Credit Assignment in Neural Networks through Deep Feedback Control

Alexander Meulemans*, Matilde Tristany Farinha*, Javier García Ordóñez,
Pau Vilimelis Aceituno, João Sacramento, Benjamin F. Grewe

Credit assignment (CA)

" How does the strength of a synapse need to be changed to improve the system's global behaviour?"

Spatial Credit Assignment

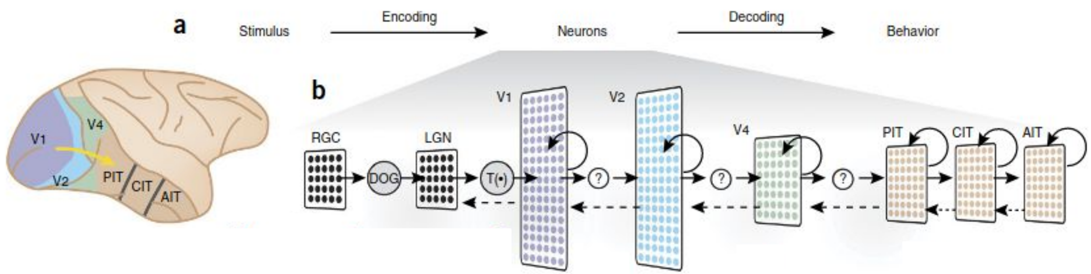

Yamins & Dicarlo 2016

Spatial credit assignment: Backpropagation

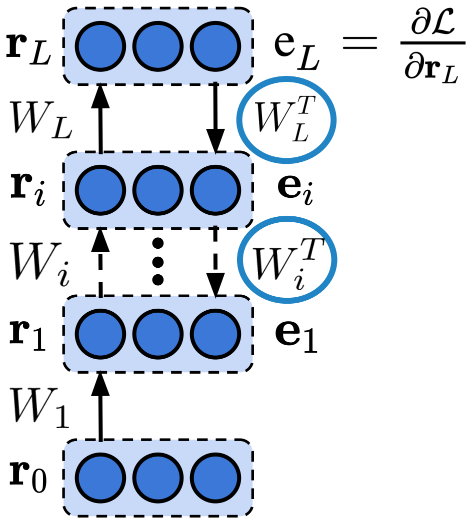

Some biological issues:

1. Weight transport

2. Feedback does not influence neuron activations only synaptic strength

Research questions

"Is principled **credit assignment without strict alignment** between the feedback path and feedforward path possible?"

"Can we use **feedback** not only **to learn** the synapses but also to change the **neural activations**?"

Outline

Part I: The intuition behind Deep Feedback Control (DFC)

Part II: Theoretical analysis of DFC

Part III: Learning the feedback weights for DFC

Part VI: Simulation results

Deep Feedback Control (DFC): intuition

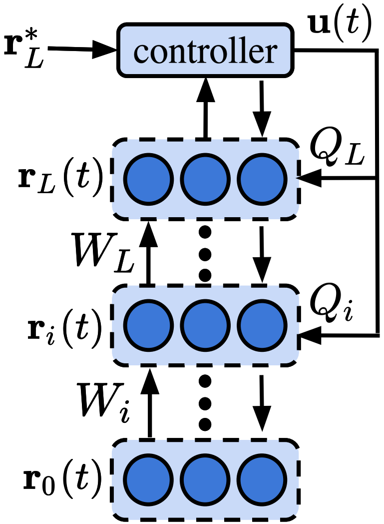

feedback

⇩ adjust

neural activations

⇩ learn

synapses

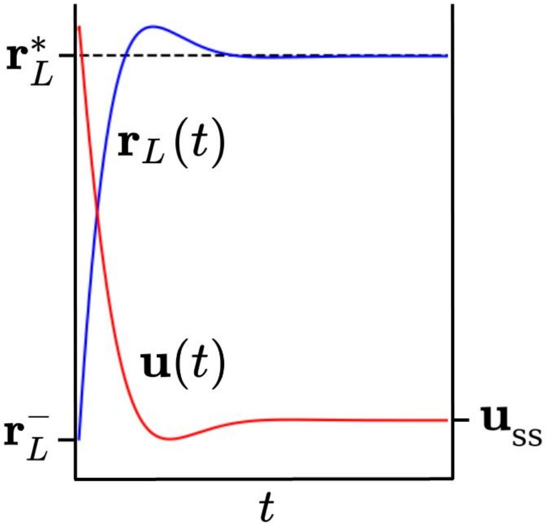

Gilra & Gerstner 2017
Alemi et al. 2017

Deep Feedback Control: dynamics

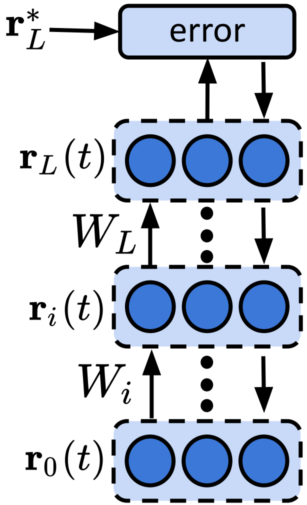

Network dynamics:

$$\tau_v \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}_i(t) = -\mathbf{v}_i(t) + W_i \phi\big(\mathbf{v}_{i-1}(t)\big)$$

$$\mathbf{r}_i = \phi(\mathbf{v}_i)$$

Steady state:

$$\mathbf{r}_{i,\mathrm{ss}} = \phi(W_i \mathbf{r}_{i-1,\mathrm{ss}})$$

Deep Feedback Control: dynamics

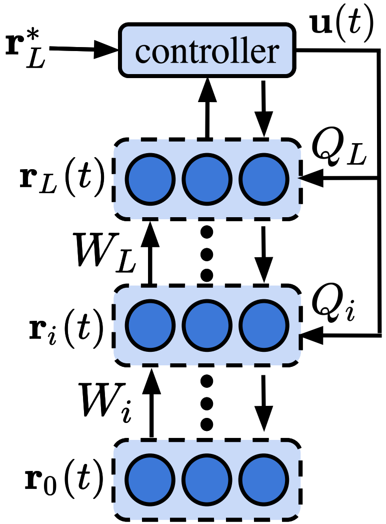

Network dynamics:

$$\tau_v \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}_i(t) = -\mathbf{v}_i(t) + W_i \phi\big(\mathbf{v}_{i-1}(t)\big) + Q_i \mathbf{u}(t)$$

$$\mathbf{r}_i = \phi(\mathbf{v}_i)$$

Deep Feedback Control: controller

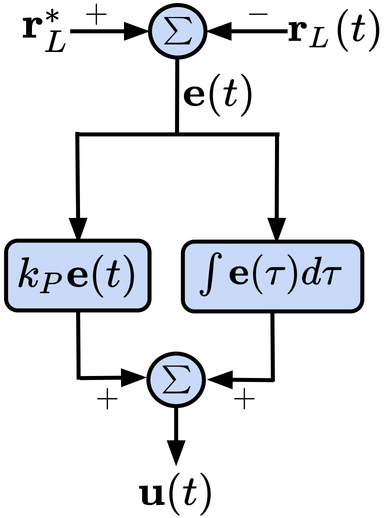

Network dynamics:

$$\tau_v \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}_i(t) = -\mathbf{v}_i(t) + W_i\phi\big(\mathbf{v}_{i-1}(t)\big) + Q_i\mathbf{u}(t)$$

$$\mathbf{r}_i = \phi(\mathbf{v}_i)$$

Proportional-Integral control

$$\mathbf{u}(t) = K_I\mathbf{u}^{\text{int}}(t) + K_P\mathbf{e}(t), \quad \tau_u \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{u}^{\text{int}}(t) = \mathbf{e}(t) - \alpha\mathbf{u}^{\text{int}}(t)$$

✔ Feedback strongly changes the neural activity

? Can this setting also be used for learning?

Deep Feedback Control: plasticity

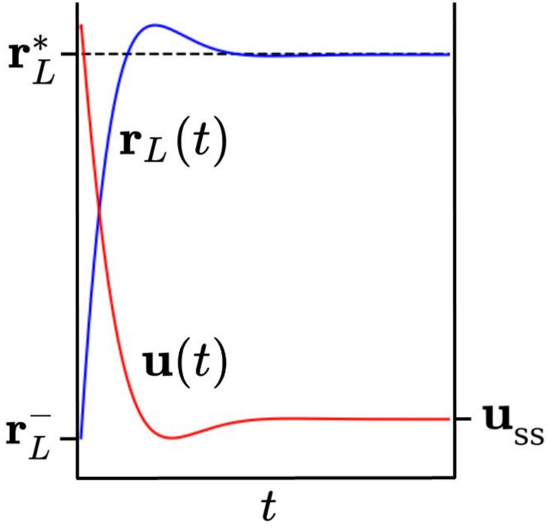

$$\mathbf{r}_L(t) = \phi\big(\mathbf{v}_L(t)\big)$$

With an active controller, the network settles to the desired output target: $\mathbf{r}_L(t) \longrightarrow \mathbf{r}_L^*$

→ Steady-state hidden layer activations can be seen as *"hidden targets"*

→ Difference between hidden targets and initial (feedforward) activations can be used as a learning signal

Deep Feedback Control: plasticity

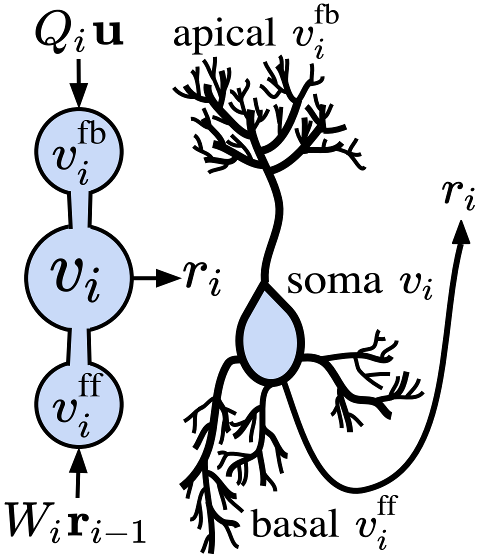

$$\tau_W \frac{d}{dt} W_i(t) = \left(\phi(\mathbf{v}_i(t)) - \phi(\mathbf{v}_i^{\text{ff}}(t))\right)\mathbf{r}_{i-1}(t)^T$$

local in **space** and **time**

Urbanczik & Senn, 2014

Summary part I

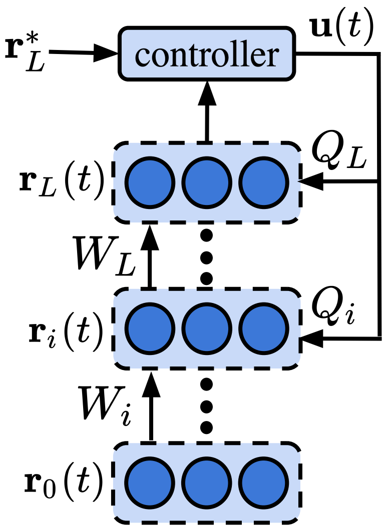

- DFC uses a feedback controller to push the network to a desired output target
- The changed neural activations are used for credit assignment
- Using a multi-compartment neuron makes the learning local in time and space

Outline

Part I: The intuition behind Deep Feedback Control (DFC)

Part II: Theoretical analysis of DFC

- Does DFC perform principled credit assignment?
- Under which conditions is DFC stable?

Part III: Learning the feedback weights for DFC

Part VI: Simulation results

DFC approximates Gauss-Newton (GN) optimization

DFC approximates Gauss-Newton optimization (minibatch size = 1)

IF

$$\mathrm{Col}(Q) = \mathrm{Col}(J^T)$$

$$J = \frac{\partial \mathbf{r}_L}{\partial \mathbf{v}}$$

AND

$$\mathbf{r}_L^* = \lim_{\lambda \to 0} \mathbf{r}_L^- - \lambda \frac{\partial \mathcal{L}}{\partial \mathbf{r}_L}$$

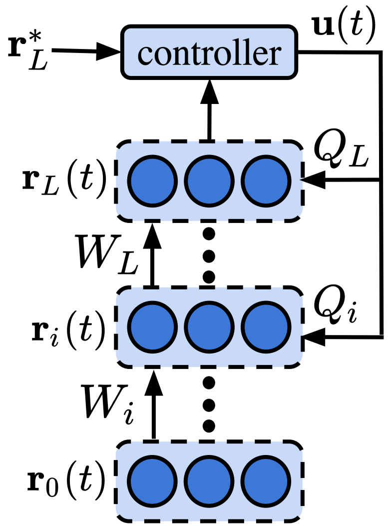

Minimum norm updates: intuition

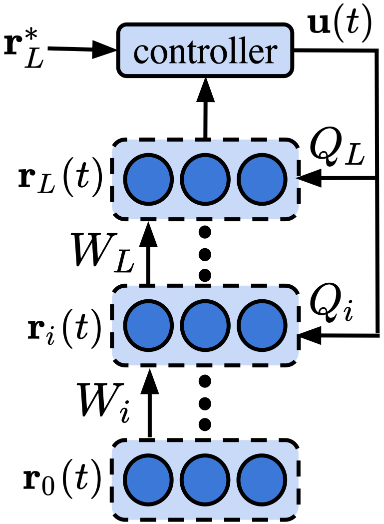

- Controller pushes the network to its output target
- Many possible configurations of the network reach exactly the output target

… Which configuration is the best?

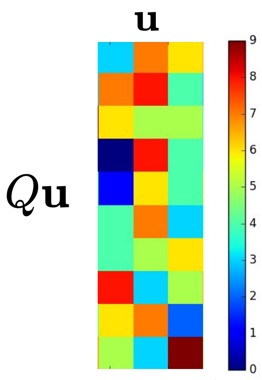

$$\mathrm{Col}(Q) = \mathrm{Col}(J^T)$$

Stability of DFC

Simplified dynamics (linearized + separation of timescales):

$$\tau_u \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{u}(t) = -JQ\mathbf{u}(t) + \boldsymbol{\delta}_L$$

⇨ Local stability if $\lambda_{\max}(-JQ) < 0$

Flexibility of the feedback path

Optimal credit assignment when

$$\mathrm{Col}(Q) = \mathrm{Col}(J^T)$$

$$\lambda_{\max}(-JQ) < \alpha$$

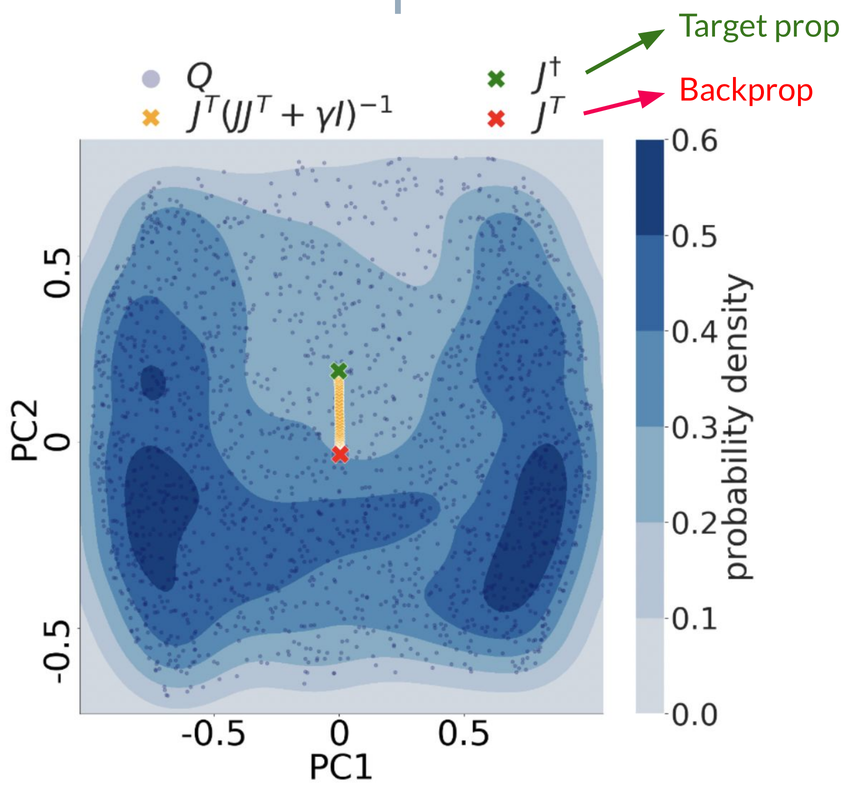

Summary part II

DFC can do principled credit assignment without the need for strict alignment

$$\mathrm{Col}(Q) = \mathrm{Col}(J^T)$$

$$\lambda_{\max}(-JQ) < \alpha$$

… How to make sure the flexible conditions are satisfied?

Learning the feedback weights in DFC

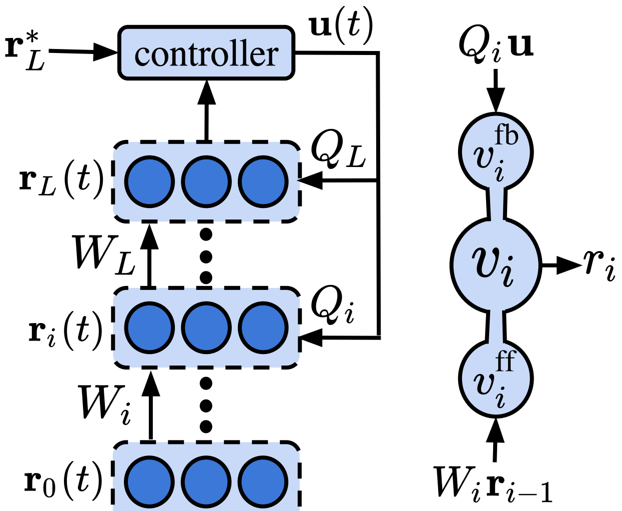

Add noise to dynamics:

$$\tau_v \frac{\mathrm{d}\mathbf{v}_i}{\mathrm{d}t} = -\mathbf{v}_i + W_i\phi(\mathbf{v}_{i-1}) + Q_i\mathbf{u} + \sigma\boldsymbol{\xi}_i$$

Feedback plasticity rule:

$$\tau_Q \frac{\mathrm{d}}{\mathrm{d}t} Q_i(t) = -\mathbf{v}_i^{\text{fb}}(t)\mathbf{u}(t)^T - \beta Q_i$$ ⇨ Anti-Hebbian

Feedback weights align with:

$$\mathbb{E}[Q_{\text{ss}}] \underset{\sim}{\propto} J^T(JJ^T + \gamma I)^{-1}$$

⇨ Satisfies column space condition
⇨ Satisfies stability condition

Outline

Part I: The intuition behind Deep Feedback Control (DFC)

Part II: Theoretical analysis of DFC

- Does DFC perform principled credit assignment?
- Under which conditions is DFC stable?

Part III: Learning the feedback weights for DFC

Part VI: Simulation results

Toy experiments

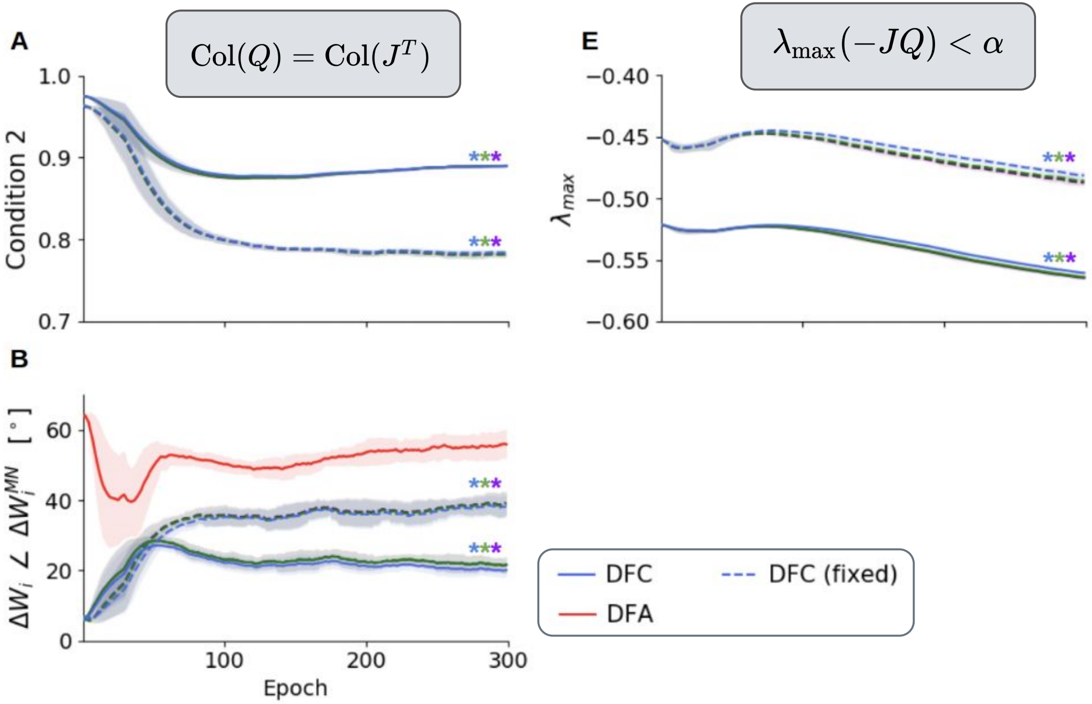

Computer vision experiments

	MNIST	Fashion-MNIST	MNIST autoencoder	MNIST (train loss)
BP	$2.08^{\pm 0.15}\%$	$10.60^{\pm 0.34}\%$	$9.42^{\pm 0.09} \cdot 10^{-2}$	$1.53^{\pm 0.19} \cdot 10^{-7}$
DFC	$2.25^{\pm 0.094}\%$	$11.17^{\pm 0.27}\%$	$11.28^{\pm 0.18} \cdot 10^{-2}$	$7.61^{\pm 0.65} \cdot 10^{-8}$
DFC (fixed)	$2.47^{\pm 0.12}\%$	$11.62^{\pm 0.30}\%$	$33.37^{\pm 0.60} \cdot 10^{-2}$	$1.30^{\pm 0.15} \cdot 10^{-6}$
DFA	$2.69^{\pm 0.11}\%$	$11.38^{\pm 0.25}\%$	$29.95^{\pm 0.36} \cdot 10^{-2}$	$7.09^{\pm 1.11} \cdot 10^{-7}$

Conclusion

- DFC uses a feedback controller to drive the network to a desired output target
- Learning rule local in time and space
- Optimal credit assignment without the need for strict alignment
- Intimate connection with cortical pyramidal neurons

Thank you!

Javier García Ordóñez

Pau Vilimelis Aceituno

João Sacramento

Benjamin Grewe

FNSNF
SWISS NATIONAL SCIENCE FOUNDATION

HFSP International Human Frontier Science Program Organization

SDSC

paper

Contact: ameulema@ethz.ch