
Uncertainty, phase and oscillatory hippocampal recall supplementary material

Máté Lengyel and Peter Dayan
 Gatsby Computational Neuroscience Unit
 University College London
 17 Queen Square, London WC1N 3AR, United Kingdom
 {lmate, dayan}@gatsby.ucl.ac.uk

A Network dynamics for memory recall

Taking the posterior distribution over patterns given the input to the network and the synaptic weights between neurons (equation 8 of the main paper) to be the distribution that a network tries to approximate (equation 4 and 5 of the main paper) yields the following update rules for neuronal firing phases, concentrations, and rates:

$$\tau \frac{d\phi_i}{dt} = \lambda(c_i, r_i) \left[F_\phi(\phi_i, \tilde{x}_i, \tilde{z}_i) + \frac{1}{\sigma_W^2} \sum_{j \neq i} \lambda(c_j, r_j) H_\phi(\phi_i, \phi_j) \right] \quad (1a)$$

$$\tau \frac{dc_i}{dt} = \gamma'(c_i) \rho(r_i) \left[F_c(\phi_i, \tilde{x}_i, \tilde{z}_i) + \frac{1}{\sigma_W^2} \sum_{j \neq i} \lambda(c_j, r_j) H_c(\phi_i, \phi_j) - L_c(r_i, c_i) \right] \quad (1b)$$

$$\tau \frac{dr_i}{dt} = \rho'(r_i) \left[\gamma(c_i) \left[F_c(\phi_i, \tilde{x}_i, \tilde{z}_i) + \frac{1}{\sigma_W^2} \sum_{j \neq i} \lambda(c_j, r_j) H_c(\phi_i, \phi_j) \right] + F_r(\tilde{x}_i, \tilde{z}_i) + \frac{1}{\sigma_W^2} \sum_{j \neq i} \rho(r_j) H_r - L_r(r_i, c_i) \right] \quad (1c)$$

where

$$F_\phi(\phi_i, \tilde{x}_i, \tilde{z}_i) = \frac{\partial}{\partial \phi_i} \log P(\phi_i) + \tilde{z}_i \frac{\partial}{\partial \phi_i} \log \tilde{P}_1(\tilde{x}_i | \phi_i) \quad (2a)$$

$$F_c(\phi_i, \tilde{x}_i, \tilde{z}_i) = \log P(\phi_i) + \tilde{z}_i \left[\log \tilde{P}_1(\tilde{x}_i | \phi_i) - \tilde{\beta}(\tilde{x}) \right] \quad (2b)$$

$$F_r(\tilde{x}_i, \tilde{z}_i) = \tilde{z}_i \left[\log \left(\frac{\eta_1 \eta_0}{(1 - \eta_1)(1 - \eta_0)} \right) - \log \tilde{P}_0(\tilde{x}_i) + \tilde{\beta}(\tilde{x}) \right] \quad (2c)$$

describe the effect of the firing phase of the neuron (in conjunction with the input) on its firing phase, concentration, and rate,

$$H_\phi(\phi_i, \phi_j) = (W_{ij} - \mu_W - \Omega(\phi_i, \phi_j)) \frac{\partial}{\partial \phi_i} \Omega(\phi_i, \phi_j) \quad (3a)$$

$$H_c(\phi_i, \phi_j) = (W_{ij} - \mu_W) (\Omega(\phi_i, \phi_j) - \beta_\Omega) - \frac{1}{2} (\Omega^2(\phi_i, \phi_j) - \beta_{\Omega^2}) \quad (3b)$$

$$H_r = \frac{1}{2} [(W_{ij} + W_{ji} - 2\mu_W) \beta_\Omega - \beta_{\Omega^2}] = \begin{cases} (W_{ij} - \mu_W) \beta_\Omega - \frac{1}{2} \beta_{\Omega^2} & \text{for symmetric } \Omega \\ -\frac{1}{2} \beta_{\Omega^2} & \text{for antisymmetric } \Omega \end{cases}$$

are the firing phase-based effects of presynaptic neuron j on postsynaptic neuron i 's firing phase, concentration, and rate, and

$$L_c(r_i, c_i) = \left(1 - \frac{\epsilon}{T}\right) \log \left(1 + \frac{T}{\epsilon} \frac{\gamma(c_i)}{1 - \gamma(c_i)}\right) + \beta \quad (4a)$$

$$L_r(r_i, c_i) = \log \frac{\rho(r_i)}{1 - \rho(r_i)} \frac{1 - p_z}{p_z} \frac{\eta_0}{1 - \eta_1} + \left(\gamma(c_i) + \frac{\epsilon}{T} (1 - \gamma(c_i))\right) \log \left(1 + \frac{T}{\epsilon} \frac{\gamma(c_i)}{1 - \gamma(c_i)}\right) + \log(1 - \gamma(c_i)) - \beta \cdot (1 - \gamma(c_i)) - \log T \quad (4b)$$

is a change in firing concentration and rate that is only a function of the firing concentration and rate of the neuron itself, and $\beta = \overline{\log P}(\star)$, $\tilde{\beta}(\tilde{x}_i) = \overline{\log \tilde{P}_1}(\tilde{x}_i | \star)$, $\beta_\Omega = \overline{\Omega}(\star, \star)$ ($\beta_\Omega = 0$ for antisymmetric Ω), and $\beta_{\Omega^2} = \overline{\Omega^2}(\star, \star)$.

B Interpretation of the dynamics, and experimental predictions

We interpret the dynamics given by Eqs. 1a-4b both in functional and biological terms.

1. Functionally:

- (a) Bar the effects of burst strength, the firing phase dynamics (Eq. 1a) are the same as for the MAP network [1], *ie* they find the maximum of the posterior $P[x_i | z_i = 1, \tilde{\mathbf{x}}, \tilde{\mathbf{z}}, \mathbf{W}]$, as they should since firing phases represent a delta-distribution approximation of the posterior, and doing a variational approximation as we do with a delta distribution is equivalent to making a MAP approximation.
 - i. The term F_ϕ (Eq. 2a) maximizes the prior probability and the likelihood of the input.
 - ii. The term H_ϕ (Eq. 3a) maximizes the likelihood of the synaptic weights of the neuron. Intuitively, it tunes ϕ_i until the synaptic weight change that would be caused by the present pre- and postsynaptic firing phases matches the difference between the actual synaptic weight and the expected synaptic weight after storing $M - 1$ patterns.
- (b) The effect of burst strength on firing phase changes (Eq. 1a) is intuitive: the bigger the burst strength of a cell, *ie* the more certain it is about the value represented by its firing phase, the bigger its influence on its postsynaptic partners should be.
- (c) The terms F_c (Eq. 2b) and H_c (Eq. 3b) in the firing concentration dynamics (Eq. 1b) evaluate the log posterior probability of the value represented by the current firing phase of the cell given that it should fire, $P[\phi_i | z_i = 1, \tilde{\mathbf{x}}, \tilde{\mathbf{z}}, \mathbf{W}]$: if it is greater than the average of the log posterior, then the certainty and hence the firing concentration is increased.
- (d) The firing concentration of a neuron is also increased by the burst strengths of its presynaptic partners (Eq. 1b), as it should be intuitively if they collectively represent certainty about non-independent quantities.
- (e) Finally, L_c (Eq. 4a) describes a negative feed-back in firing concentration. This is a necessary consequence of the variational approximation used here: the network tries to find the distribution with maximal entropy (represented by minimal firing concentrations) that best matches the mode of the true posterior.

2. Biologically:

- (a) Since the firing phase-dependent terms in the firing phase dynamics are the same as for the MAP network, their biological plausibility is unaffected by this extension.
 - i. The prior- and input-dependent term F_ϕ (Eq. 2a) describes gradual phase locking with an external oscillation. Analytic calculations and numerical simulations of biophysical neurons show that phase locking of this sort can easily occur in response to current injection that is in-phase with the external oscillation.
 - ii. The interaction term H_ϕ (Eq. 3a) is a phase-coupling function and has been tested experimentally *in vitro* [2]. There is one slight change in the way it is presented

here compared to the MAP network. Formerly it has been argued that the term involving $-\Omega(x_i, x_j)$ is not biologically plausible and hence needs to be substituted by its average. However, this was more obviously true for firing rate-based only networks, not for firing phase-based dynamics. Note, that the effect of this term becomes important if the synaptic weight W_{ij} is small, otherwise the term involving the synaptic weight dominates. Indeed, we have some preliminary evidence that for weak stimulation (corresponding to small synaptic weight) the form of the interaction seems to be closer to that of the term involving $-\bar{\Omega}(x_i, x_j)$. Bar this term, the magnitude of the phase-interaction between two neurons is scaled by the synaptic weight between them, as would be expected biologically.

- (b) The firing phase-based interaction is also scaled by the burst strength of the presynaptic neuron (Eq. 1a), which seems plausible: the more spikes a neuron fires in a burst, and the more rapidly it fires them, the bigger postsynaptic current it entails, and the bigger the resulting postsynaptic phase shifts should be. Preliminary *in vitro* results show that PRCs recorded in response to burst stimulation are not qualitatively different from single spike-induced PRCs. It remains to be tested whether their magnitude scales appropriately.
- (c) The way the firing phases of the presynaptic partners of a neuron influence its firing concentration (H_c , Eq. 3b) and rate (Eq. 1c) depends on the specific learning rule. Interestingly, in general, it seems that the advancing / delaying effect of presynaptic firing can freely combine with its effect in increasing / decreasing postsynaptic firing rate, even in counter-intuitive ways, *eg* there can be firing phase differences when presynaptic firing delays postsynaptic firing ($H_\phi > 0$) but increases firing rate ($H_c > 0$), and *vica versa*. In principle, it should be possible to test these prediction *in vitro*.
- (d) Importantly, the firing rates of presynaptic neurons also scale their effect on the firing rate of the postsynaptic neuron (Eq. 1c), which again seems perfectly plausible, and *in vitro* testable.
- (e) Finally, the term L_r (Eq. 4b) simply describes leaky firing rate decay and thus seems biologically justified. It may also be tested *in vitro*.

References

- [1] Lengyel M, Dayan P. In Advances in Neural Information Processing Systems 17, 769, Cambridge, MA, 2005. MIT Press.
- [2] Lengyel M, et al. Nat Neurosci 8:1677, 2005.