Change Rate Inference In Dynamic Environments

Adrian Radillo[1], Alan Veliz-Cuba[2], Krešimir Josić[1,◕], and Zachary Kilpatrick[3,◕]
1. U. of Houston ; 2. U. of Dayton ; 3. U. of Colorado Boulder.
◕. Equal contribution. Correspondence: adrian@math.uh.edu

Introduction

- When an environment is constantly changing and one can only observe it partially, what is the best way to make a perceptual decision?
- What do organisms do?
- We study formally one of the simplest settings in which an observer needs to learn a hidden statistic from her environment.

Setting

- A 2-state environment alternates in time according to a Markov process.

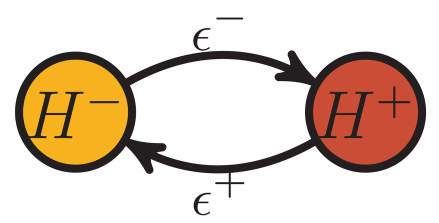

- An observer makes sequential, noisy observations of the environment and must decide, at a fixed time, what the present state is:

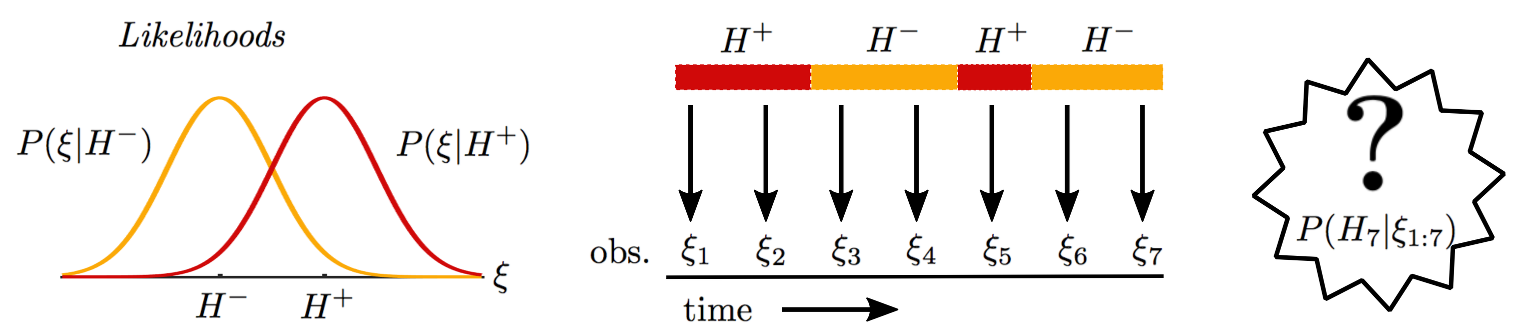

- An ideal observer must *learn the change rates* $\epsilon^{\pm}$ (hazard rate).
- This paradigm applies to several well-known experimental settings, such as a variant of the *random dots task* where the direction of motion switches within trials [1].

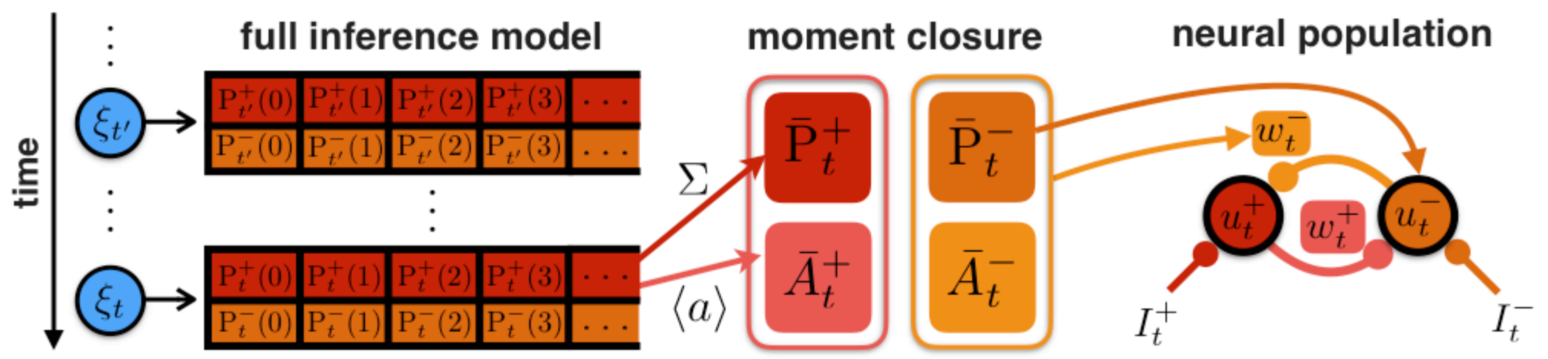

References

[1] Glaze et al (2015) *eLife*, 4, p1; [2] Radillo et al (preprint) *Neural Comp*, arXiv:1607.08318v2; [3] Veliz-cuba et al (2015) *Arxiv*, arXiv:1505.04195v1; [4] Wilson et al (2010) *Neural Comp*, 22, p2452.

An impractical normative model

- The optimal on-line algorithm is costly in the discrete-time setting when the rates are asymmetric.

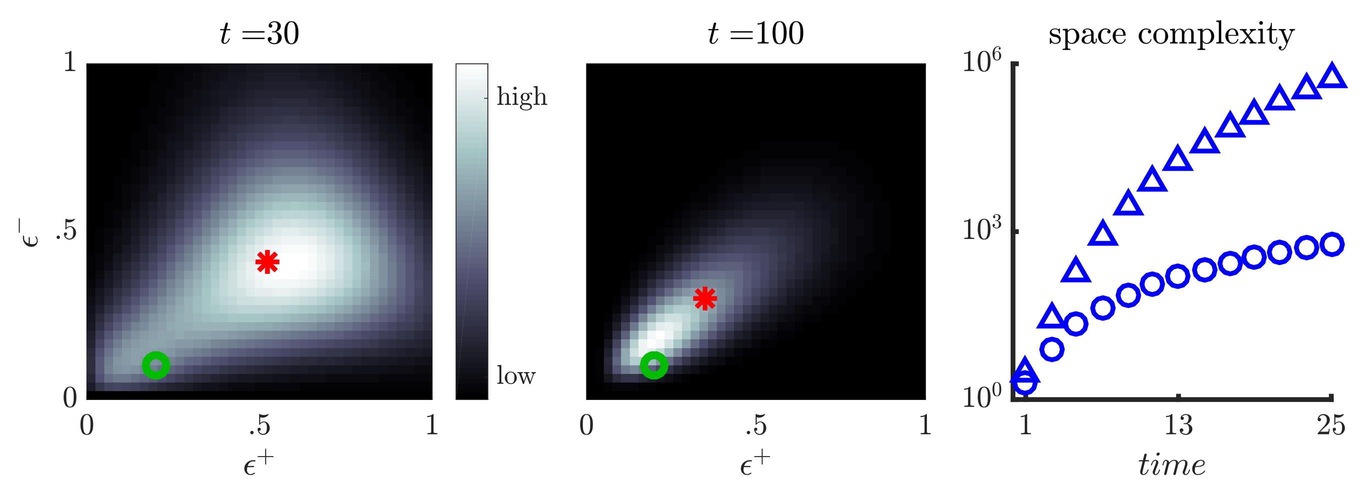

Approximate system of smaller dimension

- We reduce the dimensionality of the continuous-time system via moment closure:

$$d\bar{P}_t^{\pm} = \bar{P}_t^{\pm} \cdot I^{\pm}(t) + \left[\bar{A}_t^{\mp} - \bar{A}_t^{\pm}\right] dt$$

$$d\bar{A}_t^{\pm} = \bar{A}_t^{\pm} \cdot I^{\pm}(t) + \left[\bar{A}_t^{\mp} - \bar{A}_t^{\pm}\right] \left(\frac{1}{t+\beta} + \bar{A}_t^{+} + \bar{A}_t^{-}\right) dt$$

- The variables $\bar{P}_t^{\pm}$ are marginals over the states.
- The change point count averages $\bar{A}_t^{\pm}$ increase when evidence for a change point is presented.

When to learn?

- We explore in which parameter regimes learning the change rate is both *possible* and *useful*.

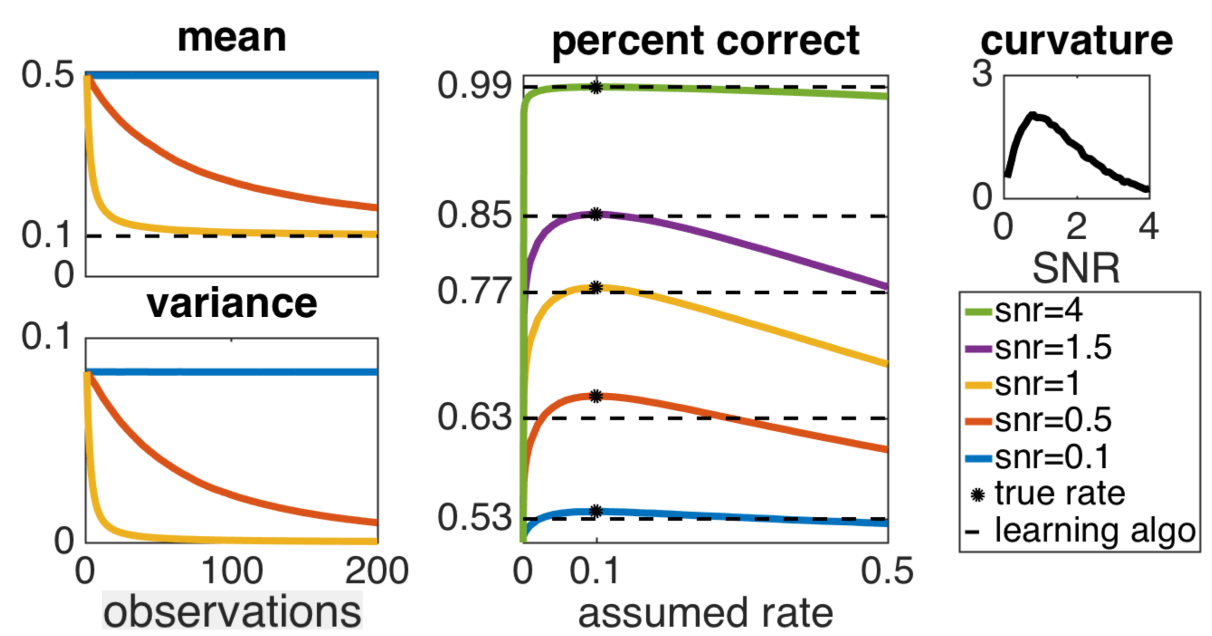

- Learning the rate is most helpful at intermediate values of signal-to-noise ratio (SNR).

Neural networks

- Moment-closure equations motivate a plastic, rate-based neural network.
- We define two excitatory neural population variables, $u_t^{\pm} = \bar{P}_t^{\pm}$, and the synaptic weights between them, $w_t^{\pm} = \bar{A}_t^{\pm}/u_t^{\pm}$.
- A *correlation-based plasticity rule* makes the weights converge to the hazard rate.

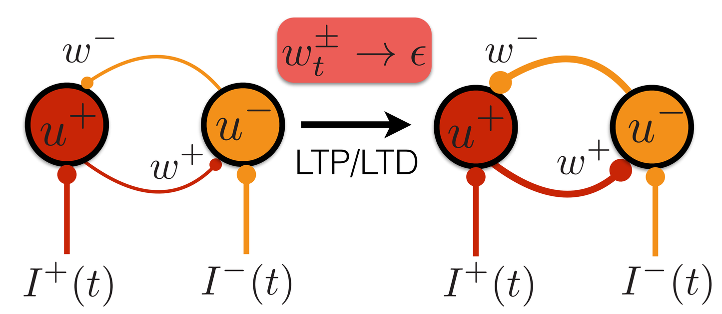

$$du_t^{\pm} = u_t^{\pm} \cdot I^{\pm}(t) + \left[w_t^{\mp} u_t^{\mp} - w_t^{\pm} u_t^{\pm}\right] dt$$

$$dw_t^{\pm} = \left[\delta\left(u_t^{+} - u_t^{-}\right) - w_t^{\pm}\right] \cdot C_t \, dt$$

- For asymmetric change rates, the learning rule becomes:

$$dw_t^{\pm} = H\left(u_{t-\tau}^{\pm} - \theta\right)\left[\delta\left(u_t^{+} - u_t^{-}\right) - w_t^{\pm}\right] \cdot C_t^{\pm} \, dt$$

(H =Heaviside, C_t =amplitude-decay, θ = activity threshold)

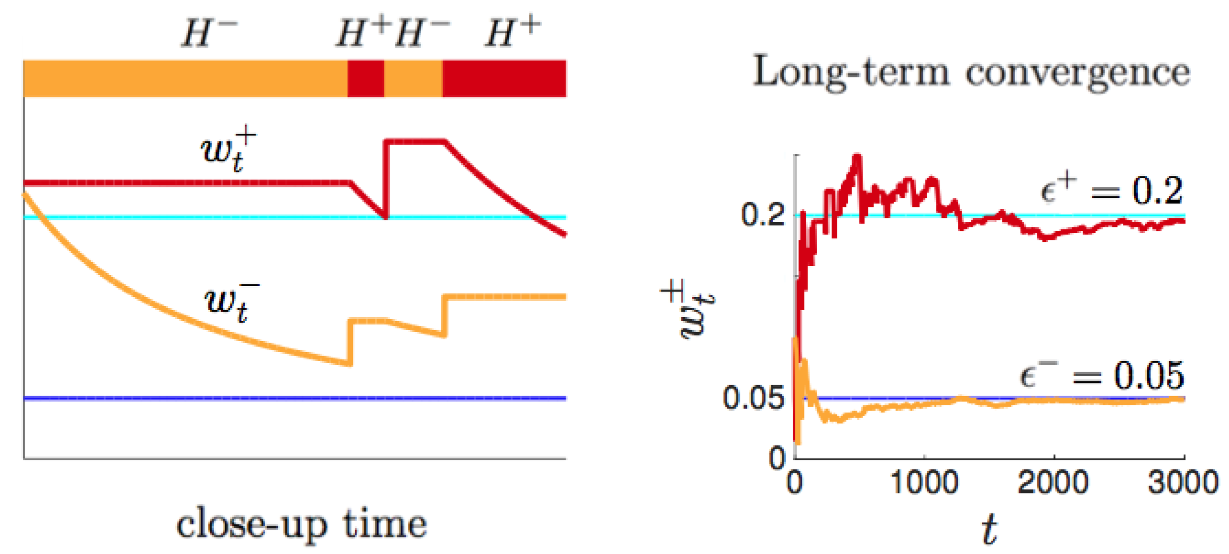

- As one population dominates, only one weight decays until it receives a pulse-increase at the change point.

Conclusion

- We derived [2] an optimal algorithm that solves the task in a vast array of conditions:

	2-state		N-state	
	discrete	continuous	discrete	continuous
symmetric	✔	✔	✔	✔
asymmetric	✔	✔	✔	✘

- The dimension of these models is generally too high.
- We implemented an approximation of the normative models in a plastic rate-based neural network.