CONTRIBUTED TALK

Parameter symmetries determine representational geometry in overparameterized nonlinear networks

Marvin Theiss

Neural Information Processing Group, University of Tübingen

International Max Planck Research School for Intelligent Systems

PREVIEW

01 Parameter symmetries doubly dissociate function and representation in overparameterized nonlinear networks.

02 Representational geometry provably splits into task-linked structure and symmetry-induced noise.

03 Minimum-norm implementation rules recover a stable link between function and representation by identifying privileged geometries.

00

Setting, terminology, and motivation

Fully-connected layers in deep nonlinear networks

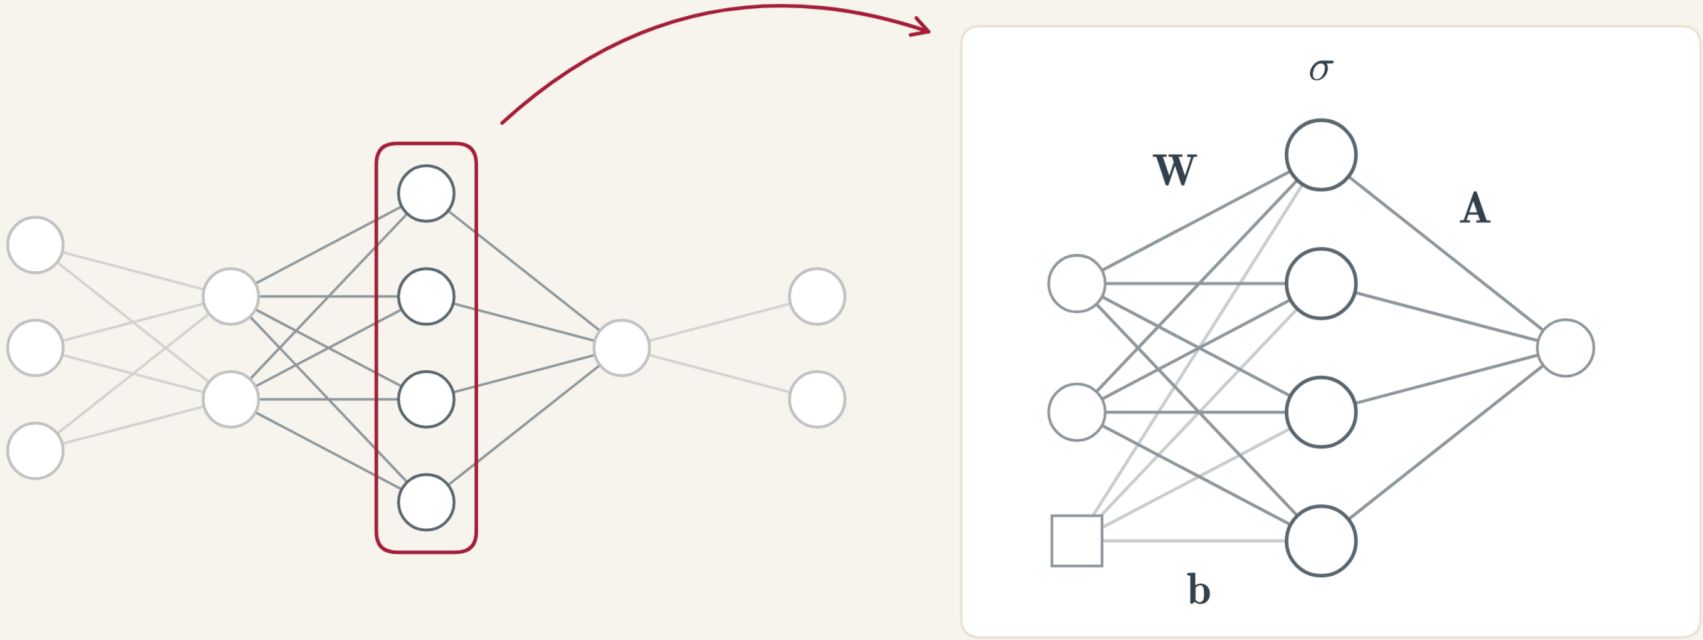

$$\boldsymbol{\theta} = (\sigma, \mathbf{W}, \mathbf{b}, \mathbf{A})$$

Function and representation

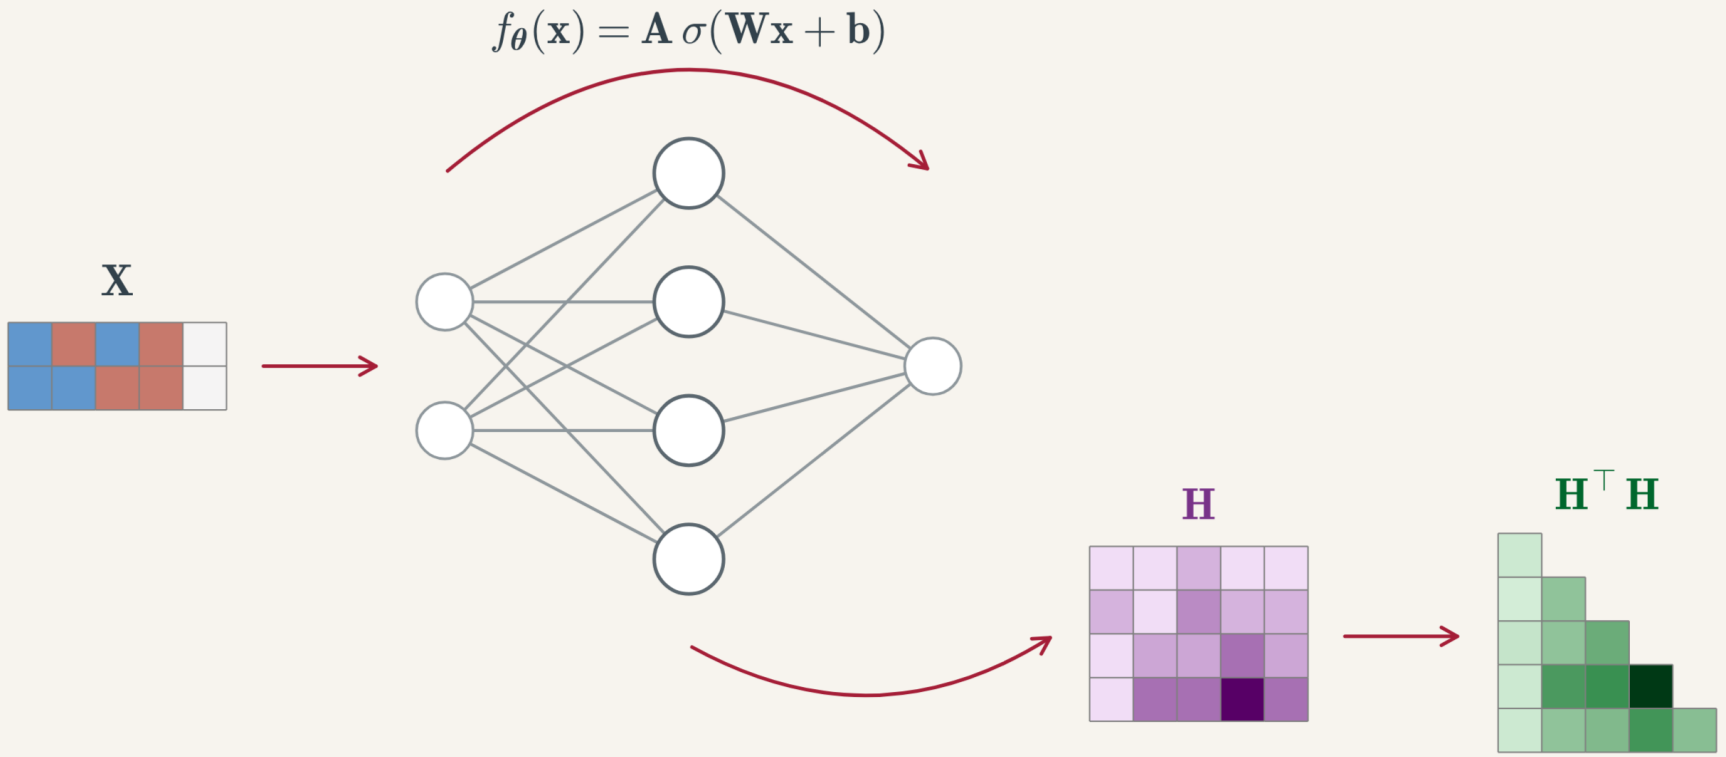

Functional parameter symmetries in deep nonlinear networks

Rescaling

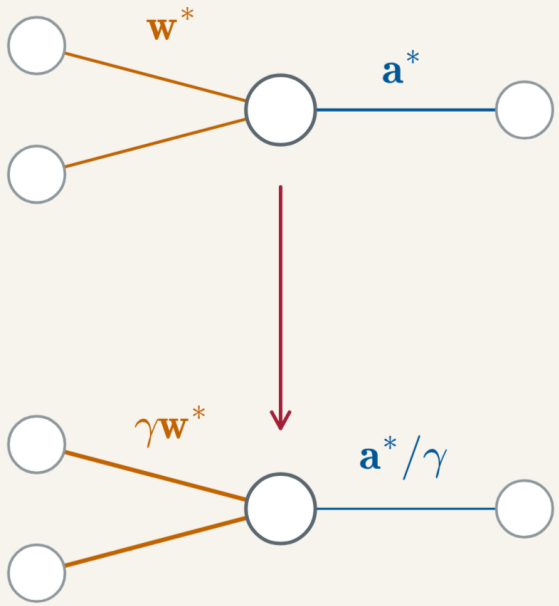

$\gamma > 0$

Duplicate group

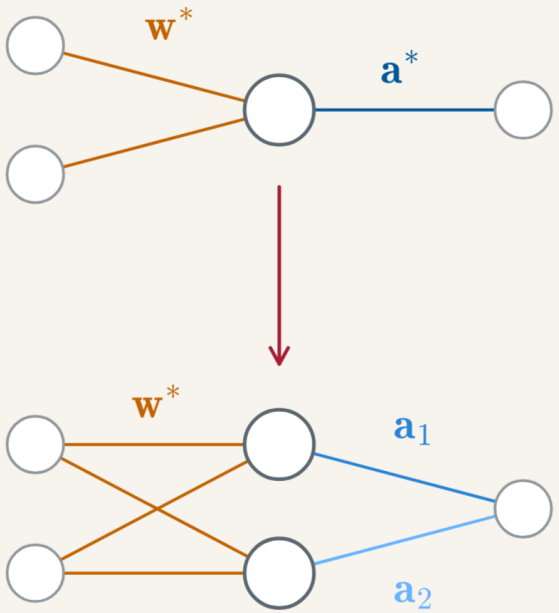

$\mathbf{a}_1 + \mathbf{a}_2 = \mathbf{a}^*$

Zero group

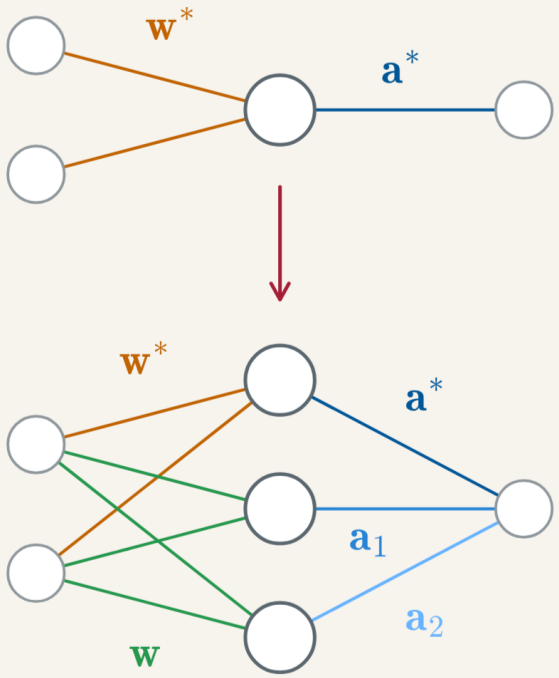

$\mathbf{a}_1 + \mathbf{a}_2 = \mathbf{0}$

Overparameterization, irreducibility, and orbit

Overparameterized

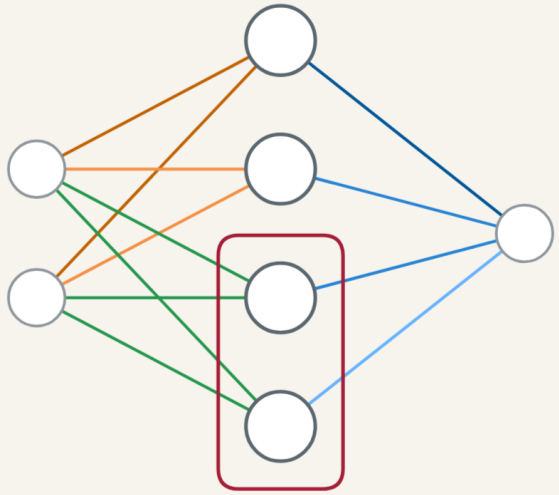

$\boldsymbol{\theta} = (\mathbf{W}, \mathbf{b}, \mathbf{A})$

Irreducible

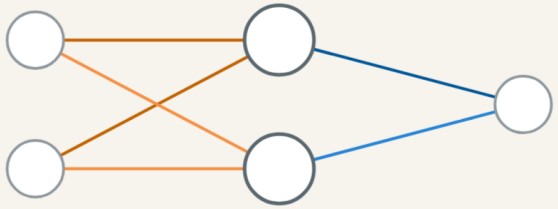

$\boldsymbol{\theta}^\star = (\mathbf{W}^\star, \mathbf{b}^\star, \mathbf{A}^\star)$

Orbit

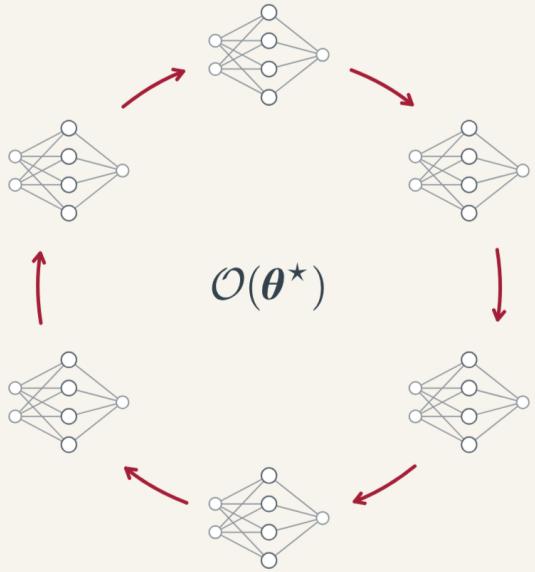

different $\boldsymbol{\theta}$, same $(f_{\boldsymbol{\theta}}, \boldsymbol{\theta}^\star)$

Studying representational geometry

- Neuroscience & cognitive science share **implicit assumption** that the geometry of neural activity reveals the computation it instantiates
- Yet, neural networks, whether artificial or biological, are **degenerate**: many distinct patterns of neural activity can support the same function
- Braun et al. (2025) demonstrate analytical **double dissociation** of function and representation in linear networks

FUNDAMENTAL QUESTION

Is representational geometry per se ill-suited to study computation in neural networks?

01

Function & representation doubly dissociate

A double dissociation of function & representation

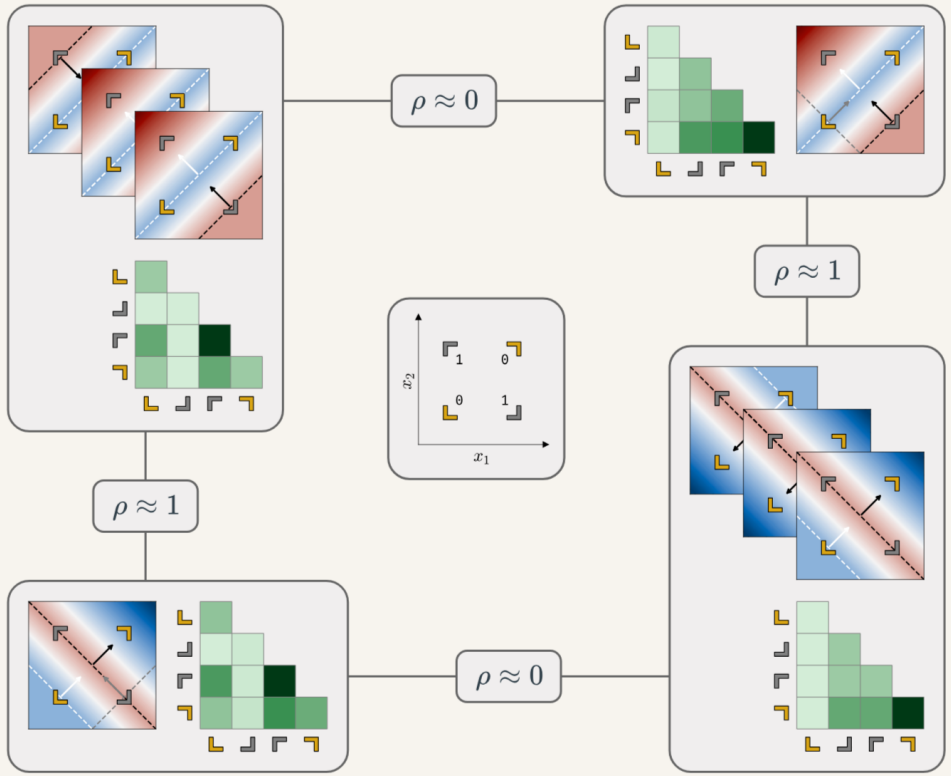

01

Parameter symmetries doubly dissociate function and representation in overparameterized nonlinear networks.

KEY TAKEAWAY · SECTION 01

02

From symmetries to feature transforms

Mapping parameter symmetries into hidden-activation space

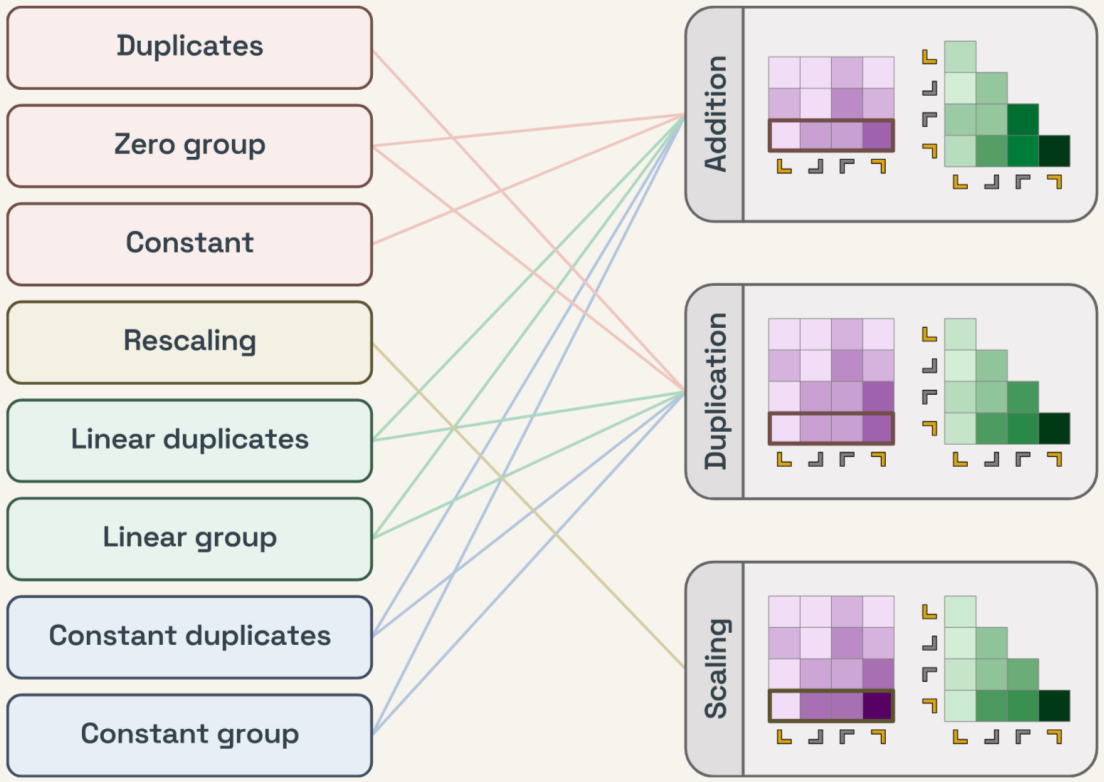

A canonical form of the hidden activation matrix

- Every parameter symmetry translates into a combination of at most two primitive feature transforms
- Multiple feature transforms performed in sequence commute …
- … and multiple transforms of the same type collapse to a single transform of that type!

A canonical form of the hidden activation matrix

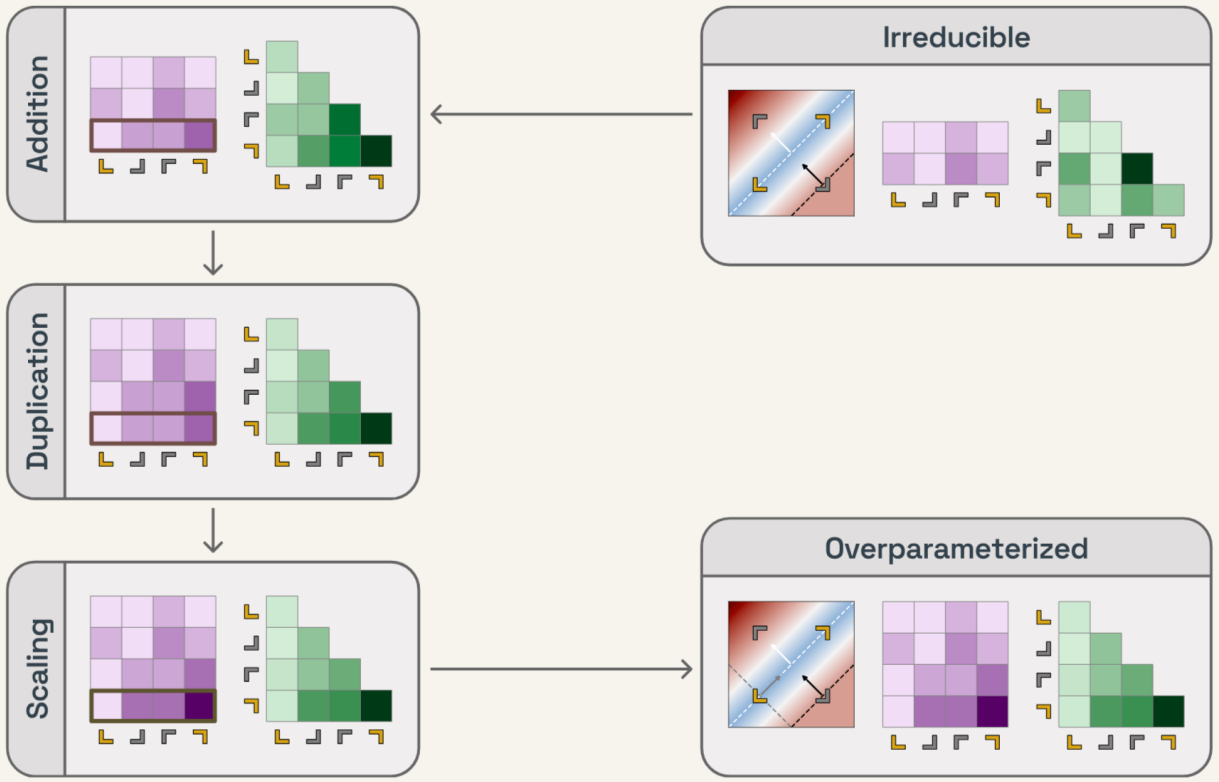

Task-linked structure and symmetry-induced noise

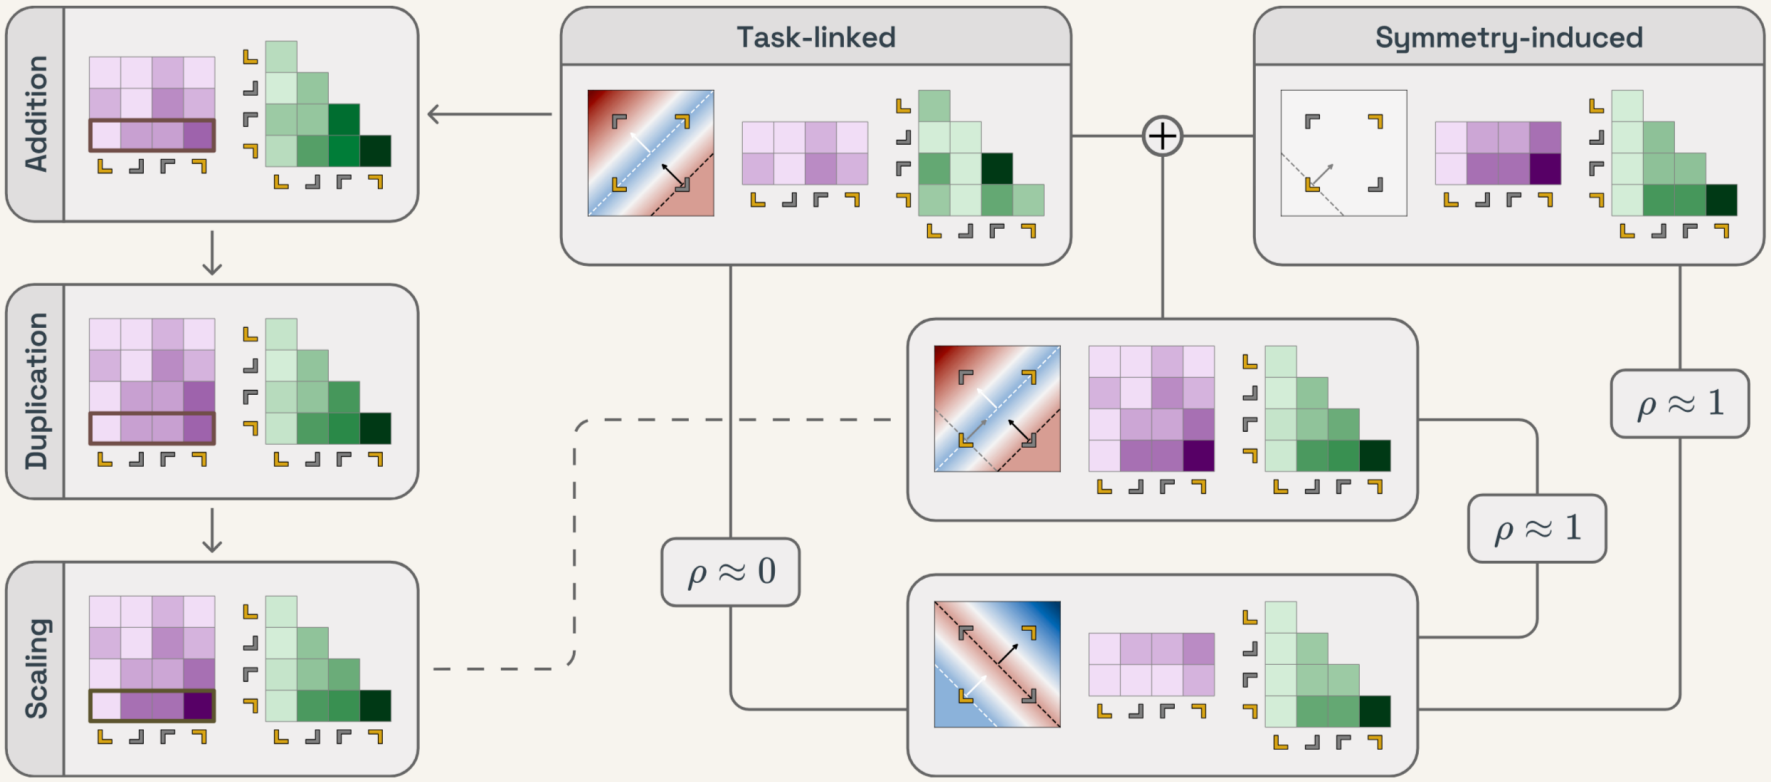

02

Representational geometry provably splits into task-linked structure and symmetry-induced noise.

KEY TAKEAWAY · SECTION 02

03

Minimum-norm implementation rules

Minimum representation-norm parameterizations

DEFINITION

A layer parameterization $\boldsymbol{\theta} = (\mathbf{W}, \mathbf{b}, \mathbf{A}) \in \mathcal{O}(\boldsymbol{\theta}^\star)$ is a **minimum representation-norm parameterization** if it minimizes

$$\|\mathbf{H}\|_F^2 + \|\mathbf{A}\|_F^2$$

over the orbit $\mathcal{O}(\boldsymbol{\theta}^\star)$.

Unique representational geometries

One function, many geometries

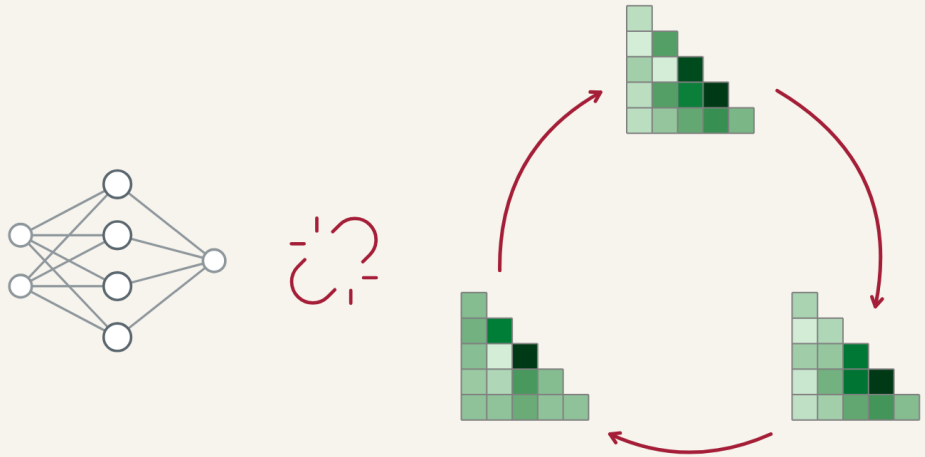

One function, one geometry

03

Minimum-norm implementation rules recover a stable link between function and representation by identifying privileged geometries.

KEY TAKEAWAY · SECTION 03

IN SUMMARY

01 Parameter symmetries doubly dissociate function and representation in overparameterized nonlinear networks.

02 Representational geometry provably splits into task-linked structure and symmetry-induced noise.

03 Minimum-norm implementation rules recover a stable link between function and representation by identifying privileged geometries.

thanks

Lukas Braun
Allen Institute
for Neural Dynamics

Andrew M. Saxe
Gatsby Unit & SWC, UCL

Erin Grant
University of Alberta & Amii

Felix A. Wichmann
University of Tübingen

FUNDING University of Tübingen · IMPRS-IS · Allen Institute · Gatsby Unit & Sainsbury Wellcome Centre, UCL · CIFAR · University of Alberta · Amii

Questions welcome!

mrvnthss.github.io ↗

PAPER

CODE

EBERHARD KARLS UNIVERSITÄT TÜBINGEN · imprs-is · GATSBY · wellcome · UCL · CIFAR · UNIVERSITY OF ALBERTA · amii · ICML International Conference On Machine Learning