Learning low-dimensional inputs for brain-machine interface control

Jorge A. Menéndez[1,2] & Peter E. Latham[1]
[1]Gatsby Computational Neuroscience Unit, [2]CoMPLEX; University College London

Observation:

learning to control certain BMI decoders...

- is fast
  Jarosiewicz et al '08; Sadtler et al. '14
- conserves population activity structure
  Hwang et al. '13; Golub et al. '18, Hennig et al. '18
- maintains memory of past tasks
  Jarosiewicz et al. '08; Ganguly et al. '11; Sadtler et al. '14

Hypothesis #1: *synaptic plasticity*

- × slow...
  Werfel et al. '03; Miconi et al. '17
- × changes activity modes
  Mastrogiuseppe & Ostojic '18
- × catastrophic forgetting
  no known local learning rules to avoid this

Hypothesis #2: **re-aiming**

- ✓ low-dimensional
  gradient-free methods will work
- ✓ network structure intact
  input lives on a fixed 2D manifold
- ? can be memorized?
  low-d variables ⇒ low-d attractors?

Model

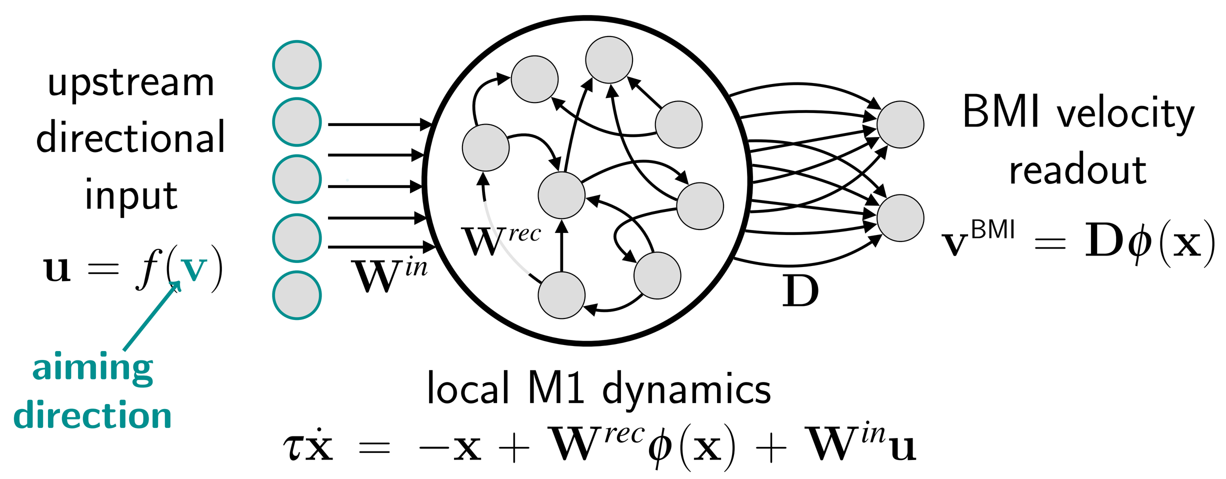

Analyze simplified control problem:

find the re-aiming direction that will drive the network to achieve a target readout at some time t

In other words, for a given network and decoder $\mathbf{D}$, solve:

$$\underbrace{\tilde{\mathbf{v}}}_{\text{optimal re-aiming direction}} = \arg\min_{\mathbf{v}} \underbrace{\|\mathbf{D}\phi\left(\mathbf{x}(t;\mathbf{v},\mathbf{x}_0)\right) - \mathbf{v}^*\|^2}_{\substack{E_t(\mathbf{v},\mathbf{v}^*)\\ \text{reaching error}}} + \gamma \underbrace{\langle u_i^2\rangle}_{\substack{C(\mathbf{v})\\ \text{upstream metabolic cost}}}$$

For **linear** dynamics $\phi(\mathbf{x}) = \mathbf{x}$ and encoding $f(\mathbf{v}) = \mathbf{M}\mathbf{v}$,

$$\mathbf{x}(t;\mathbf{v},\mathbf{x}_0) = e^{(\mathbf{W}^{rec}-\mathbf{I})\frac{t}{\tau}}\mathbf{x}_0 + \mathbf{F}(t)\mathbf{v}$$

$$\Rightarrow \tilde{\mathbf{v}} = \mathbf{G}\underbrace{\left(\mathbf{v}^* - \mathbf{D}e^{(\mathbf{W}^{rec}-\mathbf{I})\frac{t}{\tau}}\mathbf{x}_0\right)}_{\delta\mathbf{v}^*}$$

intrinsic manifold: network activity $\mathbf{x}(t)$ is restricted to live on the plane defined by the columns $\{\mathbf{f}_i\}_{i=1}^2$ of

$$\mathbf{F}(t) = \left(e^{(\mathbf{W}^{rec}-\mathbf{I})\frac{t}{\tau}} - \mathbf{I}\right)(\mathbf{W}^{rec}-\mathbf{I})^{-1}\mathbf{W}^{in}\mathbf{M}$$

where $\mathbf{G} = \left(\mathbf{\Lambda}^T\mathbf{\Lambda} + \gamma\mathbf{M}^T\mathbf{M}\right)^{-1}\mathbf{\Lambda}^T$ depends on the 2×2 *alignment matrix* $\mathbf{\Lambda}$:

$$\mathbf{\Lambda} = \mathbf{D}\mathbf{F}(t) \Rightarrow \Lambda_{ij} = \mathbf{d}_i^T\mathbf{f}_j \begin{cases} \mathcal{O}(1) & \text{if } \mathbf{d}_i, \mathbf{f}_j \text{ random} \\ \mathcal{O}(\sqrt{n}) & \text{if aligned} \end{cases}$$

In fact, for $\delta\mathbf{v}^*$ uniformly distributed on a circle of radius r, the average optimal error depends directly on $\mathbf{\Lambda}$ through its singular values s_i:

$$E^{\text{opt}} \equiv \langle E_t(\tilde{\mathbf{v}},\mathbf{v}^*)\rangle_{\delta\mathbf{v}^*} = \frac{r^2}{2}\sum_{i=1}^{2}\frac{1}{(s_i^2/\gamma+1)^2}\begin{cases} = 0 & \text{if } \gamma = 0 \\ \to r^2 & \text{as } s_i^2 \to 0 \\ \to 0 & \text{as } s_i^2 \to \infty \end{cases}$$

"Neural constraints on learning"

Is re-aiming always a viable strategy?

NO: for $\gamma > 0$, if s_i's are too small, then error will be large on average

This will happen whenever Λ_{ij} are small, i.e. when the decoder and intrinsic manifold are misaligned!

Consider decoders from Sadtler et al. '14:

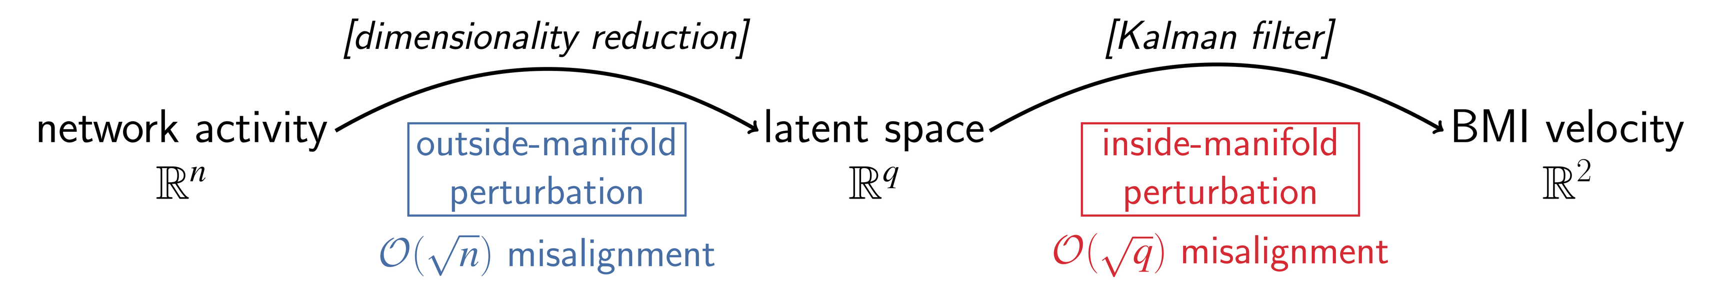

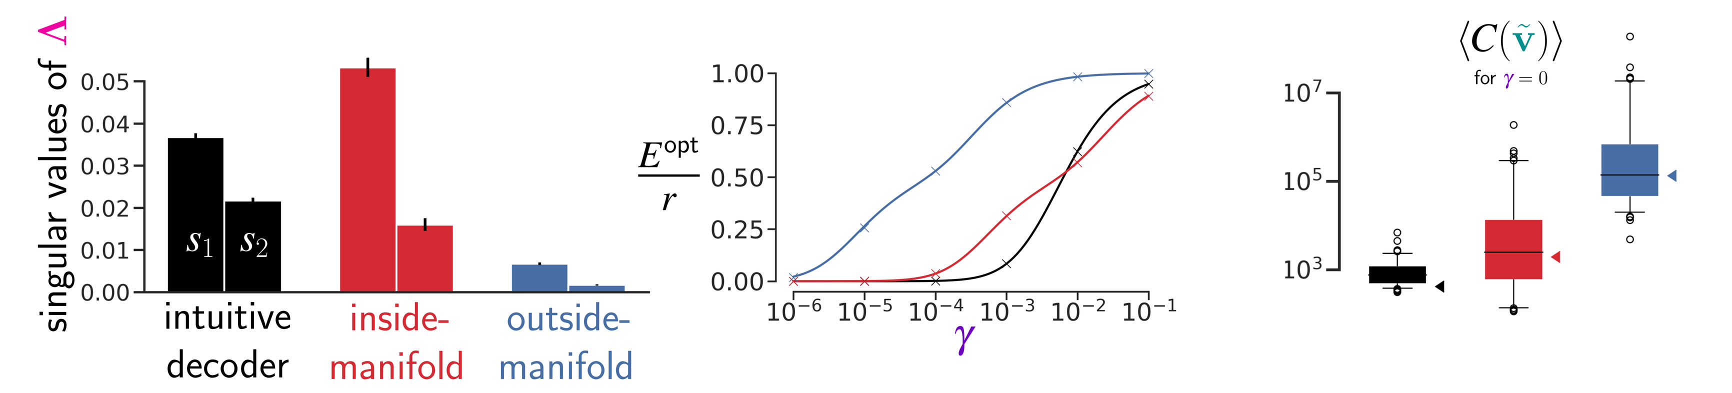

Alternative strategy (re-scaling + re-aiming): reduce *speed* of target velocity $\mathbf{v}^* \to k\hat{\mathbf{v}}^*$ so that $C(\tilde{\mathbf{v}}) \leq C^{max}$. Is this possible for any $\hat{\mathbf{v}}^*$?

NO: given $\mathbf{x}_0$, $C(\tilde{\mathbf{v}}) > C^{min}(\mathbf{x}_0)$, so this is only possible if $C^{min}(\mathbf{x}_0) < C^{max}$!

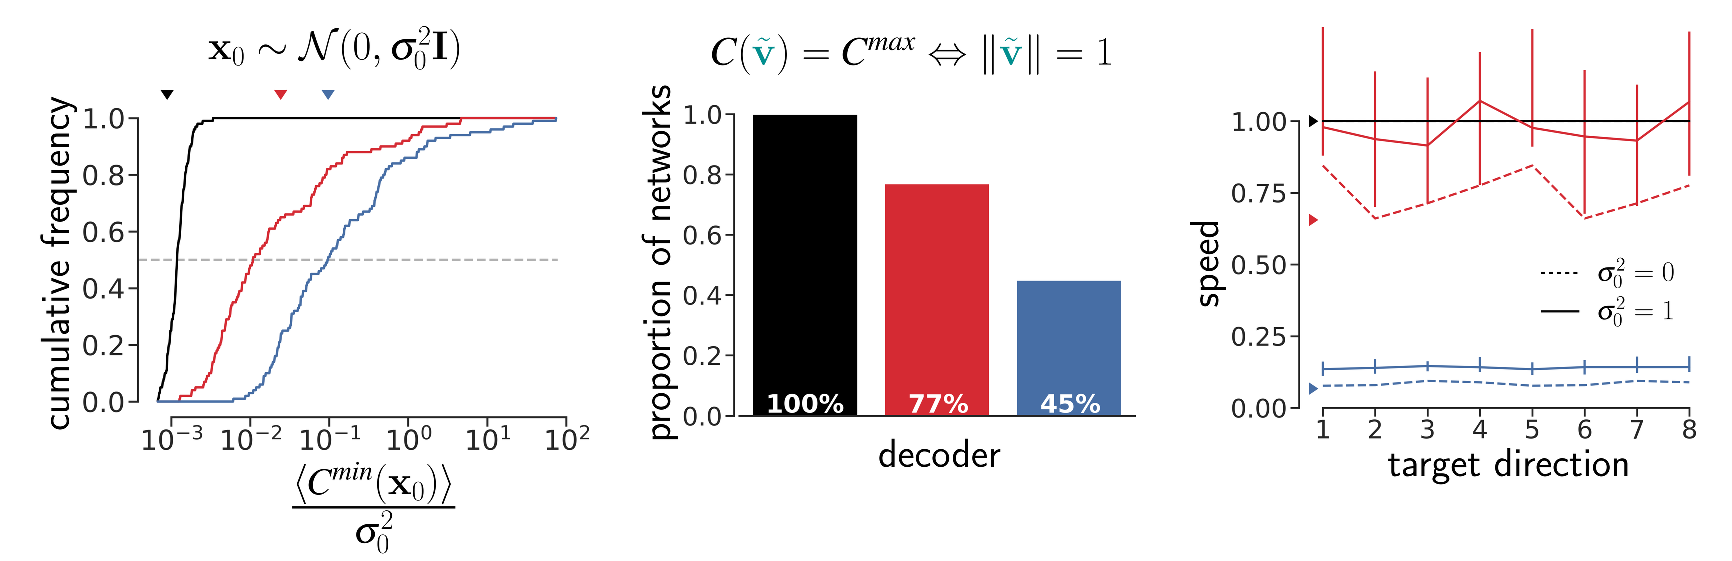

"Learning by neural reassociation"

[Golub et al. '18; Hennig et al. '18]

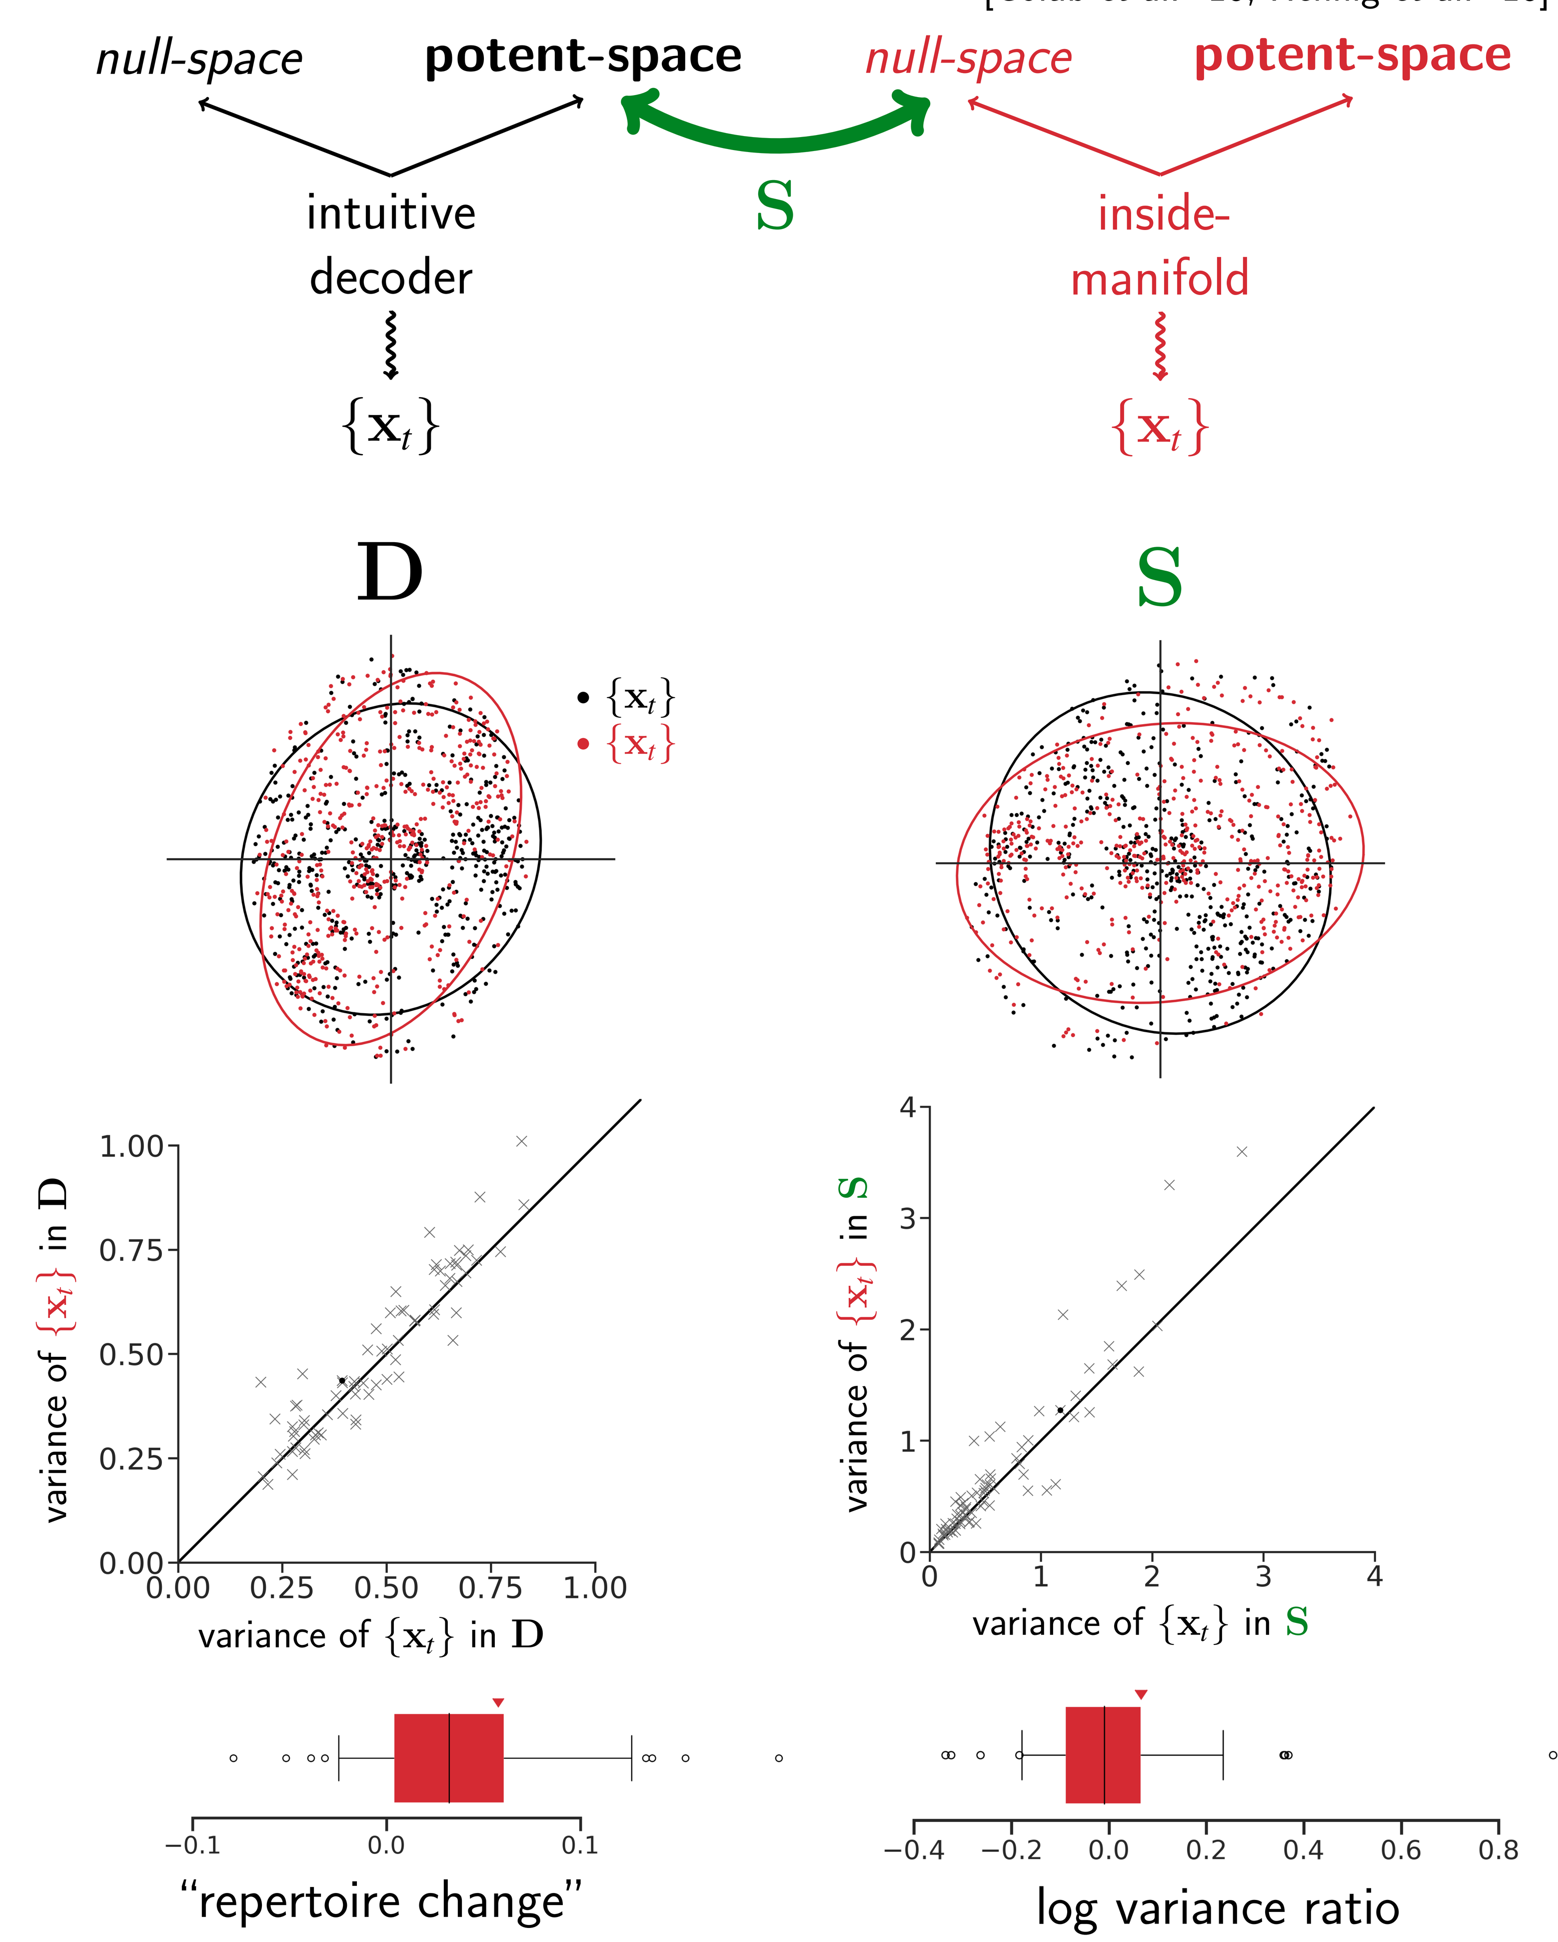

Global response modulation

[Ganguly et al. '11; Hwang et al. '13]

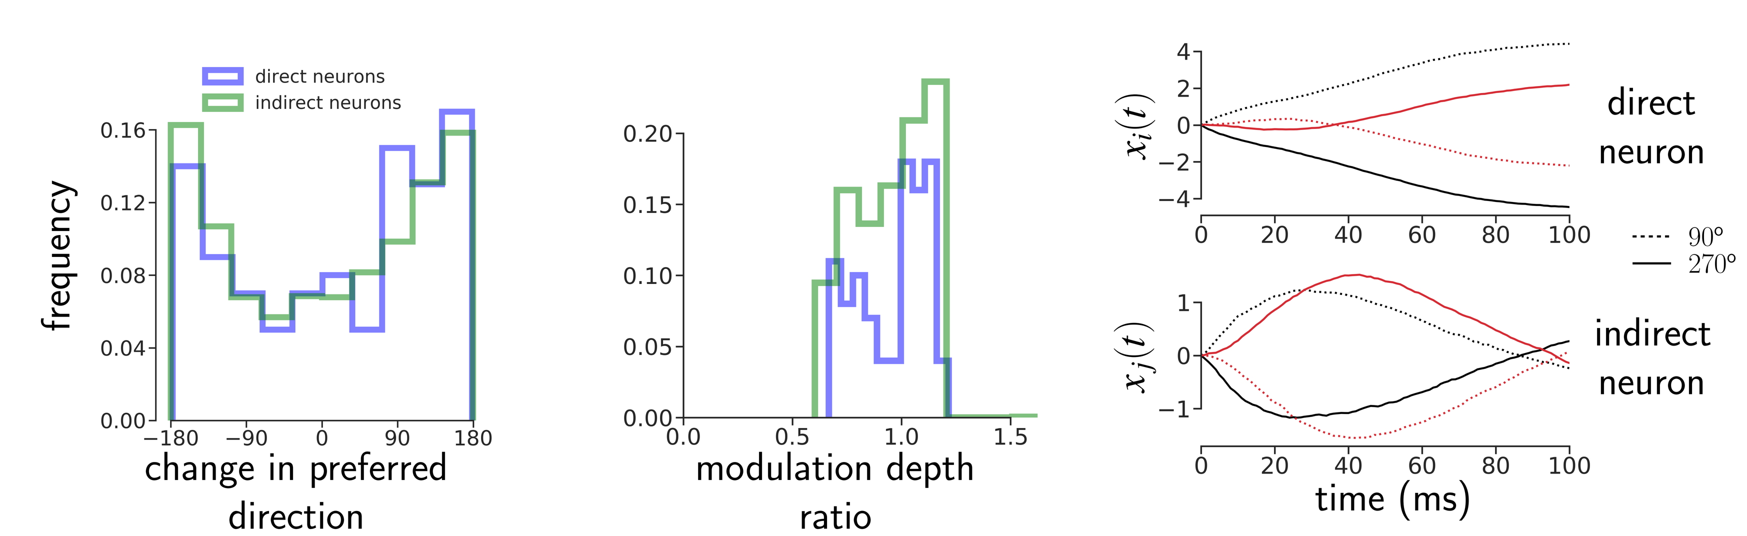

Summary & Future Directions

A **re-aiming + re-scaling** strategy can account for several observations of **fast-timsecal learning** of BMI control

This strategy requires **learning low-dimensional parametric inputs**
⇒how are these learnt?

How might this process interact with **slow-timescale learning?**

Extension to **non-linear** and **structured** networks: what is $\mathbf{F}(t)$?

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

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