Matching tutor to student

rules and mechanisms for efficient two-stage learning in neural circuits

Tiberiu Teşileanu[1], *Bence Ölveczky*[2], *Vijay Balasubramanian*[1]

[1] Initiative for the Theoretical Sciences, CUNY Graduate Center and David Rittenhouse Laboratories, University of Pennsylvania

[2] Center for Brain Science and Department of Organismic and Evolutionary Biology, Harvard University

Overview

Learning in the brain often occurs in stages: one brain region (the **tutor**) learns first before information is transferred to a downstream circuit (the **student**).

We investigate a two-stage learning model in which one area processes a reward signal to generate an input that guides learning in the downstream area.

We find:

- fast learning requires a match between tutor signal and student plasticity rule.
- the matched tutor requires a memory trace of the output error.
- results apply in both rate-based and spiking networks.
- during learning, firing in student neurons becomes burstier, and connections are pruned, as seen in songbirds.

Model

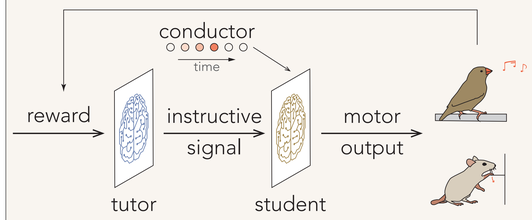

conductor provides timebase

student neurons project to motor areas

tutor signal modulates conductor–student plasticity

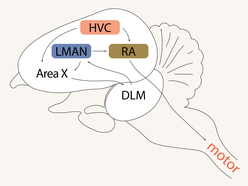

model inspired by vocal production system in songbirds

HVC=conductor

RA=student

LMAN=tutor

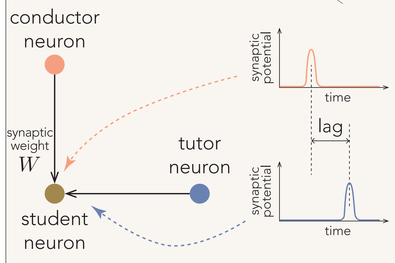

conductor–student synaptic strengths controlled by tutor signal

given a conductor burst and a tutor burst separated by a lag, the synaptic weight changes according to the plasticity rule:

$$\Delta W \sim \frac{\alpha}{\tau_1} e^{-\mathrm{lag}/\tau_1} - \frac{\beta}{\tau_2} e^{-\mathrm{lag}/\tau_2}$$

We require that the weights change in the direction of steepest decrease in error.

- tutor signal must average the motor error over a certain timescale τ_{tutor}.
- optimal timescale τ^*_{tutor} depends on the structure of the student plasticity rule:

$$\tau^*_{\mathrm{tutor}} = \frac{\alpha\tau_1 - \beta\tau_2}{\alpha - \beta}$$

plasticity parameters → tutor memory timescale

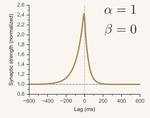

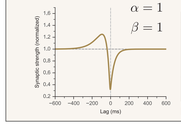

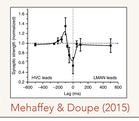

Mehaffey & Doupe (2015)

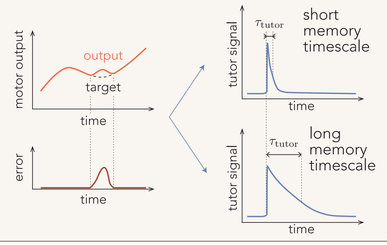

Results

Rate-based

example traces

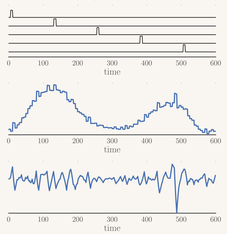

conductor

student

tutor

Spiking

example traces/spiking rasters

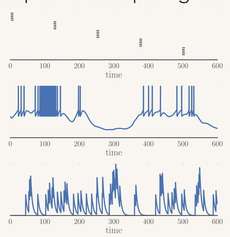

learning proceeds efficiently when tutor and student are matched

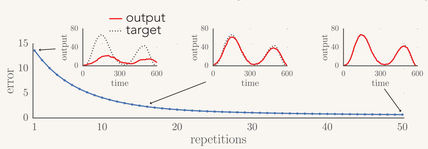

$\alpha = 0\,, \quad \beta = -1\,, \quad \tau_{\mathrm{tutor}} = 40$

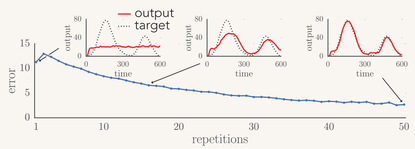

$\alpha = 1\,, \quad \beta = 0\,, \quad \tau_{\mathrm{tutor}} = 80$

learning slows down or breaks down when tutor and student are mismatched

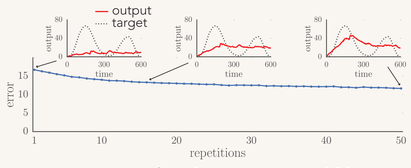

$\alpha = 1\,, \quad \beta = 0\,, \quad \tau_{\mathrm{tutor}} = 1280$

(matched tutor would have $\tau^*_{\mathrm{tutor}} = 80$)

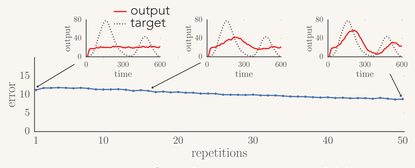

$\alpha = 1\,, \quad \beta = 0\,, \quad \tau_{\mathrm{tutor}} = 1280$

(matched tutor would have $\tau^*_{\mathrm{tutor}} = 80$)

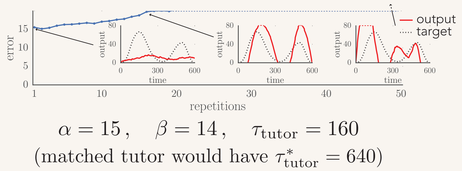

$\alpha = 15\,, \quad \beta = 14\,, \quad \tau_{\mathrm{tutor}} = 160$

(matched tutor would have $\tau^*_{\mathrm{tutor}} = 640$)

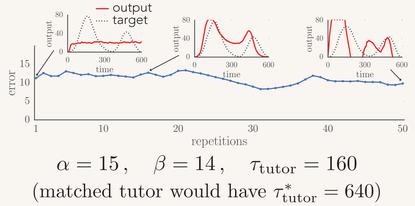

$\alpha = 15\,, \quad \beta = 14\,, \quad \tau_{\mathrm{tutor}} = 160$

(matched tutor would have $\tau^*_{\mathrm{tutor}} = 640$)

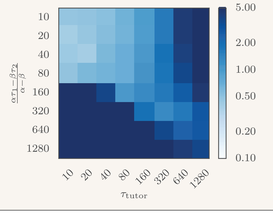

reproduction error after 250 steps depends strongly on student–tutor matching

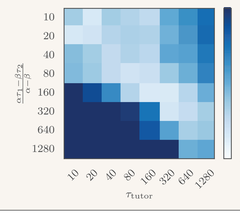

Results

learning leads to burstier firing from student neurons, in agreement with recordings from songbird LMAN

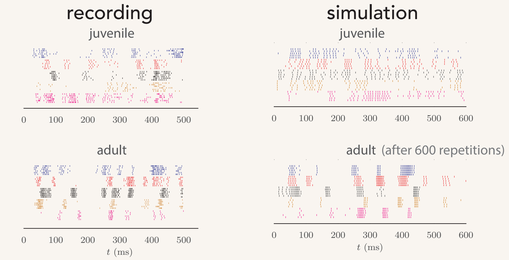

learning is sequential when tutor memory is very long

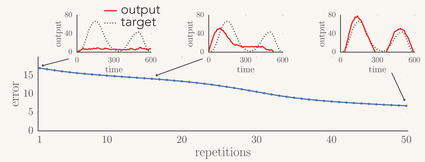

this is due to the tutor output reaching phisiologically-realistic limits

(tutor firing rates were constrained to lie between 0 and 160 Hz)

$\alpha = 24\,, \quad \beta = 23\,, \quad \tau_{\mathrm{tutor}} = \tau^*_{\mathrm{tutor}} = 1000$

tutor can generate its signal through reinforcement learning

learning is much slower but still accurate

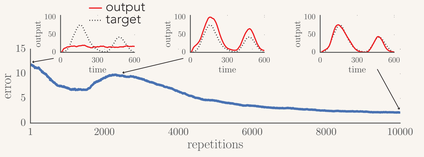

learning drives pruning of connections

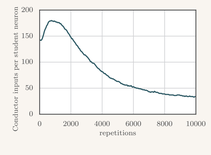

Conclusion

We built a framework for investigating information transfer between brain regions. We used a gradient descent approach on a rate-based model to predict the structure of the teaching signal that best matches the synaptic plasticity rule in the student circuit.

Using computer simulations, we showed that departures from our matching rule can lead to slowed or even abolished learning, and that our results hold in spiking as well as in rate-based networks. Finally, we showed how the tutor circuit can generate its signal using a reinforcement-learning strategy.

For details, see our preprint on the *arXiv* at https://arxiv.org/abs/1608.08040.

References

Andalman and Fee (2009). PNAS 106(30).

Fiete *et al.* (2007). J. Neurophysiology 98(4).

Mehaffey and Doupe (2015). Nature Neuroscience 18(9).