Weight Space Representation Learning via Neural Field Adaptation

Zhuoqian Yang ▪ Mathieu Salzmann ▪ Sabine Süsstrunk

CVPR IN PARIS · CVPR JUNE 3-7, 2026 DENVER COLORADO

Take-Aways

1 Properly constrained weights exhibit **semantic structure** and serve as **effective data representations**;

2 Weight-space geometry **correlates** strongly with diffusion **generation quality**.

Project Page

1 Can Weights Serve as Representations?

▪ **A fitted neural field already encodes its data.** An image or shape lives entirely in its parameters.

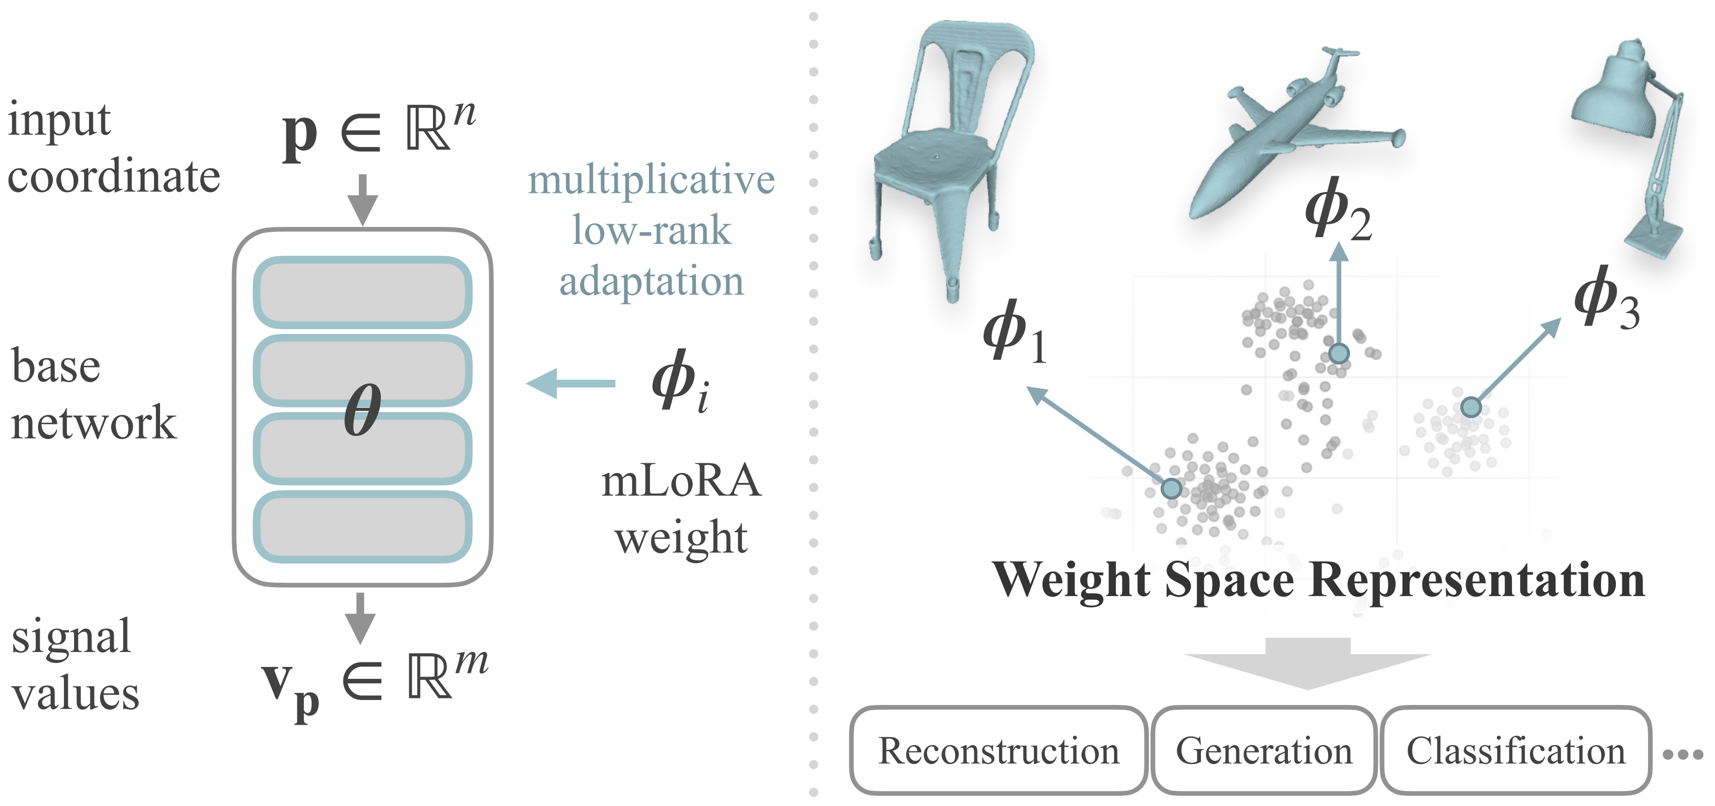

▪ **Representation Paradigm** A frozen base neural field is adapted by per-instance mLoRA weights $\boldsymbol{\phi}_i$; those weights *are* the data representation.

2 Challenge: Weight Space Symmetries

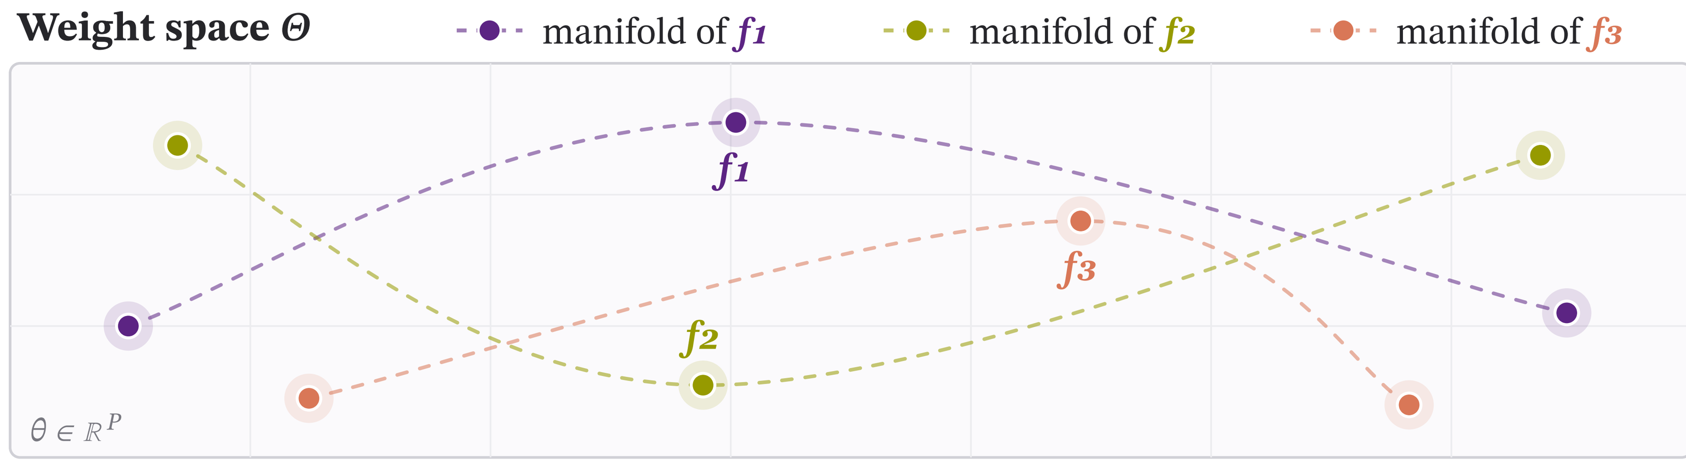

▪ **Symmetries complicate weight space structure.** Transformations on the network weights (such as permuting neurons) sends **functionally identical** networks **arbitrarily far** apart in weight space. The distribution becomes wildly **multi-modal**.

3 Enforcing Structure in Weight Space

▪ Each instance is a **multiplicative LoRA** (mLoRA) update over a shared base network: $\mathbf{W}' = \mathbf{W} \odot \mathbf{BA}$.

▪ We constrain the neural field optimization to **break symmetries**.

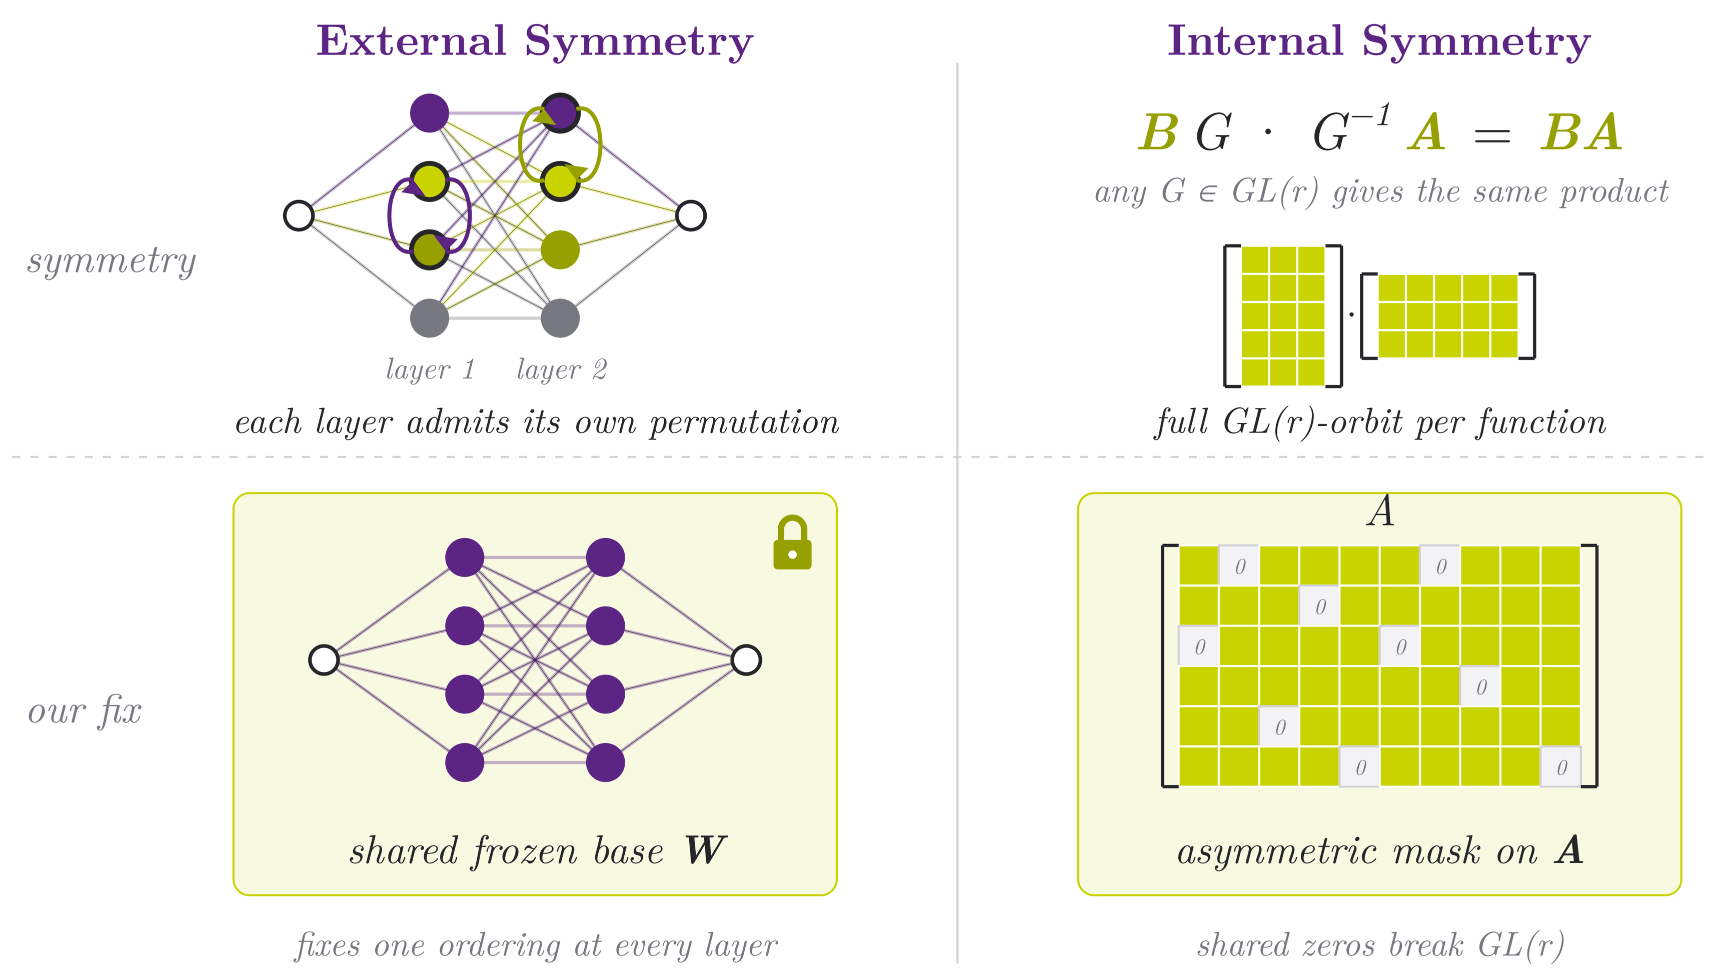

4 A Structured, Learnable Geometry

▪ Across very different initializations, mLoRA-Asym fits converge to the **same linear mode**: high weight similarity, near-zero loss barrier. This smooth geometry is what makes the weight space **learnable** (right).

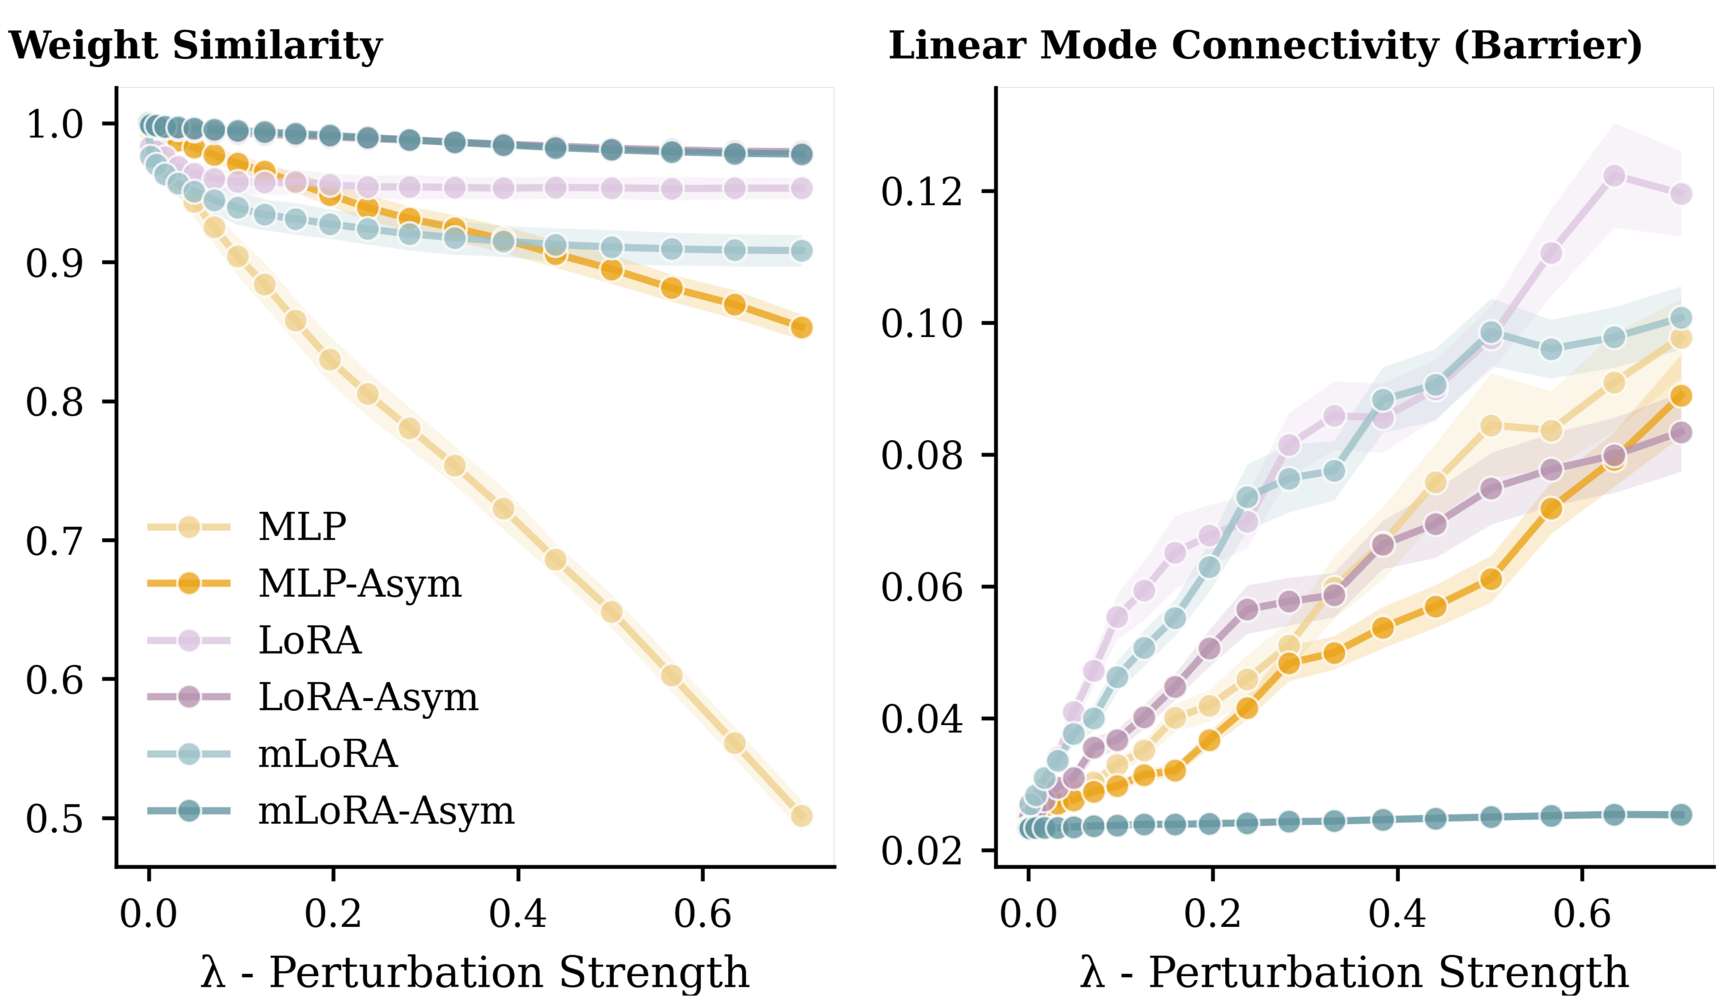

5 Generating New Neural Fields

▪ A **diffusion model** trained directly on the mLoRA weight set samples new weights, and therefore new neural fields: new images and new shapes.

First successful weight-space generation of **natural images** at 128×128.

6 Weights Encode Semantics

▪ A linear classifier on raw mLoRA weights reaches **90% accuracy** over ten ShapeNet categories. The weight space carries **semantic structure**.

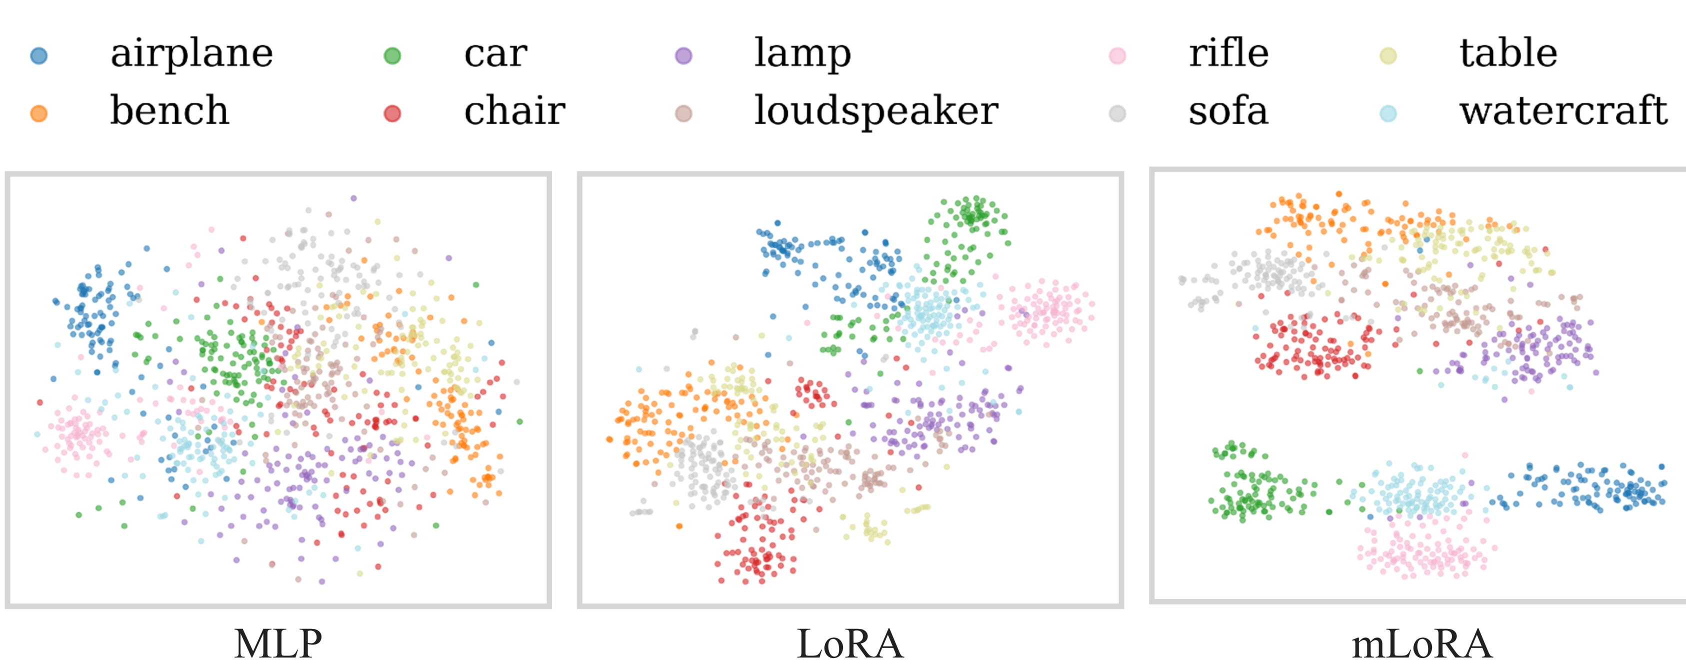

t-SNE of each weight space, colored by ShapeNet category.