore the implications for neural implementations of MCMC methods.

We simulate populations of leaky integrate-and-fire neurons and we construct population codes to approximate functions with a given distribution. We examine the integrated autocorrelation time of the decoded signal and we report the KS-test statistics and quantile/quantile plots.

- A prerequisite for neural implementation of MCMC is specification of a neural signal that approximately follows a given distribution *f*.
- Assume linear dendrites, instantaneous axons, and leaky integrate-and-fire somatic dynamics.

$$\tau_{RC}\frac{dV^{(i)}}{dt} = (E_L - V^{(i)}) + I^{(i)}_{\text{ext}}(t)$$

- Assume that the neural signal is voltage potential on the soma of a neuron that is fully connected to a population.

$$x(t) = \sum_{ij} \omega_i h_i(t - T_{ij}) \qquad h(t) \propto a_1 e^{-t/\tau_1} - a_2 e^{-t/\tau_2}$$

- Assume that current injected into a member of the population is an affine transformation of a network signal *g* (Eliasmith&Anderson 2003).

$$I^{(i)}_{\text{ext}} = a^{(i)} s(t) + b^{(i)}$$

- Assume that the distribution of *g* is known.
- Find synaptic weights such that: $x(t_1) \sim f$ holds approximately,

Methods

- Assume *f* is uniform distribution. This distribution has maximal entropy in the set of all distributions with a fixed support.
- Use ridge regression to find synaptic weights that minimise the following error:

$$E = \int_0^\infty \left(\phi(s(t)) - \sum_{ij} \omega_i h(t - T_{ij})\right)^2 \mathrm{d}t$$

Solution 1

- Assume $s(t) \sim g$ is Gaussian band limited noise with cutoff frequency *L*.
- Choose

$$\phi(s(t)) = F^{-1}(G(s(t)))$$

- *$F^{-1}(G(s(T))$* is marginally distributed as *f*. This is known as inverse sample transformation (Devroye 1986).

Solution 2

- Assume $\theta \sim g$ is a band limited bimodal mixture of Gaussians with cutoff frequency *L*.
- Choose

$$s(\theta) = (\sin\theta, \cos\theta)$$

- Choose $\phi(s(\theta)) = F^{-1}(G(\arctan(sin(\theta)/cos(\theta)))$

- This is marginally distributed as *f*.

Observations

a)

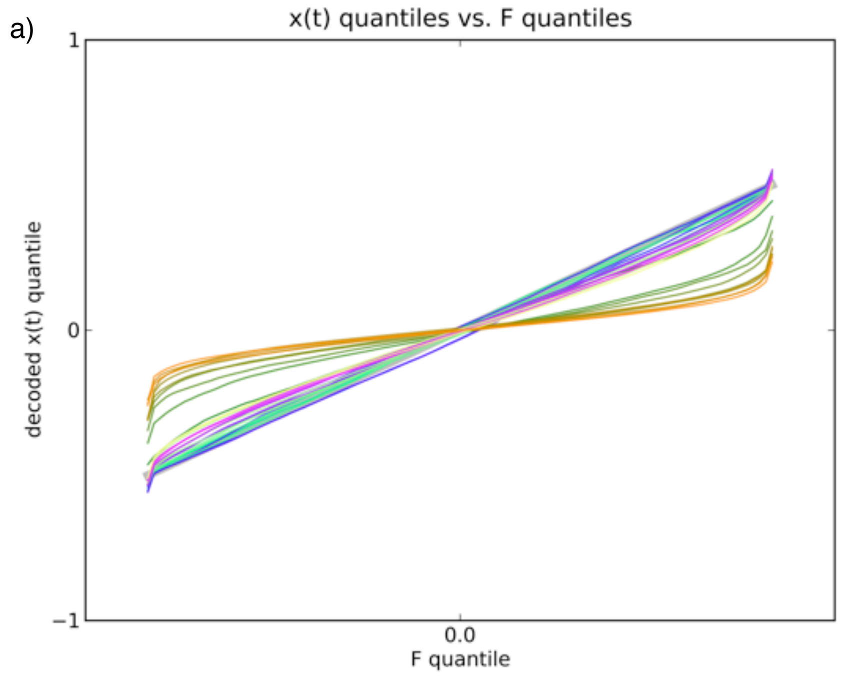

c)

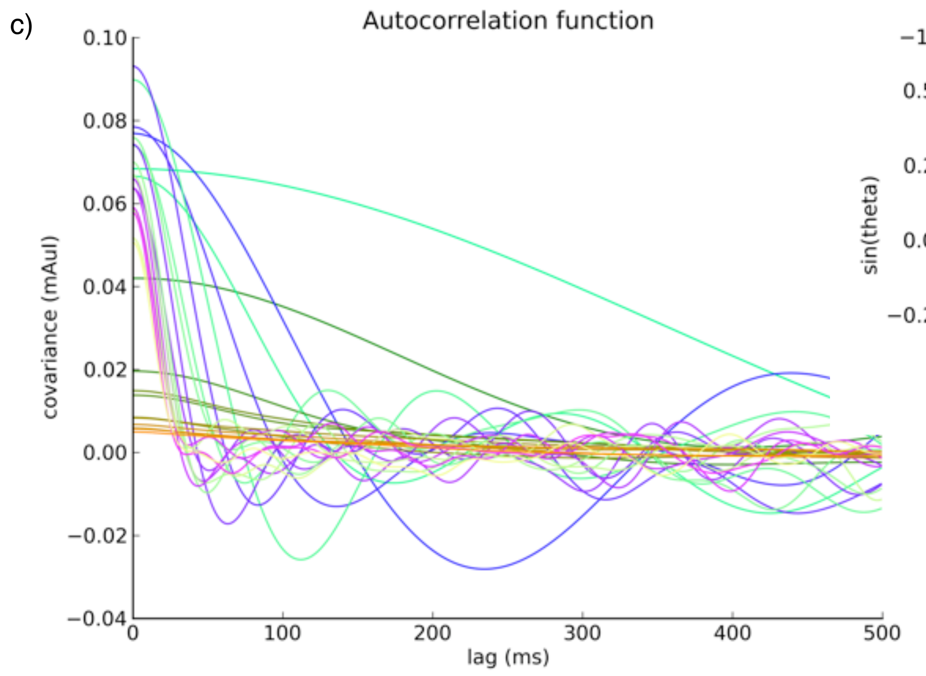

b)

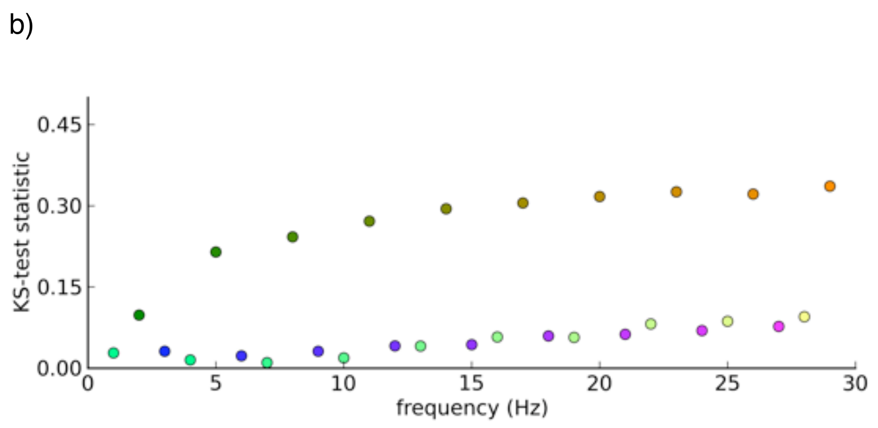

d)

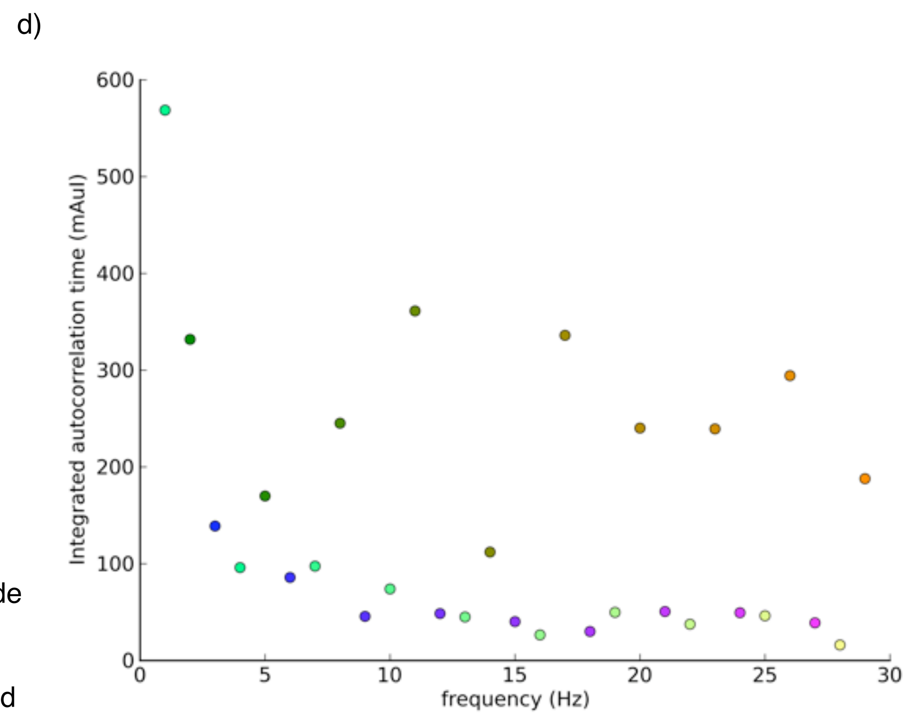

Solution 1:
We compare 3 synaptic filter types, and vary *L*. Trials span *T=100s* and network size is 500 neurons. a) Quantile/quantile plot of decoded signal against uniform distribution. NMDA synapses provide a poor approximation. Distortion accumulates at tails of distribution. b) KS-test statistic of decoded signals shows some trends across frequency and synaptic filter. c) Autocorrelation functions of decoded signals. Characteristics of synaptic filters are visible. d) Integrated autocorrelation time of decoded signal. Trend across frequency is predicted by network wide correlation.

e)

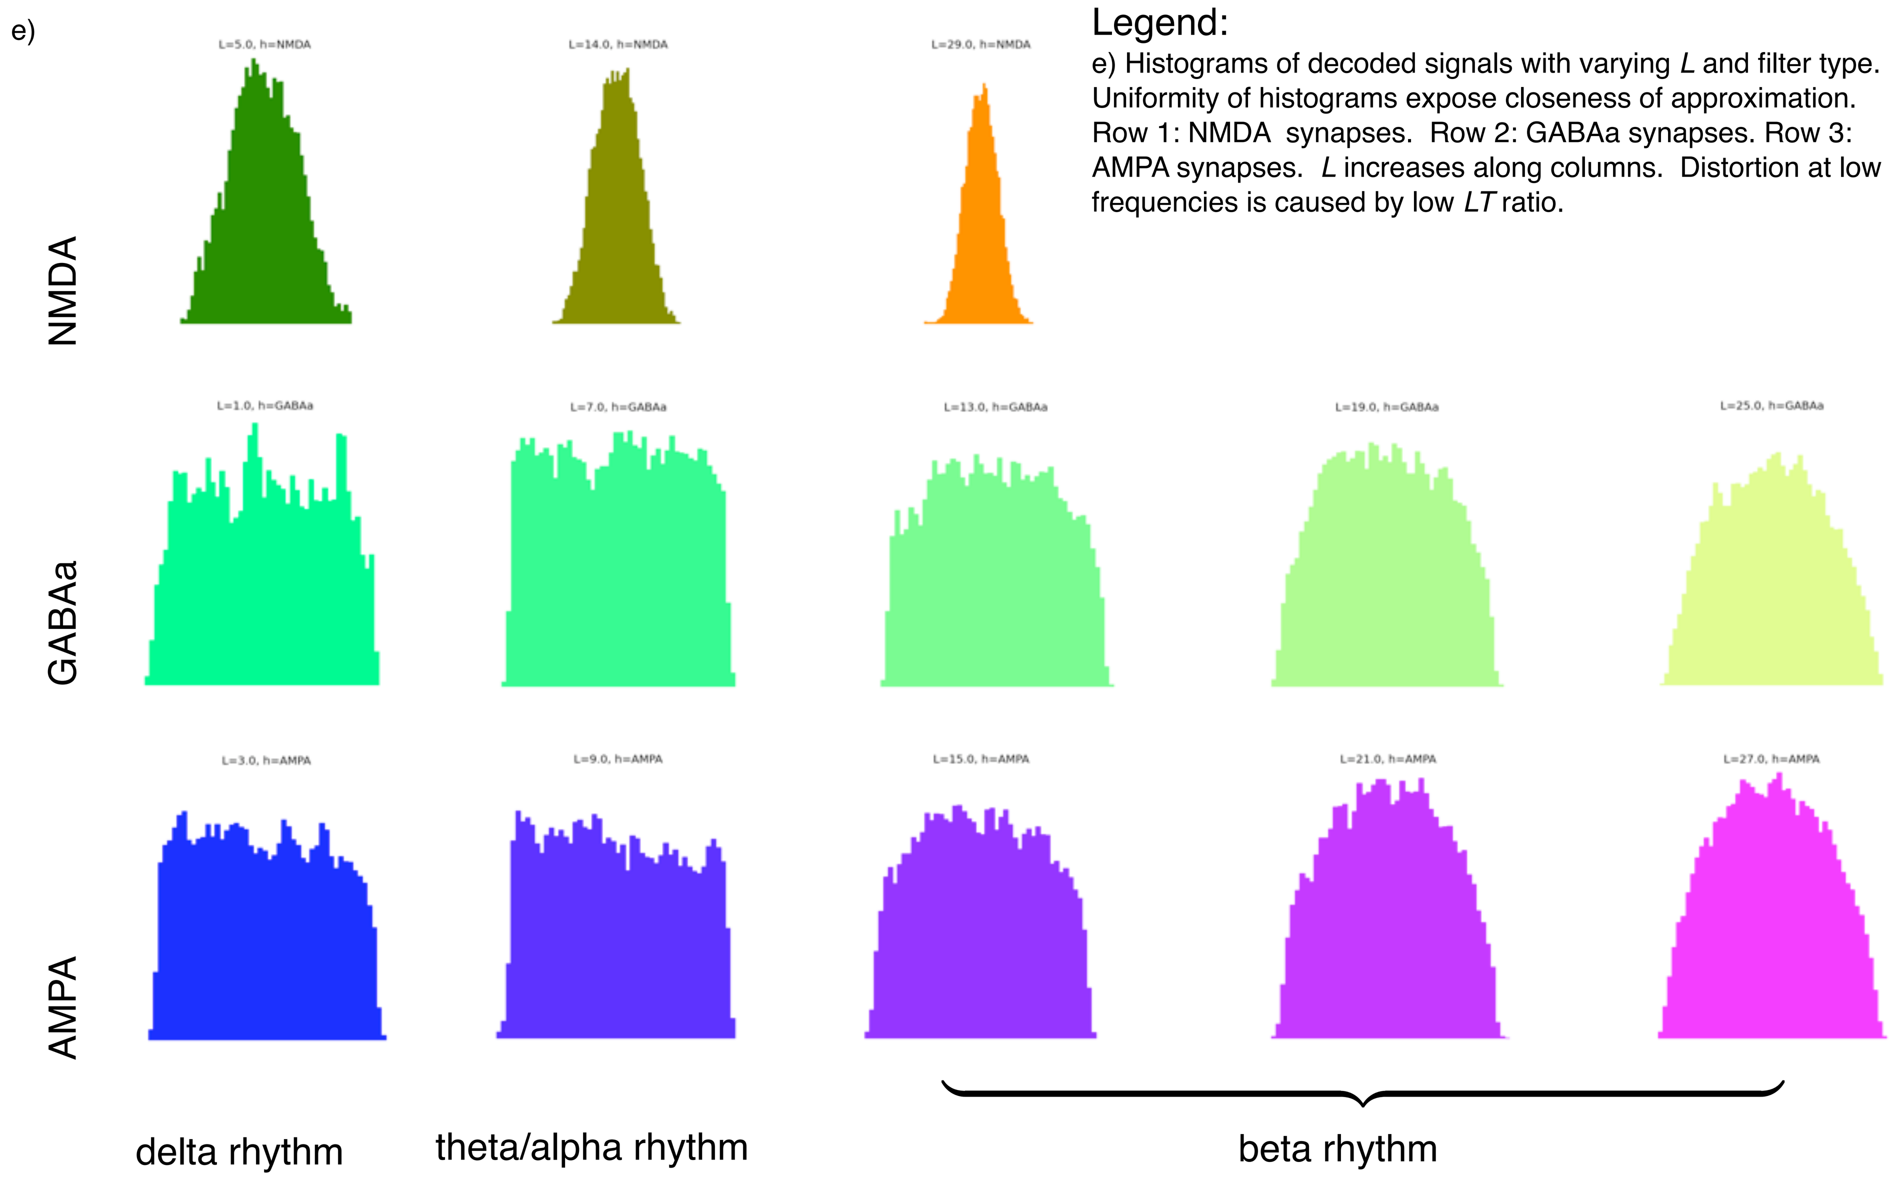

Legend:
e) Histograms of decoded signals with varying *L* and filter type. Uniformity of histograms expose closeness of approximation. Row 1: NMDA synapses. Row 2: GABAa synapses. Row 3: AMPA synapses. *L* increases along columns. Distortion at low frequencies is caused by low *LT* ratio.

a)

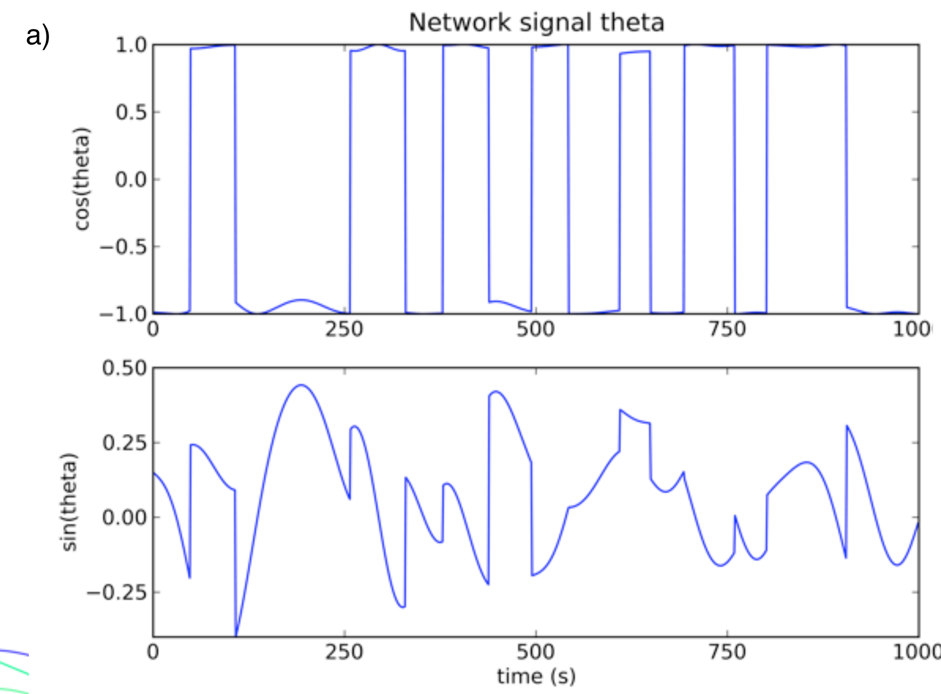

b)

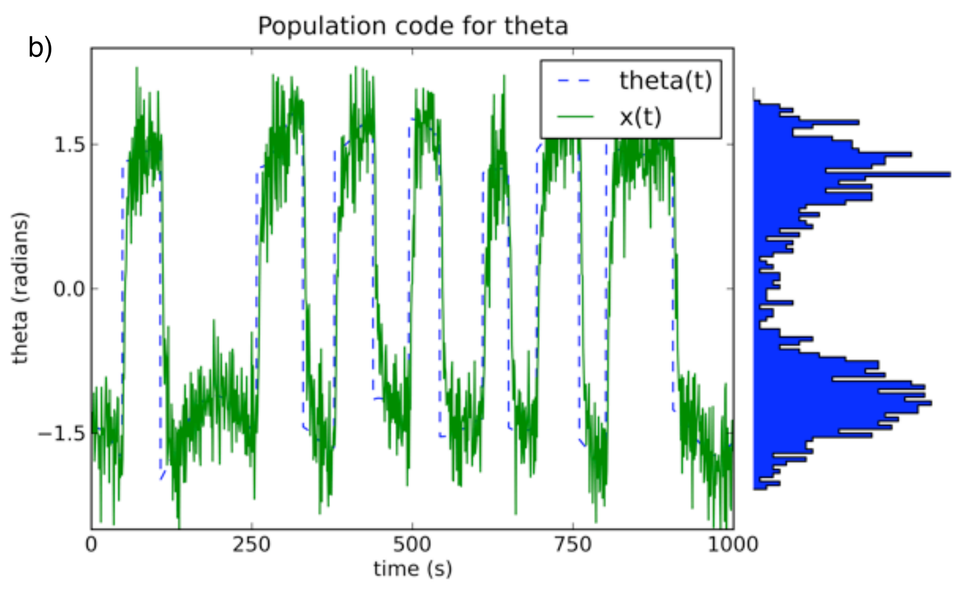

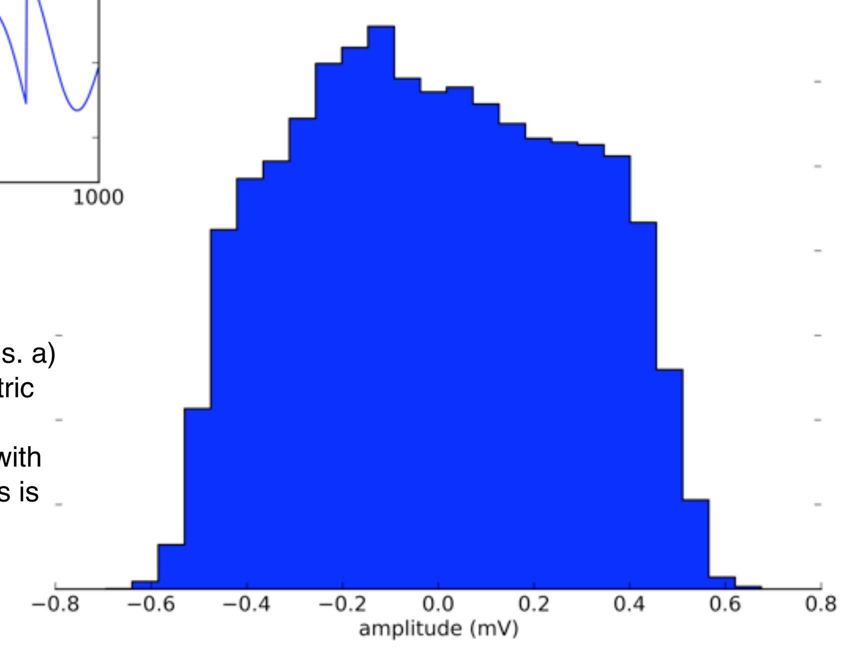

Solution 2:
Simulation spans T=20s. Network size is 500 neurons. a) the network signal is projected onto circle. b) Symmetric tuning curves allow accurate decoding of multimodal distributions. c) Histogram of decoded signal filtered with synaptic dynamics with time constant tau=0.005. This is close to uniform distribution.

Results

- qq plot divergence is dominated by tails.
- The best approximation is obtained in the delta, theta, alpha rhythm(*L* ~ 1Hz to 12Hz).
- NMDA receptors produce poor approximations of uniform distributions.
- Injection of network signal into two dimensional space allows transformation of bimodal distributions.

Discussion

We have shown a relationship between integrated autocorrelation time in marginally uniform neural signals and neural oscillation. Integrated autocorrelation time yields the variance of naïve MCMC:

$$\text{Var}\left[\int_0^T x(t)\mathrm{d}t\right] \to 1 + 2\int_0^\infty C_t/C_0 \mathrm{d}t$$

From this equation, it is possible to derive the variance of a neural implementation of an MCMC method that uses population codes and inverse sample transformations as a source of random variables. We expect to see delta, theta or alpha regime oscillation and GABAa or AMPA receptors in areas supporting such implementation. Faster oscillations lose entropy after synaptic dynamics.

Future Work

- Implement more sophisticated MCMC methods.
- Derive experimental predictions.
- Increase dimensionality of sample space.

Acknowledgements

This research was funded by NSERC grant NSERC - PGSM 374370-09. Simulations were conducted in `nengo` (http://nengo.ca) on `SHARCNET`. We would like to thank Maneesh Sahani for helpful discussion.

References

Devroye, L., 1986, 'Non-uniform random variate generation,' Springer-Verlag.
Todorov, E., 2004, 'Optimality principles in sensorimotor control,' Nat. neurosci. (9), 907-915.
Körding, K.P., Wolpert, D.M., 2004, 'The loss function of sensorimotor learning', Pnas 101(26), 9839
Eliasmith, C., Anderson, C.H., 2003, 'Neural engineering,' MIT press.