Disentangling the roles of dimensionality and cell populations in neural computations

Adrian Valente[1], Alexis Dubreuil[1], Francesca Mastrogiuseppe[2], Srdjan Ostojic[1]

1. LNC, Ecole Normale Supérieure, PSL Research University, Paris, France
2. Gatsby Computational Neuroscience Unit, UCL, London

Introduction

- Low-dimensional dynamics appear to underlie computations in the brain and in RNNs. Low-rank RNNs provide a natural model for this, both flexible and interpretable [1].
- We train low-rank RNNs for several tasks to understand how dimensionality relates to computation, and how low-dimensional dynamics arise from a specific connectivity.
- We observe that for some tasks, different populations with heterogeneous connectivity statistics are necessary.
- **When are heterogeneous populations needed?**
- We present the mixture-of-gaussians low-rank RNN as a minimally parametrized model of low-dimensional network activity with multiple populations.
- **For each task we find the minimal rank, and minimal number of populations needed. This gives us a reduced network model, and a mechanistic explanation for each task.**

Approach

Train minimal rank RNN → Reverse-engineer: Study dynamics [2] ⇅ Find minimal number of populations → Construct a class of minimally parametrized solutions

Tasks studied [5]

Behavioral task	Cognitive operations	Minimal rank / # of cell classes
Random dot motion task	stimulus integration	K = 1 P = 1
Context-dependent integration [3]	stimulus integration contextual gating	K = 1 P = 2
Multisensory integration [4]	stimulus integration accross modalities	K = 1 P = 1
Delayed comparison task	parametric working memory comparison	K = 2 P = 1
Delayed Match-to-Sample (two items {A,B})	object working memory comparison	K = 2 P = 2

Increasing rank with a single population

Random dots motion task — rank=1, #pops = 1

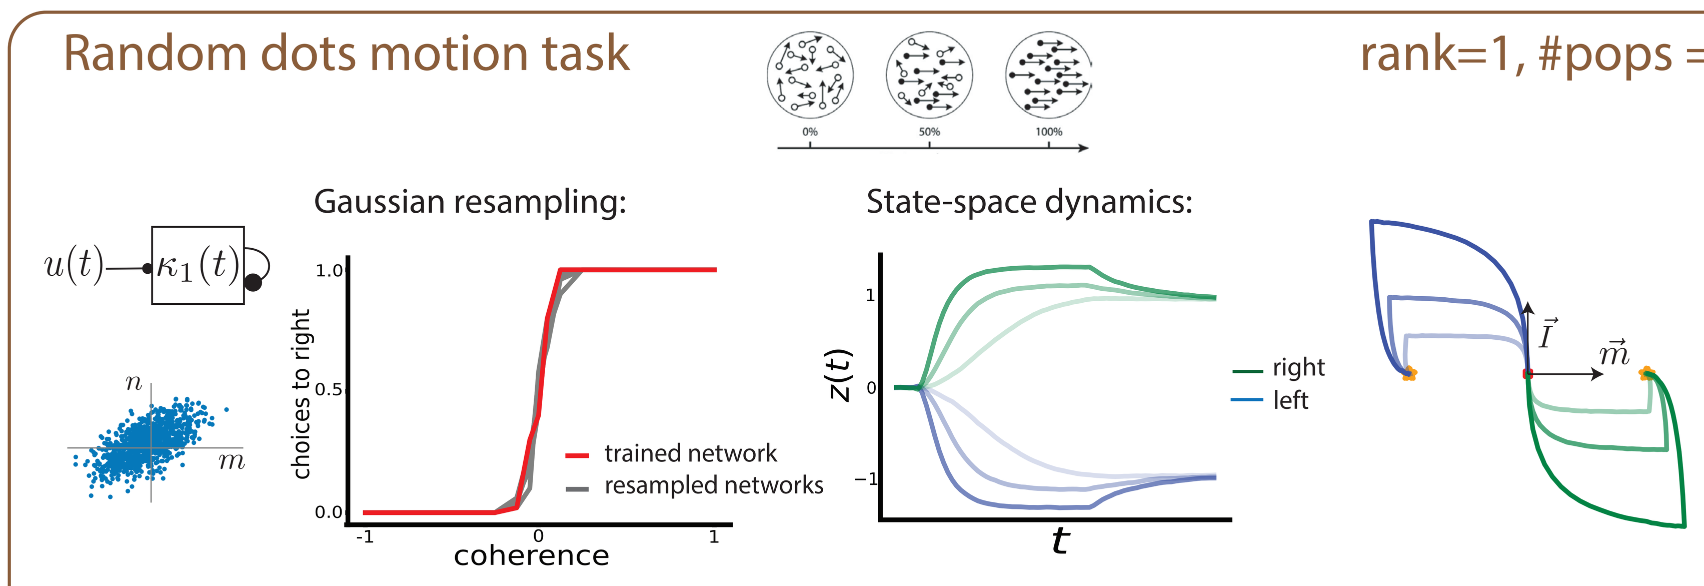

➤ **One internal variable : accumulation of evidence**

Delayed comparison task — rank=2, #pops=1

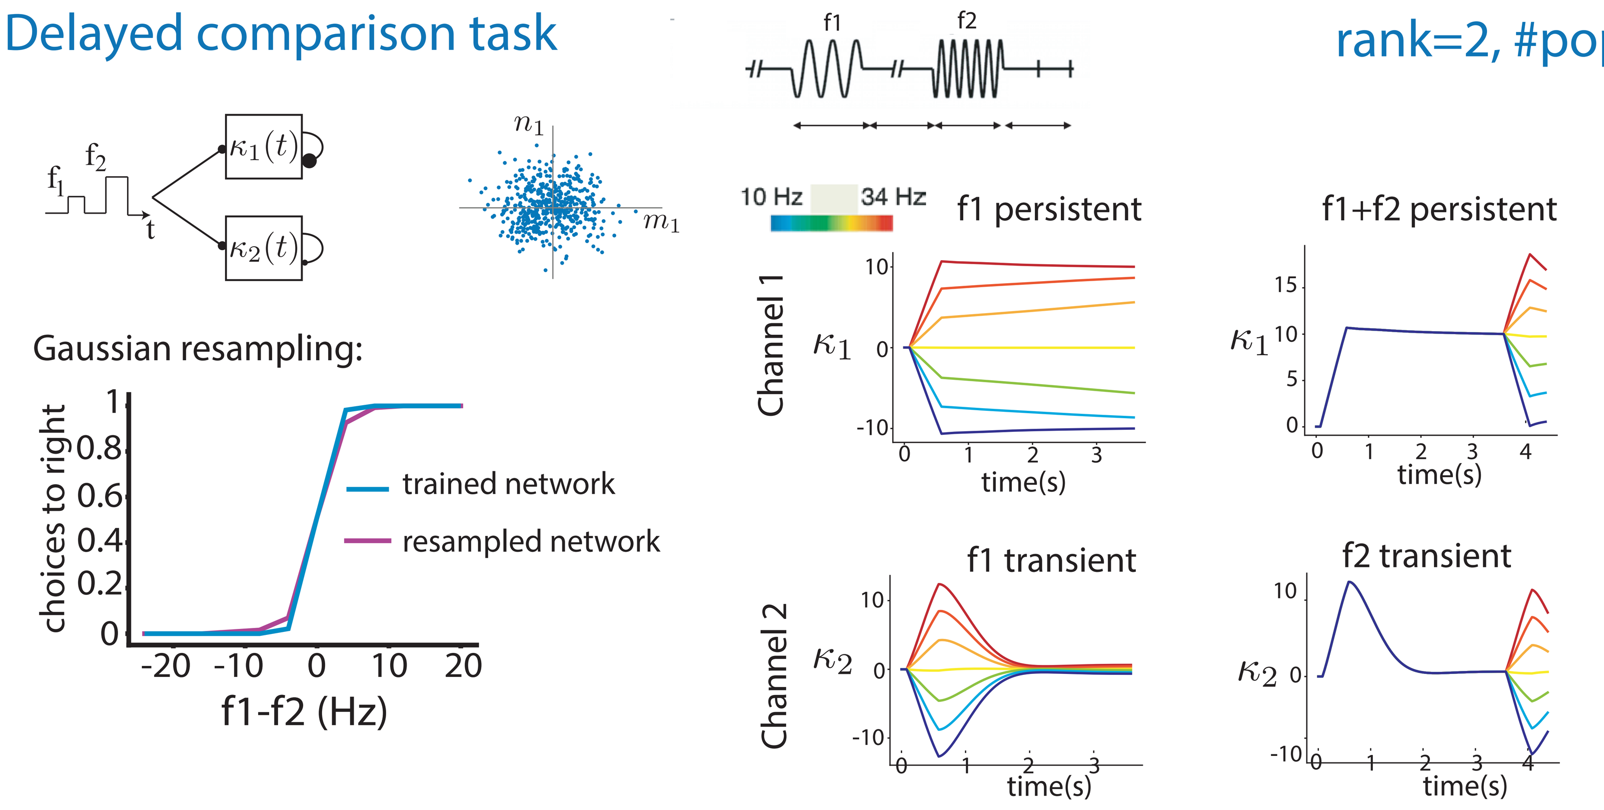

➤ **Two internal variables : memory + comparison**

Implementing multiple tasks

Computation in orthogonal spaces

$W_{rec} = \vec{m}_1^T \vec{n}_1 + \vec{m}_4^T \vec{n}_4 + \vec{m}_2^T \vec{n}_2 + \vec{m}_3^T \vec{n}_3$

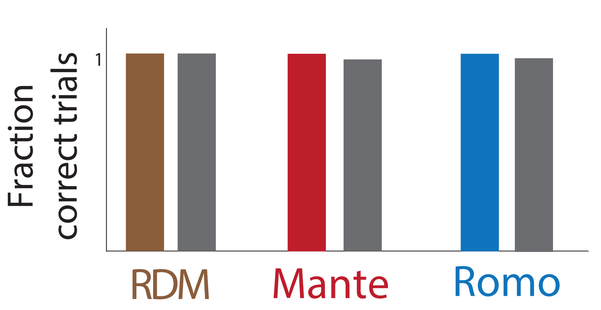

Computations with sub-populations

Mante tasks with P modalities

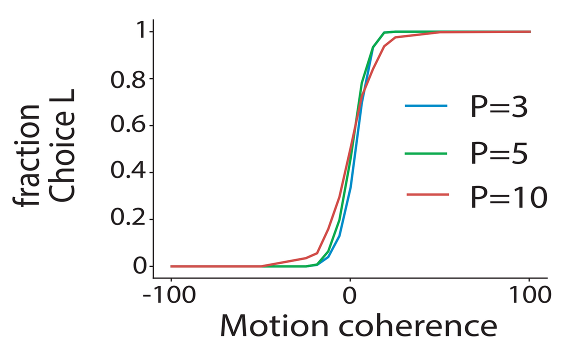

➤ **Low-rank implementations can be superposed**

Classes of networks

flexibility ↔ interpretability

Full-rank RNN: $\mathcal{O}(N^2)$ parameters

Trainable $\mathbf{W}_{rec}, \vec{I}, \vec{w}_{out}$

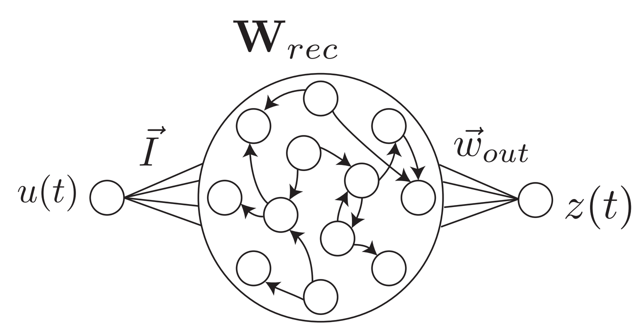

Low-rank RNN: $\mathcal{O}(N)$ parameters

$W_{rec} = \frac{1}{N} \sum_{k=1}^{K} \vec{m}^k \vec{n}^{kT}$

Trainable $\vec{m}^k, \vec{n}^k, \vec{I}, \vec{w}_{out}$

Interpretability: dynamics on K-dimensional subspace

$\vec{x} = \kappa_1 \vec{m}^1 + \cdots + \kappa_K \vec{m}_K$

The κ_i are internal variables

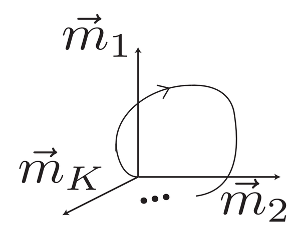

Multi-population networks: $\mathcal{O}(1)$ parameters

The $\vec{m}, \vec{n}, \vec{I}$ vectors sampled from joint mixture of gaussians distribution, e.g.

$$(m_i^1, n_i^1, I_i, w_{out}) \sim \sum_{p=1}^{P} \mathcal{N}(\vec{\mu}_p, \Sigma_p)$$

Trainable $(\vec{\mu}_p, \Sigma_p)$

Interpretability: analytically tractable dynamics (mean-field theory)

$$\dot{\kappa}_1 = -\kappa_1 + \tilde{\sigma}_{n_1 m_1} \kappa_1 + \tilde{\sigma}_{n_1 m_2} \kappa_2 + \tilde{\sigma}_{n_1 W_{in}} u(t)$$
$$\dot{\kappa}_2 = -\kappa_2 + \tilde{\sigma}_{n_2 m_1} \kappa_1 + \tilde{\sigma}_{n_2 m_2} \kappa_2 + \tilde{\sigma}_{n_2 W_{in}} u(t)$$

with functional connectivities:

$$\tilde{\sigma}_{ab} = \sum_{p=1}^{P} \sigma_{ab}^{p} \langle \phi' \rangle_p$$

Flexible tasks require multiple populations

Context-dependent integration — rank=1, #pops=2

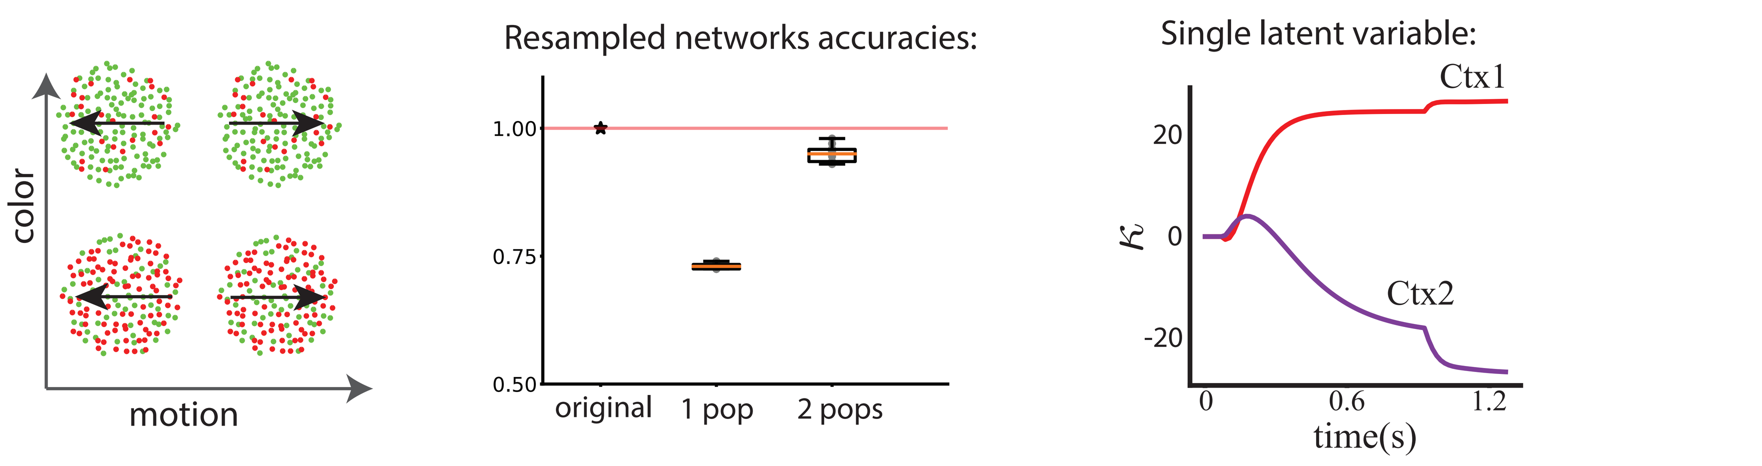

➤ **2 populations needed for switching**

Mechanism

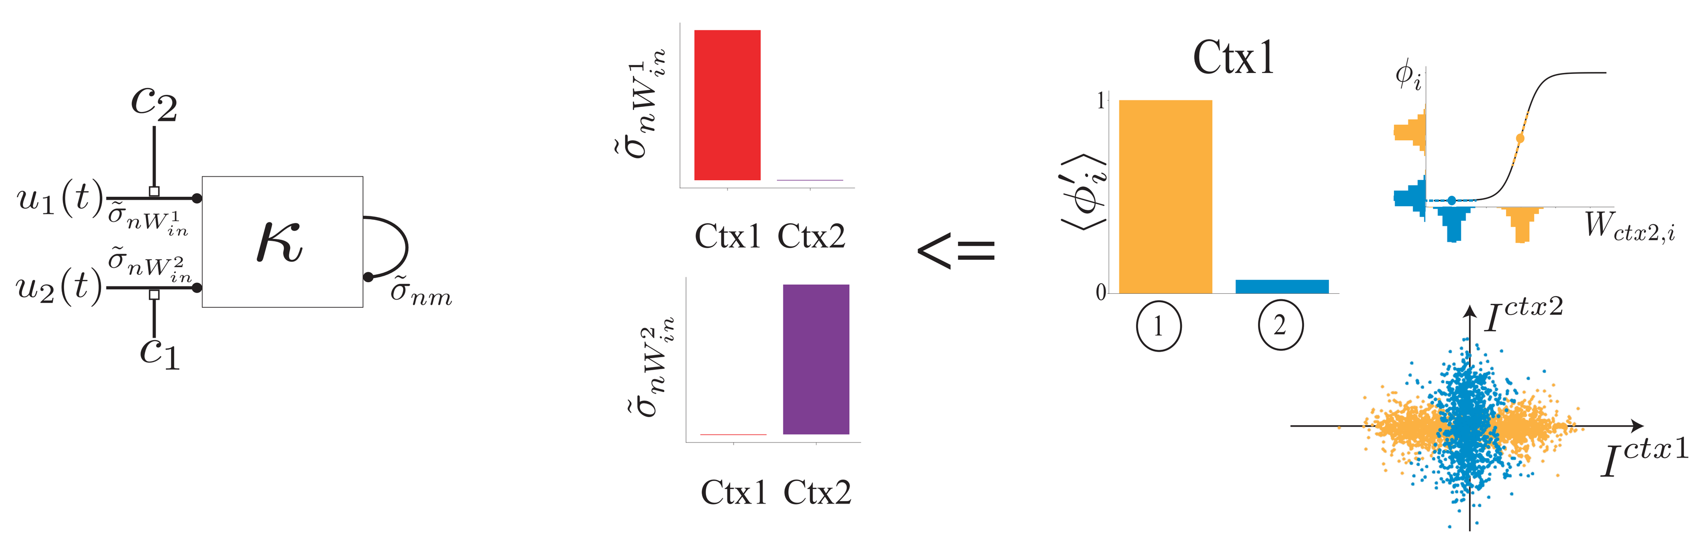

➤ **Nonlinear gating of integrated inputs**

Delayed Match-to-Sample — rank=2, #pops=2

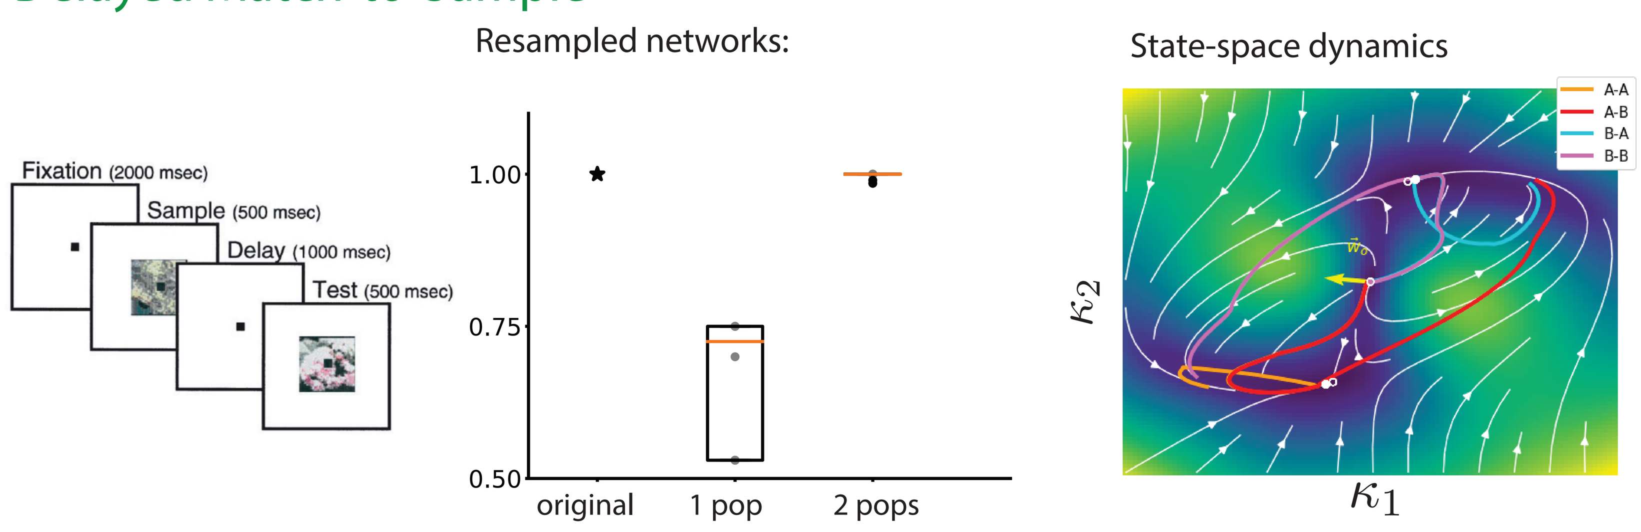

Mechanism

population-dependent control of the dynamics

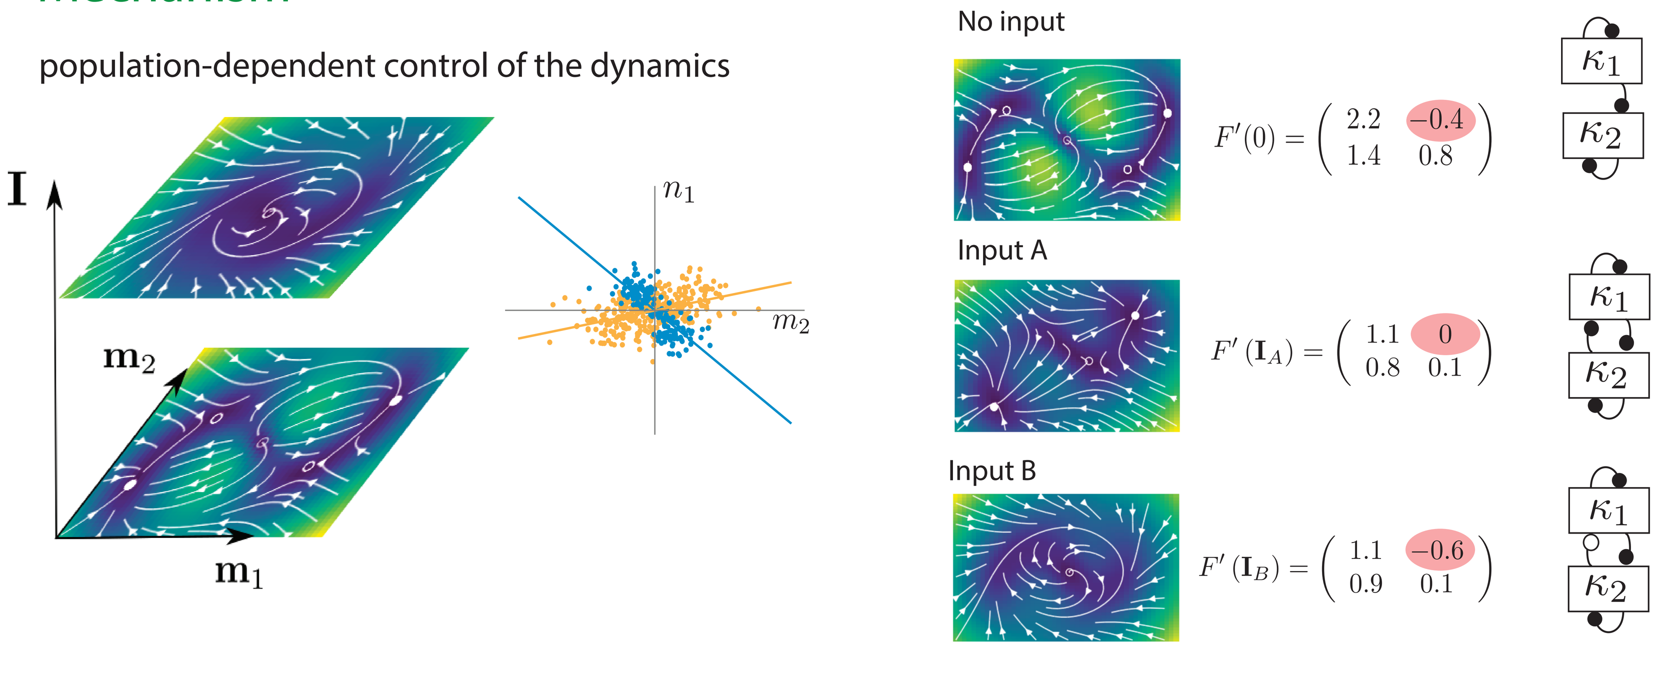

➤ **Nonlinear gating of recurrent interactions**

References

[1] F. Mastrogiuseppe, S. Ostojic, Linking connectivity, dynamics, and computations in low-rank recurrent neural networks, *Neuron*, 2018

[2] D. Sussillo, O. Barak, Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks, *Neural Comp.*, 2012

[3] Mante et al., Context-dependent computation by recurrent dynamics in prefrontal cortex, *Nature*, 2013

[4] Raposo et al., A category-free neural population supports evolving demands during decision-making, *Nature Neuroscience*, 2014

[5] Yang et al., Task representations in neural networks trained to perform many cognitive tasks, *Nature Neuroscience*, 2019