Computing with rates vs spikes: insights from two solutions to an integrator network

Jorge A. Menéndez[1,2,3] **& Peter E. Latham**[1]

[1]Gatsby Computational Neuroscience Unit, [2]Sainsbury-Wellcome Centre, [3]CoMPLEX, University College London

How does spiking affect the dynamics of a functional network?

- We find that two different integrator networks - one spiking and one not - implement drastically different algorithms for integration
- Particularly, the error-encoding membrane potential approach [1, 2, 3] to building spiking integrator networks yields networks with idiosyncratic dynamical properties
- These properties endow such networks with several features that make it highly advantageous over a linear rate network performing the same computation [4, 5]
- However, they make strong experimental predictions that may not hold in biological brains

A spiking and non-spiking integrator network

Our goal is to build a recurrent network with readout signals $r_1(t), \ldots, r_N(t)$ from each neuron that can be linearly summed to produce an estimate $\hat{\mathbf{x}}(t)$ tracking the integral of some d-dimensional time-varying input signal $\mathbf{c}(t)$:

$$\sum_{i=1}^{N} \mathbf{D}_i r_i(t) = \mathbf{D}\mathbf{r}(t) = \hat{\mathbf{x}}(t) \approx \mathbf{x}(t) = \int_0^t \mathbf{c}(t')\mathrm{d}t'$$

How can we wire up a network to do this? We consider two previously proposed recipes:

Linear rate network [4]

- Without any input, network activity naturally decays to 0. But by tuning the recurrent weights to provide a precise level of positive feedback [5], the network can reliably control its estimate to perform integration
- Because activity outside the low-dimensional coding space doesn't change the network estimate, the space of solutions allows for arbitrary dynamics in this "nullspace".

Leaky integrate-and-fire (LIF) network [1, 3, 2]

- We want each neuron to spike whenever its contribution $\mathbf{D}_i$ to $\hat{\mathbf{x}}$ reduces the estimate error $\mathbf{e}(t) = \mathbf{x}(t) - \hat{\mathbf{x}}(t)$
- This can be achieved by having each neuron's membrane potential $V_i(t)$ track the alignment of $\mathbf{e}(t)$ with its decoding vector $\mathbf{D}_i$
- Whenever the alignment is beyond firing threshold, the neuron will then fire and its contribution to the estimate will cancel out the error

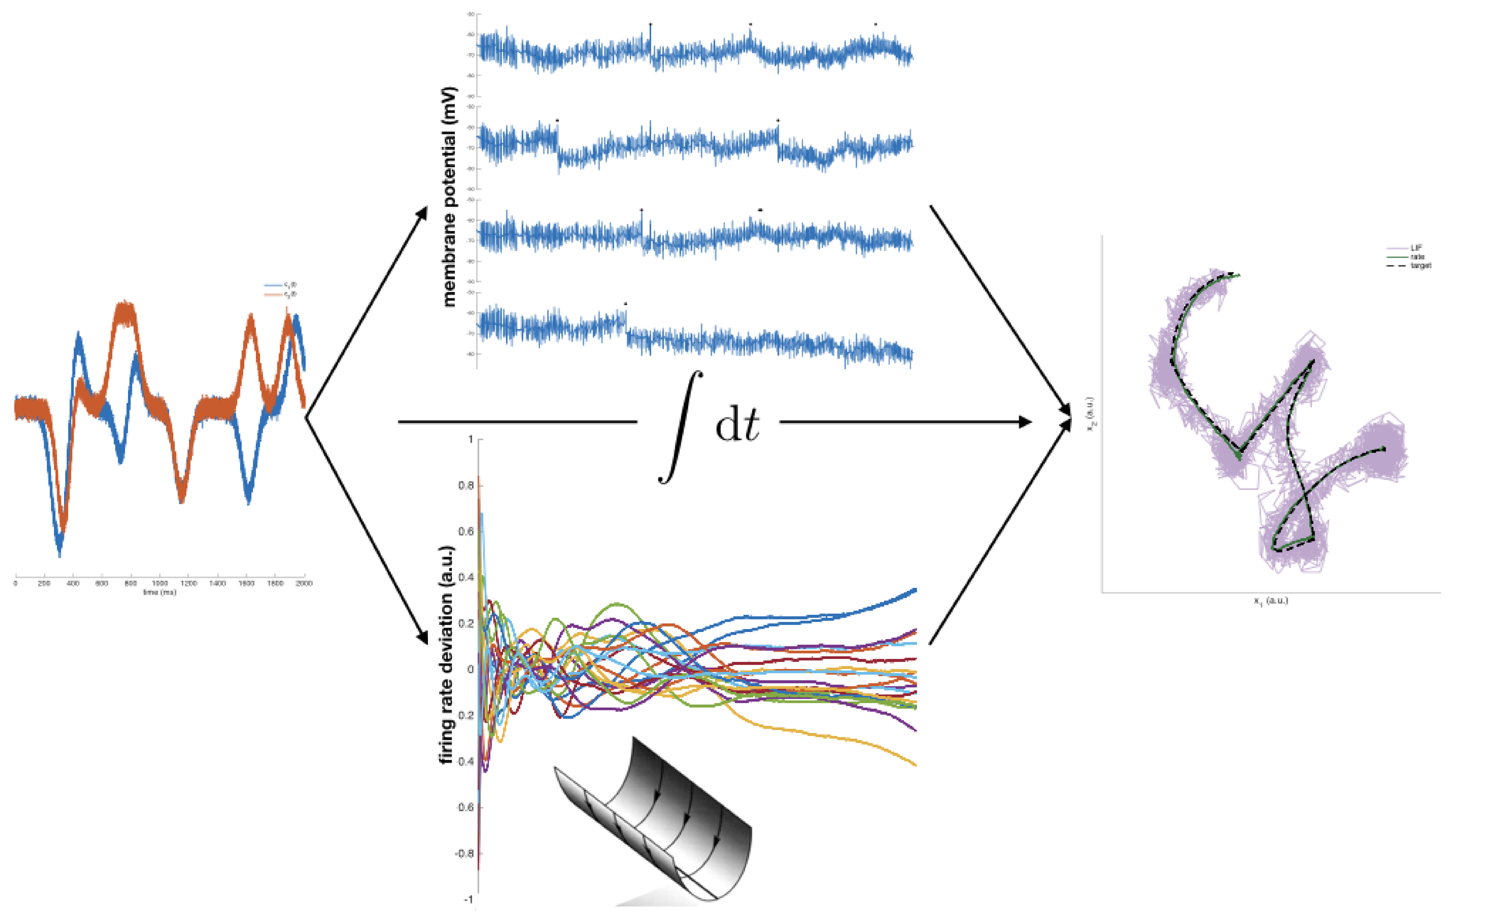

Two fundamentally different solutions to integration

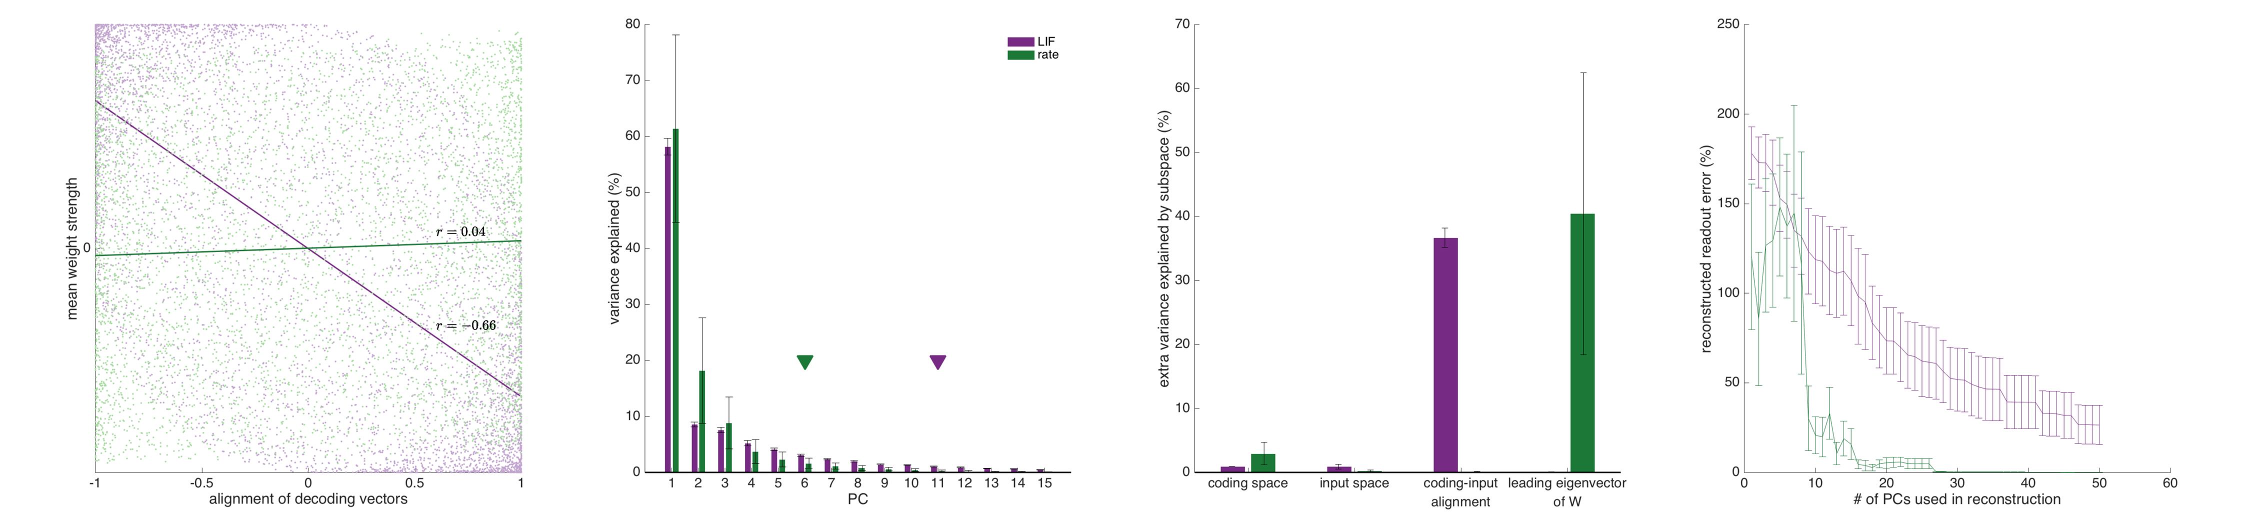

Are these properties useful?

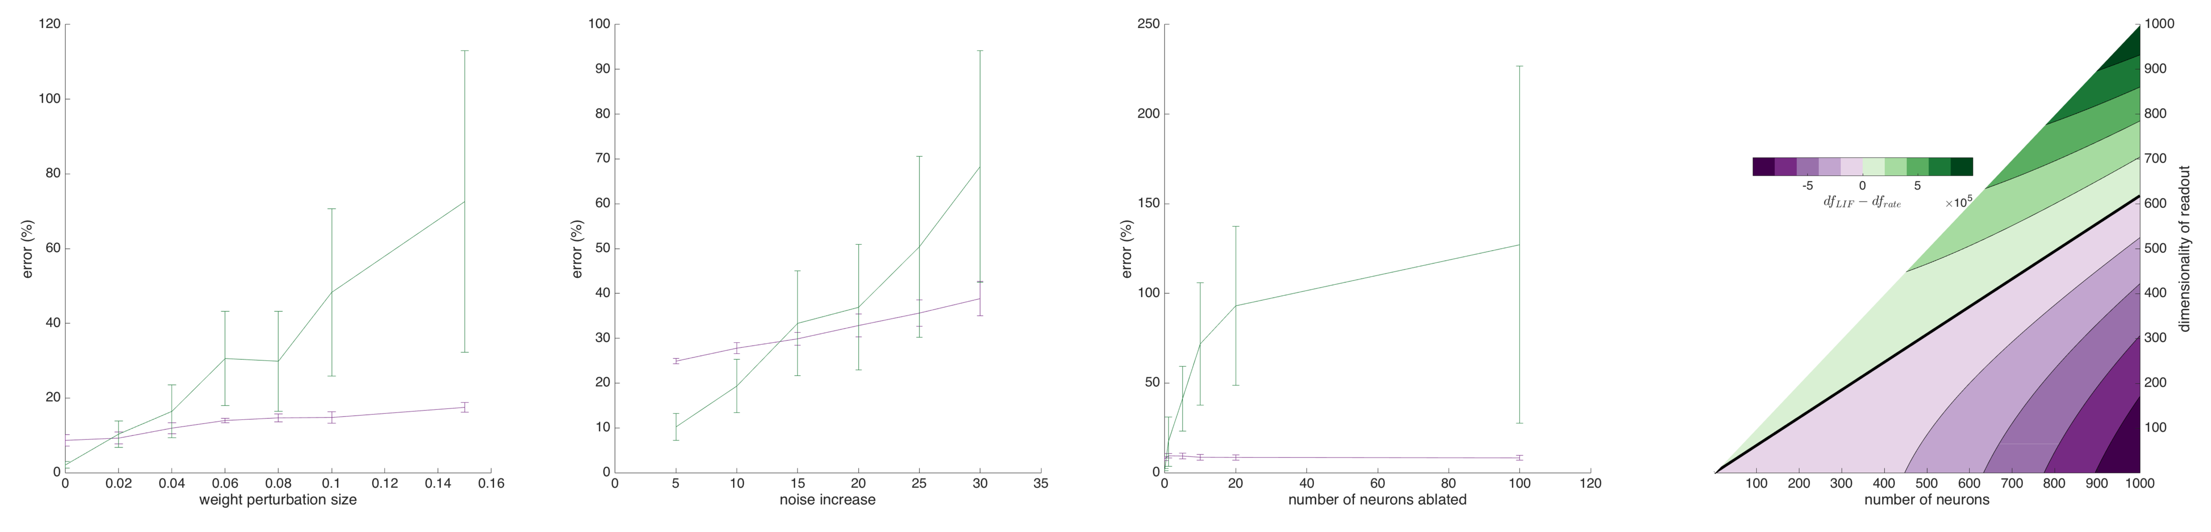

How can we distinguish between these experimentally?

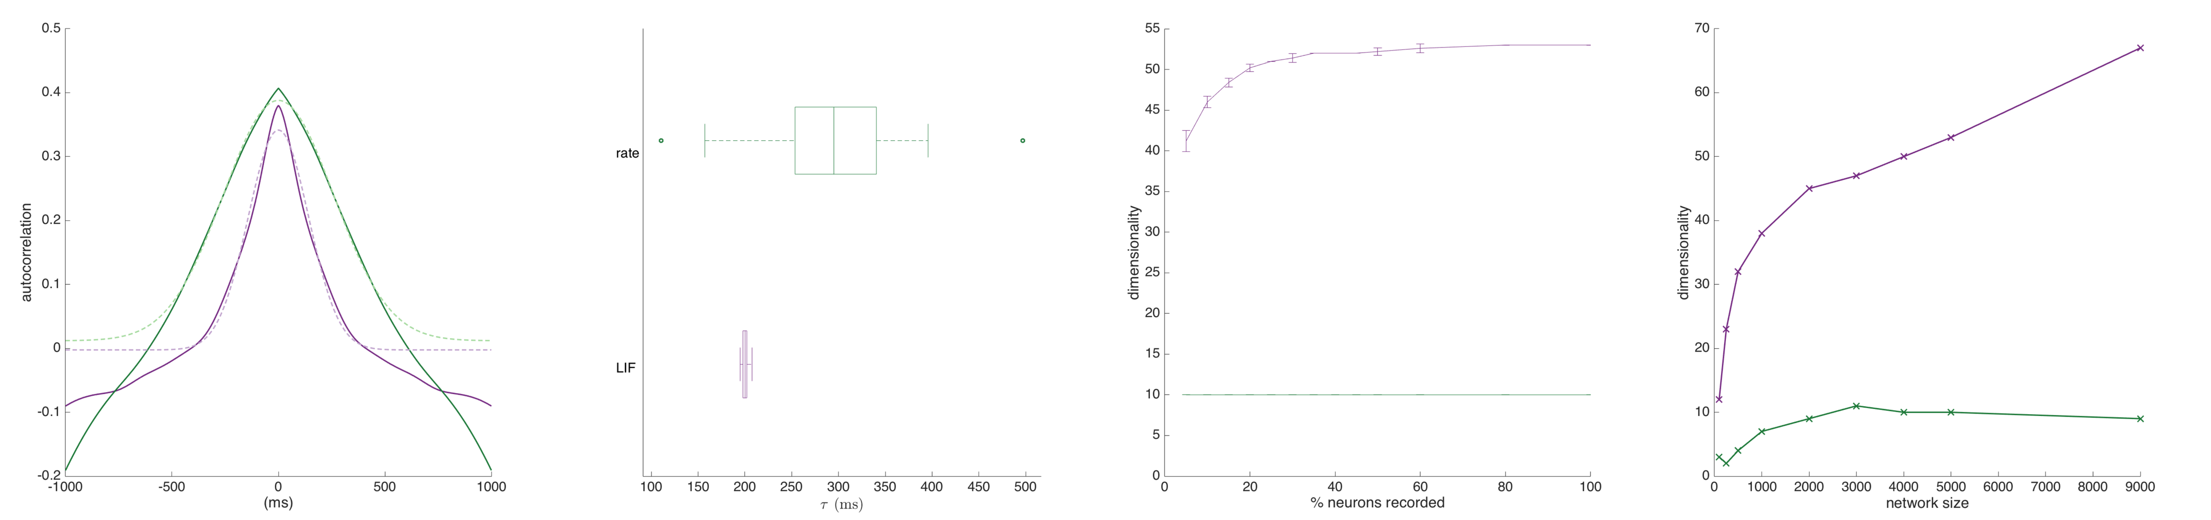

Network Equations

We want to wire up the LIF network to achieve the following:

$$V_i(t) = \mathbf{F}_i^T\mathbf{e}(t) = \mathbf{F}_i^T(\mathbf{x}(t) - \hat{\mathbf{x}}(t)) = \mathbf{F}_i^T\mathbf{x}(t) - \mathbf{F}_i^T\mathbf{D}\mathbf{r}(t)$$

where $\mathbf{F}_i\mathbf{D}_i > 0$, thus enforcing that V_i roughly encode the alignment of $\mathbf{e}(t)$ and $\mathbf{D}_i$. Interpret the terms as input weights $\mathbf{F}_i$ and recurrent weights $\mathbf{W}_i = -\mathbf{F}_i^T\mathbf{D}$, we obtain the following dynamics:

$$c_m\dot{V}_i = -g_L(V_i(t) - V_r) + \sum_{j=1}^{d} F_{ij}c_j(t) + \sum_{j=1}^{n} W_{ij}\dot{r}_j(t) + \sigma\xi(t)$$

The PSC shape is thus constrained to be the derivative of the readout signal.
In the linear rate network, we assume dynamics

$$\tau\dot{\mathbf{r}} = -\mathbf{r} + \mathbf{W}\mathbf{r} + \mathbf{F}\mathbf{c} + \sigma\xi$$

By precisely picking the synaptic weights, we can easily obtain a network whose readout is the integral of its input:

$$\mathbf{F} = \tau\mathbf{D}^T(\mathbf{D}\mathbf{D}^T)^{-1} + \mathbf{F}_\perp \qquad \mathbf{W} = \mathbf{D}^T(\mathbf{D}\mathbf{D}^T)^{-1}\mathbf{D} + \mathbf{W}_\perp$$

where $\mathbf{D}\mathbf{F}_\perp = \mathbf{D}\mathbf{W}_\perp = \mathbf{0}$.

Conclusions

- We find that the two networks implement drastically different integration algorithms:
  - The linear rate network induces a strong correlation structure to tightly control activity in the coding space, making it susceptible to small perturbations that alter this structure
  - The LIF network actively decorrelates neurons with similar decoding vectors in a chaotic fashion, endowing it with particular robustness to perturbations
- This is evidenced in the significantly higher dimensionality of the dynamics of the LIF network, which is explained by its consistently shorter autocorrelation timescale [6]
- Consistent with the chaotic and decorrelative nature of the algorithm, we find that the dynamics of the LIF integrator network increase in dimensionality with increasing size, suggesting that if such a network existed in the brain it should be easily identified by uncharacteristically high-dimensional dynamics (in the *linear* sense)
- On the other hand, the linear rate network has the property that any subsample of its neurons will show dynamics with the same dimensionality as the network as a whole, which can also be tested experimentally.

Further Questions

- What about learning? The LIF network seems to more easily accomodate integration of higher-dimensional signals, as maintains a large space of solutions in this regime.
- Can the spiking algorithm be implemented with a non-linear rate network? Or is the spiking essential for a network to implement this robust and efficient algorithm?
- How would the presence of more realistic synaptic and conductance-based dynamics in the LIF network affect its dynamics? [2, 7]
- Does this surprising relationship between network size and dimensionality of dynamics hold for other computations under the error-encoding framework? [1, 3]

References

- Martin Boerlin, Christian K Machens, and Sophie Denève. Predictive coding of dynamical variables in balanced spiking networks. *PLoS computational biology*, 9(11):e1003258, 2013.
- Matthew Chalk, Boris Gutkin, and Sophie Deneve. Neural oscillations as a signature of efficient coding in the presence of synaptic delays. *Elife*, 5:e13824, 2016.
- Dominik Thalmeier, Marvin Uhlmann, Hilbert J Kappen, and Raoul-Martin Memmesheimer. Learning universal computations with spikes. *PLoS computational biology*, 12(6):e1004895, 2016.
- Shaul Druckmann and Dmitri B Chklovskii. Neuronal circuits underlying persistent representations despite time varying activity. *Current Biology*, 22(22):2095–2103, 2012.
- H Sebastian Seung. How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences*, 93(23):13339–13344, 1996.
- Peiran Gao, Eric Trautmann, Byron M. Yu, Gopal Santhanam, Stephen Ryu, Krishna Shenoy, and Surya Ganguli. A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv*, 2017.
- Michael A Schwemmer, Adrienne L Fairhall, Sophie Denéve, and Eric T Shea-Brown. Constructing precisely computing networks with biophysical spiking neurons. *Journal of Neuroscience*, 35(28):10112–10134, 2015.

This work was funded by a UCL Graduate Research Scholarship and Overseas Research Scholarship.

Contact: jorge.menendez.15@ucl.ac.uk