University of Zurich[UZH]

ETH zürich

Operative dimensions in unconstrained connectivity of recurrent neural networks

Presenter: Renate Krause
rekrau@ini.uzh.ch

Institute of Neuroinformatics
University of Zurich and ETH Zurich

What happens inside an RNN?

Unconstrained, vanilla RNN

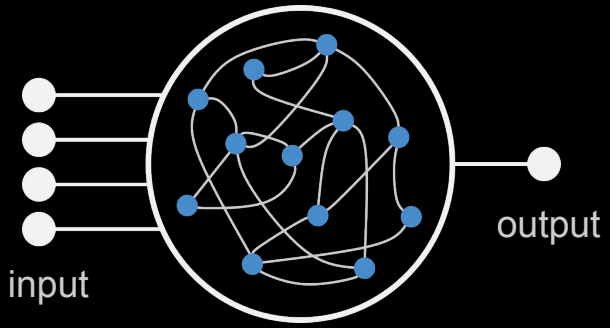

?

Network function

Network structure

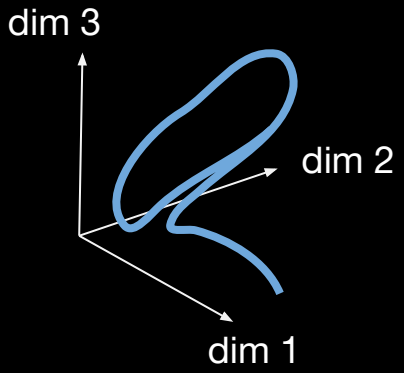

$$W = \begin{pmatrix} 4 & 2 & 2 \\ 2 & 6 & 1 \\ 8 & 1 & 4 \end{pmatrix}$$

Our main contribution: Operative dimensions as a novel tool to identify task-specific subspaces in recurrent weight matrices

Definition of operative dimensions

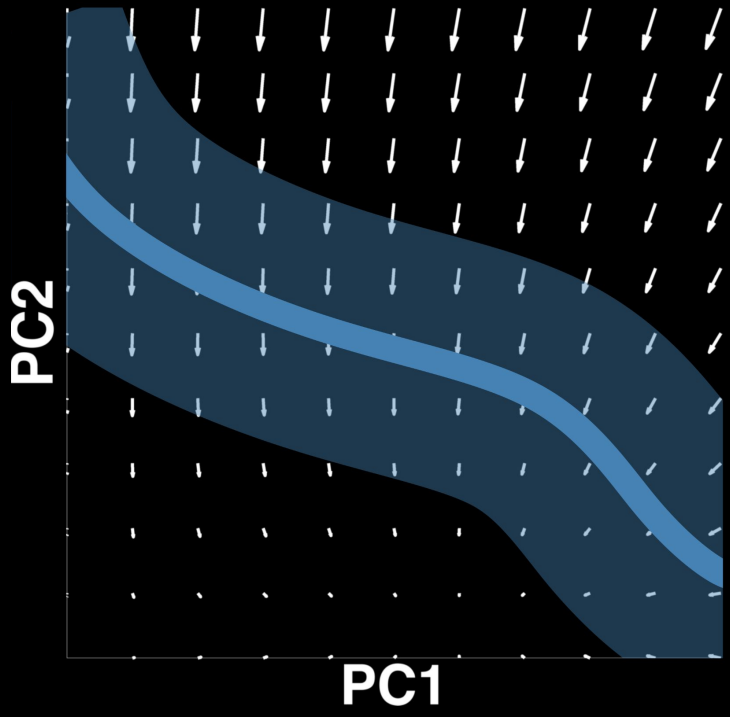

recurrent flow field

activity trajectory

Which dimension should be remove from W to change the relevant recurrent flow field the most?

$$\text{argmax}_{\text{dimension}} \; |\Delta \text{flow field}|$$

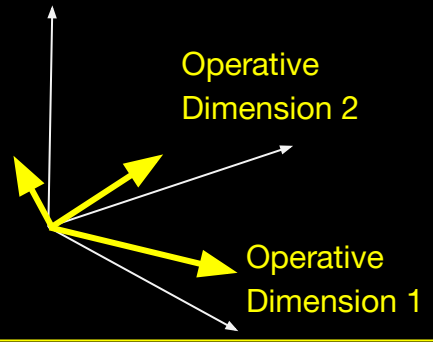

Operative dimensions identify functionally relevant subspace in recurrent weight matrix

⇒ Low-dimensional subspace is sufficient for original performance

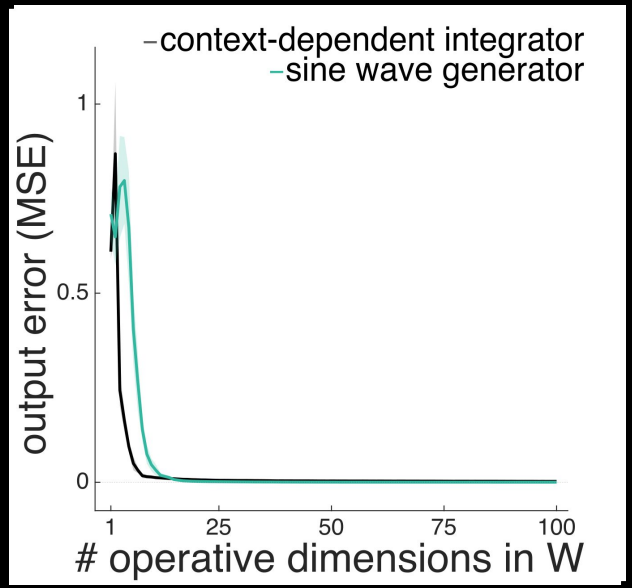

⇒ Functionally relevant subspace does <u>not</u> explain most of the variance in the weight matrix

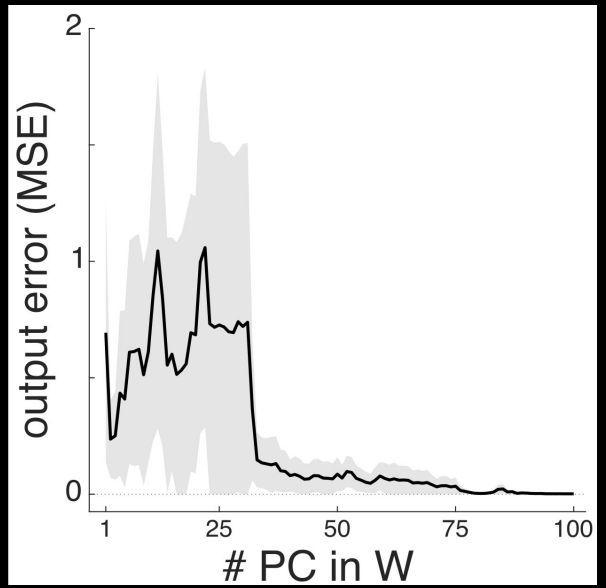

Thank you for your attention!

Contact:
rekrau@ini.uzh.ch
valerio@ini.uzh.ch

Acknowledgement

Valerio Mante &
NCC group

Matthew Cook &
CO2 group

Giacomo Indiveri
& NCS group

Institute of
Neuroinformatics
UZH/ETH Zurich