
Beyond Structural Symmetries: Linear Mode Connectivity via Neuron Identifiability

Vincent Bürgin^{*1} Daniel Herbst^{*1} Ya-Wei Eileen Lin¹ Stefanie Jegelka^{1,2}

Abstract

Many striking phenomena in deep learning, such as linear mode connectivity and the structured behavior of training dynamics, are closely tied to parameter symmetries: transformations that leave the realized function unchanged. Despite growing attention to parameter symmetries, the exact interplay between parameters, data, and representations remains underexplored. To investigate this, we develop a theoretical framework of effective function classes, i.e., the set of functions a neuron can realize on its input support, and the norm cost of realizing them. We then formalize *effective symmetry breaking* via neuron identifiability across independent training runs. Our analysis shows that neural networks can admit large families of approximately equivalent solutions even in *structurally asymmetric* models. We further show that neuron identifiability enables representation merging *without prior alignment*, and characterize when such merging admits a linear low-loss path. These findings highlight the role of effective function classes in affecting the loss landscape.

1. Introduction

Modern neural networks are overparameterized and have highly nonconvex loss landscapes, yet independently trained models often achieve similar predictions and generalization (Goodfellow et al., 2015; Nguyen et al., 2019; Zhang et al., 2017; Li et al., 2018). One reason is that architectures admit large parameter symmetry groups (Hecht-Nielsen, 1990; Zhao et al., 2026), which affect optimization geometry and

^{*}Equal contribution ¹Technical University of Munich; School of CIT, MCML, MDSI ²MIT; Dept. of EECS, CSAIL. Correspondence to: Vincent Bürgin <vincent.buergin@tum.de>, Daniel Herbst <daniel.herbst@tum.de>. Code: github.com/vuenc/neuron-identifiability. This is a short version of the work with the same title at the 43rd International Conference on Machine Learning (ICML 2026).

Workshop on Weight-Space Symmetries, held in conjunction with the 43rd International Conference on Machine Learning, Seoul, South Korea, 2026. Copyright 2026 by the author(s).




Figure 1. Effect of neuron identifiability on the loss landscape.

model merging. The most ubiquitous are permutations of hidden units (Simsek et al., 2021; Zhao et al., 2022; 2024), which are central to Linear Mode Connectivity (LMC) (Frankle et al., 2020): independently trained networks can often be linearly interpolated with little loss degradation only after accounting for hidden-unit permutations (Entezari et al., 2022; Singh & Jaggi, 2020; Ainsworth et al., 2023).

Recent work has shown that breaking structural symmetries can yield unaligned LMC (Lim et al., 2024b; Ziyin et al., 2025). However, this outcome is not universal. Since both data and representations are often close to low-rank (Ansuini et al., 2019), transformations may remain approximate, data-dependent symmetries on actual inputs to a layer while differing substantially elsewhere. In this work, we ask:

When does a given parameter symmetry breaking mechanism select a consistent assignment of features to neurons across runs, and how do the data and representation distribution control its effectiveness?

Our Contributions. We develop a theoretical framework for *neuron identifiability*, i.e., the consistent assignment of features to neurons across random training seeds. In each layer, we view neurons functionally on their input support. Given a symmetry-breaking mechanism and a regularization budget, we characterize their corresponding *effective function classes*. This perspective shows that effective symmetry breaking is controlled by which architectural perturbations are observable on the input, and explains when hidden-layer representations can be merged without alignment, leading to unaligned LMC. Overall, we show that symmetry breaking is not binary, but governed by the interaction between architecture, data geometry, and effective function classes.

2. Preliminaries

Permutation Symmetries. Consider a 2-layer MLP $f_{\theta}(x) := \mathbf{W}_2 \eta(\mathbf{W}_1 x)$, $\theta = (\mathbf{W}_2, \mathbf{W}_1)$, where η is an elementwise nonlinearity. For any permutation $\pi \in S_m$ with permutation matrix $\mathbf{P}_{\pi} \in \{0, 1\}^{m \times m}$, the mapping $(\mathbf{W}_2, \mathbf{W}_1) \mapsto (\mathbf{W}_2 \mathbf{P}_{\pi}^{\top}, \mathbf{P}_{\pi} \mathbf{W}_1)$ is a parameter symmetry. In other words, hidden units are interchangeable, i.e., permuting them and compensating in the adjacent layer does not change the computed function. Depending on η , this group can also be significantly larger (e.g., $\text{GL}(m)$ for linear η). Similar symmetries appear in modern models, e.g., channels in CNNs, attention heads, MoEs, or through any computation graph automorphism (Lim et al., 2024a;b).

Symmetry Breaking. Symmetry breaking refers to any modification of a network’s forward pass that reduces the effective parameter symmetry group so that fewer distinct parameters represent the same function. In this work, we study an architectural intervention that replaces each trainable weight matrix \mathbf{W} with an effective matrix

$$\mathbf{W}_{\text{eff}} = \mathbf{F} + \mathbf{D} \odot \mathbf{W}, \quad (1)$$

where $\mathbf{F}, \mathbf{D} \in \mathbb{R}^{m \times d}$ are considered fixed parts of the architecture, \odot denotes elementwise multiplication, and $\mathbf{W} \in \mathbb{R}^{m \times d}$ contains the trainable parameters that are randomly initialized independently across runs and then optimized. $\mathbf{F}, \mathbf{D} \in \mathbb{R}^{m \times d}$ can be fixed upfront or sampled once per architecture. As summarized in Table 1, this subsumes several existing parameter symmetry breaking schemes.

Table 1. Symmetry breaking schemes covered by (1).

	\mathbf{F}	\mathbf{D}
\mathbf{W} -asym. (Lim et al., 2024b)	$\mathbf{M} \odot \mathbf{B}$ where $\mathbf{B} \sim \mathcal{N}(0, \sigma_{\mathbf{B}}^2)$ i.i.d., $\mathbf{M} \in \{0, 1\}^{m \times d}$	$\mathbf{1}\mathbf{1}^{\top} - \mathbf{M}$
syre (Ziyin et al., 2025)	$\mathcal{N}(0, \sigma_{\mathbf{F}}^2)$ i.i.d.	$\text{Unif}(1-\varepsilon, 1+\varepsilon)^{-1/2}$ i.i.d.
Linear residual	\mathbf{I}_d	$\mathbf{1}\mathbf{1}^{\top}$
Sparse network	0	$\mathbf{M} \in \{0, 1\}^{m \times d}$

For \mathbf{W} -asymmetric networks (Lim et al., 2024b), a binary mask \mathbf{M} selects a subset of coordinates that $\mathbf{F} := \mathbf{M} \odot \mathbf{B}$ fixes to constant random values, while unmasked coordinates remain trainable. syre (Ziyin et al., 2025) directly draws \mathbf{F} and \mathbf{D} i.i.d. Linear residual connections and sparse networks, both of which are known to break hidden-unit permutation symmetries (Lim et al., 2024b; Zhao et al., 2025), can also be interpreted in this framework.

Setup and Assumptions. Throughout most of the work, we will consider an asymmetric layer (1) in isolation. Concretely, we study the layer map $\mathbf{H} : \mathbb{R}^{m \times d} \times \mathbb{R}^d \rightarrow \mathbb{R}^m$,

$$\mathbf{H}(\mathbf{W}; \mathbf{x}) := \eta((\mathbf{F} + \mathbf{D} \odot \mathbf{W})\mathbf{x}). \quad (2)$$

Let $\mathbf{f}_i, \mathbf{d}_i$, and \mathbf{w}_i be the i -th rows of \mathbf{F}, \mathbf{D} , and \mathbf{W} . The layer’s i -th neuron in (2) can be written as $\mathbf{h}_i(\mathbf{w}_i; \mathbf{x}) := \eta((\mathbf{f}_i + \mathbf{d}_i \odot \mathbf{w}_i)^{\top} \mathbf{x})$. Let the input \mathbf{x} be supported on $\mathcal{X} \subset \mathbb{R}^d$ and drawn from a distribution \mathbb{P} on \mathcal{X} . Define $\mathcal{H}(\mathcal{X}) := \{\mathbf{H}(\mathbf{W}; \cdot) \mid \mathbf{W} \in \mathbb{R}^{m \times d}\}$ and analogously

$\mathcal{H}_i(\mathcal{X})$ as the set of functions a layer (or neuron) can implement. \mathcal{X} is typically far from filling the ambient space \mathbb{R}^d . For our analysis, we assume that \mathbb{P} is supported on a k -dim. subspace $\mathcal{U} \subset \mathbb{R}^d$ and fix an orthonormal basis $\mathbf{U} \in \mathbb{R}^{d \times k}$ of \mathcal{U} (§ C). While this is idealized (§ E.1), it enables exact projections onto \mathcal{U} and provides explicit coordinates via \mathbf{U} .

3. Neuron Identifiability

We now relate the effectiveness of a symmetry breaking intervention according to (1) to its ability of equipping neurons with distinguishable identities that lead to consistent assignments across training runs. Within one layer, each neuron can be viewed as implementing a particular feature. Across training runs, a layer often realizes essentially the same set of features, and runs differ mainly by a permutation that matches neurons. In a symmetric layer, this assignment is determined by the initialization. By *effective symmetry breaking* we mean that, instead, architectural asymmetries bias training towards a consistent assignment. We analyze this through each neuron’s *effective function class*, i.e., the functions it can realize on the input support together with their *realization costs*, measured by the minimum weight norm required to implement them. This is well-motivated under explicit weight decay and the implicit norm bias of gradient-based optimization. The underlying premise is that uniformly large reassignment costs over nontrivial permutations are equivalent to there being a unique minimum-complexity assignment of the feature set to neurons.

We additionally assume $\text{span}(\text{supp}(\mathbb{P})) = \mathcal{U}$ if η is injective (e.g., tanh, LeakyReLU). If $s \mapsto (\eta(s), \eta(-s))$ is injective (e.g., ReLU, GELU), assume there is $T \subseteq \text{supp}(\mathbb{P})$ s.t. $T = -T$ and $\text{span}(T) = \mathcal{U}$. This lets us write realizability of a feature directly in subspace coordinates: Setting $\mathbf{v}_i := \mathbf{U}^{\top} \mathbf{f}_i \in \mathbb{R}^k$, $\mathbf{M}_i := \mathbf{U}^{\top} \text{Diag}(\mathbf{d}_i) \in \mathbb{R}^{k \times d}$, $\mathbf{h} \in \mathcal{H}_i(\mathcal{X})$ can be written as $\mathbf{h}(\mathbf{x}) = \eta((\mathbf{U}\mathbf{a})^{\top} \mathbf{x})$, where

$$\mathbf{a} := \mathbf{U}^{\top} (\mathbf{f}_i + \text{Diag}(\mathbf{d}_i)\mathbf{w}) = \mathbf{v}_i + \mathbf{M}_i \mathbf{w} \in \mathbb{R}^k. \quad (3)$$

The map $\mathbf{h} \mapsto \mathbf{a}$ is a bijection that identifies $\mathcal{H}_i(\mathcal{X})$ with $\mathbf{v}_i + \text{im}(\mathbf{M}_i) \subseteq \mathbb{R}^k$ (cf. Prop. C.4). In particular, $\mathcal{H}_i(\mathcal{X})$ is an affine subspace of \mathbb{R}^k of dimension $\text{rk}(\mathbf{M}_i) \leq k$, which allows for a geometric comparison of neuron function classes. This implies that if all \mathbf{M}_i have full rank, all neurons can represent exactly the same functions, i.e., $\mathcal{H}_1(\mathcal{X}) = \dots = \mathcal{H}_m(\mathcal{X}) \cong \mathbb{R}^k$. We mainly focus on this case.

Realization Costs. Having characterized which functions a neuron can represent, we now quantify at which cost. We define the *realization cost* of $\mathbf{h} : \mathcal{X} \rightarrow \mathbb{R}$ w.r.t. $\mathcal{H}_i(\mathcal{X})$ as

$$\|\mathbf{h}\|_{\mathcal{H}_i(\mathcal{X})} := \inf \{\|\mathbf{w}\|_2 \mid \mathbf{h} = \mathbf{h}_i(\mathbf{w}; \cdot)\}, \quad (4)$$

with the convention $\|\mathbf{h}\|_{\mathcal{H}_i(\mathcal{X})} := +\infty$ for $\mathbf{h} \notin \mathcal{H}_i(\mathcal{X})$. Define $\|\cdot\|_{\mathcal{H}(\mathcal{X})}$ analogously w.r.t. the Frobenius norm $\|\cdot\|_F$.¹

¹ $\|\cdot\|_{\mathcal{H}(\mathcal{X})}$ is not a norm ($\mathbf{H}(0; \cdot)$ can be $\neq 0$ but has cost 0).

(4) can be computed in closed form as a Mahalanobis seminorm: Defining the Gram matrix $\mathbf{S}_i := \mathbf{M}_i \mathbf{M}_i^\top = \mathbf{U}^\top \text{Diag}(\mathbf{d}_i)^2 \mathbf{U}$, $\mathbf{h} = \eta((\mathbf{U}\mathbf{a})^\top \cdot)$ has realization cost

$$\|\eta((\mathbf{U}\mathbf{a})^\top \cdot)\|_{\mathcal{H}_i(\mathcal{X})} = \|\mathbf{a} - \mathbf{v}_i\|_{\mathbf{S}_i^\dagger}, \quad (5)$$

see [Lem. C.5](#). Here, $\|\mathbf{x}\|_{\mathbf{A}} := \sqrt{\mathbf{x}^\top \mathbf{A} \mathbf{x}}$ and \dagger indicates the pseudoinverse. This shows that the realization costs w.r.t. neuron i depend crucially on the anisotropy of its Gram matrix \mathbf{S}_i . If $\mathbf{S}_i \approx c \mathbf{I}_k$, the cost is near-Euclidean, whereas anisotropy makes some features much more expensive than others. In the case of the diagonal asymmetry \mathbf{d}_i being sampled i.i.d. with mean $\mu_{\mathbf{D}}$, $\mathbb{E}[\mathbf{S}_i] = \mu_{\mathbf{D}} \mathbf{I}_k$, and the deviation $\|\mathbf{S}_i - \mu_{\mathbf{D}} \mathbf{I}_k\|_{\text{op}}$ can be controlled by the *subspace coherence* ([Candes & Tao, 2005](#); [Candes & Recht, 2009](#))

$$\nu(\mathcal{U}) := \max_{\ell \in [d]} (\mathbf{U}\mathbf{U}^\top)_{\ell\ell} \in [k/d, 1]. \quad (6)$$

We make this precise in [Thm. C.6](#). Intuitively, $\nu(\mathcal{U})$ quantifies how much of the k -dim. variability of the input can be seen from any ambient coordinate. High coherence (≈ 1) makes the effect of coordinate masking on symmetry breaking pronounced; low coherence ($\approx k/d$) can diffuse it across ambient dimensions. This means that the impact of \mathbf{D} is not severe if the input is sufficiently incoherent, which is often the case in data or hidden representations. Hence, one should not automatically expect \mathbf{D} alone to effectively break symmetries. We now discuss the interplay of \mathbf{D} and \mathbf{F} .

Realization Cost Sensitivity. For a layer $\mathbf{H} : \mathbb{R}^d \rightarrow \mathbb{R}^m$, define a group action of $(\pi, \tau) \in S_m \times S_d$:

$$((\pi, \tau) \cdot \mathbf{H})_i := \mathbf{h}_{\pi(i)}(\mathbf{P}_\tau^\top \cdot), \quad (7)$$

and define the (square) realization cost sensitivity

$$\Delta_{\pi, \tau}(\mathbf{H}; \mathcal{X}) := \|(\pi, \tau) \cdot \mathbf{H}\|_{\mathcal{H}(\mathbf{P}_\tau \mathcal{X})}^2 - \|\mathbf{H}\|_{\mathcal{H}(\mathcal{X})}^2 \quad (8)$$

(we will often omit \mathcal{X} for brevity). It decomposes as

$$\Delta_{\pi, \tau}(\mathbf{H}; \mathcal{X}) = \underbrace{\Delta_{\pi, \text{id}}((\text{id}, \tau) \cdot \mathbf{H}; \mathbf{P}_\tau \mathcal{X})}_{=: \Delta_\pi^{\text{out}}} + \underbrace{\Delta_{\text{id}, \tau}(\mathbf{H}; \mathcal{X})}_{=: \Delta_\tau^{\text{in}}} \quad (9)$$

This allows us to separately consider the effects of permuting the inputs (via τ) and outputs (via π) of a layer, for which we will write Δ_τ^{in} , Δ_π^{out} respectively. In both cases, the sensitivity consists of a displacement term of $\mathbf{F}\mathbf{U}$ plus terms from the trainable weights and anisotropy \mathbf{S}_i (cf. [Lem. C.8](#)). We focus on the *center-dominated regime*, where the former dominates and $\mathbf{S}_i \approx \mu_{\mathbf{D}} \mathbf{I}_k$ are close to isotropic.

Neuron Reassignments (Δ_π^{out}). Assume that all \mathbf{S}_i have full rank, so all neurons can realize the same functions on \mathcal{X} . Then, $\pi \in S_m$ only changes which neuron pays the realization cost for which feature. In the center-dominated regime, the leading effect is the projected center displacement

$$\delta_\pi^{\text{out}} := (\mathbf{P}_\pi - \mathbf{I}_m) \mathbf{F}\mathbf{U} \in \mathbb{R}^{m \times k}. \quad (10)$$

When the Gram matrices are approximately isotropic, $\mathbf{S}_i \approx \mu_{\mathbf{D}} \mathbf{I}_k$, the realization cost sensitivity satisfies $\Delta_\pi^{\text{out}}(\mathbf{H}) \sim \mu_{\mathbf{D}}^{-1} \|\delta_\pi^{\text{out}}\|_F^2$ up to lower-order terms from the trainable part and \mathbf{S}_i (see [Thm. C.10](#)). Thus, the smallest neuron

reassignment cost is

$$\min_{\pi \neq \text{id}} \Delta_\pi^{\text{out}}(\mathbf{H}) \sim \mu_{\mathbf{D}}^{-1} \gamma_{\text{out}}^2, \quad \gamma_{\text{out}} := \min_{i \neq j} \|\delta_{(ij)}^{\text{out}}\|_F. \quad (11)$$

Large γ_{out} gives neurons stable identities, while small γ_{out} leaves approximate symmetries. For sampled \mathbf{F} with centered i.i.d. entries of std. / scale $\sigma_{\mathbf{F}}$,

$$\gamma_{\text{out}} = \Theta(\sigma_{\mathbf{F}} \sqrt{k} m^{-2/k}) \quad (12)$$

with high probability ([Thm. C.12](#)). Thus, the gap decays slowly with width when the intrinsic dim. k is moderate, but can collapse for low-dim. data or hidden representations.

Input Reindexings (Δ_τ^{in}). Input permutations $\tau \in S_d$ change how the fixed weights align with the input subspace. The corresponding projected center displacement is

$$\delta_\tau^{\text{in}} := \mathbf{F}(\mathbf{P}_\tau - \mathbf{I}_d) \mathbf{U} \in \mathbb{R}^{m \times k}. \quad (13)$$

In the same center-dominated and near-isotropic regime, $\Delta_\tau^{\text{in}}(\mathbf{H}) \sim \mu_{\mathbf{D}}^{-1} \|\delta_\tau^{\text{in}}\|_F^2$, again up to trainable part and Gram anisotropy effects ([Thm. C.13](#)). For transpositions (ab) ,

$$\min_{a \neq b} \Delta_{(ab)}^{\text{in}}(\mathbf{H}) \sim \mu_{\mathbf{D}}^{-1} \gamma_{\text{in}}^2, \quad \gamma_{\text{in}} := \min_{a \neq b} \|\delta_{(ab)}^{\text{out}}\|_F. \quad (14)$$

For i.i.d. centered subgaussian fixed weights,

$$\gamma_{\text{in}} = \Theta(\sigma_{\mathbf{F}} \sqrt{m}) \min_{a \neq b} \|\mathbf{U}_{b,\cdot} - \mathbf{U}_{a,\cdot}\|_2 \quad (15)$$

with high probability ([Thm. C.15](#)). $\|\mathbf{U}_{b,\cdot} - \mathbf{U}_{a,\cdot}\|_2$ measures how distinguishable a and b are on the input. If swapping acts almost trivially on \mathcal{X} , sensitivity is small. Since $\|\mathbf{U}_{a,\cdot} - \mathbf{U}_{b,\cdot}\|_2 \leq 2\sqrt{\nu(\mathcal{U})}$, low coherence can mitigate sensitivity; asymmetry matters only along input-visible directions.


4. Linear Mode Connectivity

We also analyze how intermediate representations evolve when we linearly interpolate between two independently trained models. LMC asks whether the loss stays small along this segment. A natural sufficient condition is that the network's hidden representations vary approximately affinely along the interpolation path; recent work has observed this co-occurs with LMC ([Zhou et al., 2023](#)). We therefore quantify the nonlinearity induced by a single layer along the path, and relate it to an upper bound on the loss barrier. Given $\mathbf{W}^A, \mathbf{W}^B \in \mathbb{R}^{m \times d}$, consider $\mathbf{W}(\lambda) := (1 - \lambda) \mathbf{W}^A + \lambda \mathbf{W}^B$ for $\lambda \in [0, 1]$ and measure nonlinearity along it by the chord deviation $\xi_{\mathbf{H}}(\lambda; \mathbf{x}) := \mathbf{H}(\mathbf{W}(\lambda); \mathbf{x}) - ((1 - \lambda) \mathbf{H}(\mathbf{W}^A; \mathbf{x}) + \lambda \mathbf{H}(\mathbf{W}^B; \mathbf{x}))$.


For $\eta = \text{ReLU}$, $\mathbf{x} \sim \mathcal{N}(0, \Sigma)$, in the center-dom. regime $\|\mathbf{d}_i \odot \mathbf{w}_i^X\|_{\Sigma} \leq \beta \|\mathbf{f}_i\|_{\Sigma}$, $X \in \{A, B\}$, $\beta < 1$, we find that

$$\sup_{\lambda \in [0, 1]} \|\xi_{\mathbf{H}}(\lambda; \cdot)\|_{L^2(\mathbb{P}; \mathbb{R}^m)} = \mathcal{O}(\beta^{3/2}) \|\mathbf{F}\Sigma^{1/2}\|_F, \quad (16)$$

i.e., the chord deviation goes to zero as $\beta \rightarrow 0$, and its scale depends on the total norm of the fixed part under the input covariance ([Thm. D.1](#)). With control over chord deviations, one can upper bound the loss along the linear path between two independently trained models (cf. [Prop. D.3](#)).



(a) \mathcal{W} -MLP, different σ_F at epoch 100 (b) \mathcal{W} -ResNet, different σ_F at epoch 100 (c) \mathcal{W} -asym. models over different epochs
 Figure 2. Act. matching objs. of optimal, identity, and random permutations w/ diff. σ_F , using post-norm, post-act. function values.



(a) ResNets on CIFAR-10, unaligned (b) ResNets on CIFAR-10, aligned
 Figure 3. Aligned and unaligned LMC interpolation.




Figure 4. LMC barriers on synthetic dataset, $\mathbf{F} = 0$, varying $\nu(\mathcal{U})$.

5. Experimental Results

Q1: Do symmetry-breaking architectures improve unaligned LMC? We compare standard, \mathcal{W} -asym., and *syre*-MLPs on MNIST, and the corresponding ResNets on CIFAR-10. We measure unaligned LMC and activation-matched LMC, using $\sigma_F = 0$ and center-dom. \mathcal{W} -asym. settings ($\sigma_F = 1$ for MLPs, 2 for ResNets).² Fig. 7 shows that large fixed weights yield unaligned LMC in MLPs and partly in ResNets, while standard and \mathcal{W} -asym. models with $\sigma_F = 0$ have high unaligned barriers. After alignment, all barriers are low, including for \mathcal{W} -asym. models with $\sigma_F = 0$ and *syre*-ResNets. Thus structural asymmetry alone is insufficient.

Q2: When do \mathcal{W} -asym. networks learn the same assignment across runs? We sweep σ_F and compare the activation-matching objective of the identity, the optimal permutation, and random permutations. Figs. 2a and 2b show that, once σ_F is large enough, the identity becomes nearly optimal while random permutations remain much worse; Fig. 2c shows the same trend over training. Hence sufficiently large fixed centers make neuron identities consistent across runs.

Q3: Do swap costs depend on fixed scale and intrinsic dimension? We estimate realization costs using (5) on the 90% PCA subspace and compute $\Delta_{(ij)}^{\text{out}}$. The theory predicts

²For asymmetric models, alignment is applied to effective weights $\mathbf{W}_{\text{eff}} = \mathbf{F} + \mathbf{D} \odot \mathbf{W}$.




Figure 5. $\Delta_{(ij)}^{\text{out}}$ on synthetic Gaussian mixture data w/ varying intrinsic dim. k , positive/near-zero/negative values.

$\Delta_{\pi}^{\text{out}} \sim \mu_{\mathbf{D}}^{-1} \|\delta_{\pi}^{\text{out}}\|_{\mathbf{F}}^2$ and $\gamma_{\text{out}} = \Theta(\sigma_F \sqrt{km}^{-2/k})$. On Gaussian mixtures with width $m = 64$ and $k \in \{2, 8, 32\}$, Fig. 5 shows many cheap swaps for $\sigma_F = 0$, but increasing separation for $\sigma_F = 1$ as k grows. The unaligned midpoint accuracy drops are 46.4, 15.7, 6.1 for $k = 2, 8, 32$, confirming that low intrinsic dim. leaves approx. swaps accessible.

Q4: How does subspace coherence affect symmetry breaking? With $\mathbf{F} = 0$, diagonal masks should matter mainly when coordinate directions are visible on the data subspace. We test this by embedding $a, b \sim \mathcal{N}(0, 1)$ into \mathbb{R}^{20} by copying each coordinate to $t \leq 10$ positions, giving $\nu(\mathcal{U}) = 1/t$, and train a one-hidden-layer \mathcal{W} -asym. MLP to fit $\text{ReLU}(a + b)$. Fig. 4 shows low barriers at high coherence and high barriers at low coherence, also under a random-frame variant. Thus, diagonal asymmetry alone is effective mainly along data-visible coordinate directions.

6. Conclusion

We studied when parameter symmetry breaking yields neuron identifiability, i.e., a consistent assignment of features to neurons. This explains when asymmetry can be present in parameter space yet ineffective on data. We also showed that effective symmetry breaking can yield unaligned LMC. We leave extensions to richer symmetry groups, exact training dynamics, and more practical diagnostics for future work.

Acknowledgements

We thank the anonymous reviewers for their insightful feedback, and Valerie Engelmayer for providing feedback on an earlier version of this manuscript. VB acknowledges support from a PhD fellowship of the Munich Center for Machine Learning (MCML). DH acknowledges support from the Munich Data Science Institute (MDSI) via the MDSI Doctoral Fellowship. YEL acknowledges support from the Schmidt Futures Israeli Women’s Postdoctoral Award. This research was supported by an Alexander von Humboldt Professorship.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

- Adilova, L., Andriushchenko, M., Kamp, M., Fischer, A., and Jaggi, M. Layer-wise linear mode connectivity. In *The Twelfth International Conference on Learning Representations*, 2024.
- Ainsworth, S., Hayase, J., and Srinivasa, S. Git re-basin: Merging models modulo permutation symmetries. In *The Eleventh International Conference on Learning Representations*, 2023.
- Ansuini, A., Laio, A., Macke, J. H., and Zoccolan, D. Intrinsic dimension of data representations in deep neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Bar, N. and Giryes, R. Pruning at initialization—a sketching perspective. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- Candes, E. J. and Recht, B. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9:717–772, 2009.
- Candes, E. J. and Tao, T. Decoding by linear programming. *IEEE transactions on information theory*, 51(12):4203–4215, 2005.
- Chen, A. M., Lu, H.-m., and Hecht-Nielsen, R. On the geometry of feedforward neural network error surfaces. *Neural computation*, 5(6):910–927, 1993.
- Crisostomi, D., Fumero, M., Baieri, D., Bernard, F., and Rodolà, E. C^2M^3 : Cycle-consistent multi-model merging. In *Advances in Neural Information Processing Systems*, 2024.
- Dai, Z., Karzand, M., and Srebro, N. Representation costs of linear neural networks: Analysis and design. *Advances in Neural Information Processing Systems*, 34, 2021.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houselby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- Draxler, F., Veschgini, K., Salmhofer, M., and Hamprecht, F. Essentially no barriers in neural network energy landscape. In *International Conference on Machine Learning*, pp. 1309–1318. PMLR, 2018.
- Drineas, P., Magdon-Ismail, M., Mahoney, M. W., and Woodruff, D. P. Fast approximation of matrix coherence and statistical leverage. *Journal of Machine Learning Research*, 13:3475–3506, 2012.
- Entezari, R., Sedghi, H., Saukh, O., and Neyshabur, B. The role of permutation invariance in linear mode connectivity of neural networks. In *International Conference on Learning Representations*, 2022.
- Evcı, U., Gale, T., Menick, J., Castro, P. S., and Elsen, E. Rigging the lottery: Making all tickets winners. In *International Conference on Machine Learning*, pp. 2943–2952. PMLR, 2020.
- Ferbach, D., Goujaud, B., Gidel, G., and Dieuleveut, A. Proving linear mode connectivity of neural networks via optimal transport. In *International Conference on Artificial Intelligence and Statistics*, pp. 3853–3861. PMLR, 2024.
- Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019.
- Frankle, J., Dziugaite, G. K., Roy, D. M., and Carbin, M. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, 2020.
- Frankle, J., Dziugaite, G. K., Roy, D. M., and Carbin, M. Pruning neural networks at initialization: Why are we missing the mark? In *The Ninth International Conference on Learning Representations*, 2021.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pp. 1050–1059. PMLR, 2016.
- Garipov, T., Izmailov, P., Podoprikin, D., Vetrov, D. P., and Wilson, A. G. Loss surfaces, mode connectivity, and fast ensembling of DNNs. *Advances in Neural Information Processing Systems*, 31, 2018.

- Goodfellow, I. J., Vinyals, O., and Saxe, A. M. Qualitatively characterizing neural network optimization problems. *The Third International Conference on Learning Representations*, 2015.
- Grigsby, E., Lindsey, K., and Rolnick, D. Hidden symmetries of relu networks. In *International Conference on Machine Learning*, pp. 11734–11760. PMLR, 2023.
- Gunasekar, S., Lee, J. D., Soudry, D., and Srebro, N. Implicit bias of gradient descent on linear convolutional networks. *Advances in Neural Information Processing Systems*, 31, 2018.
- Hecht-Nielsen, R. On the algebraic structure of feedforward network weight spaces. In *Advanced Neural Computers*, pp. 129–135. Elsevier, 1990.
- Imfeld, M., Graldi, J., Giordano, M., Hofmann, T., Anagnostidis, S., and Singh, S. P. Transformer fusion with optimal transport. In *The Twelfth International Conference on Learning Representations*, 2024.
- Ito, A., Yamada, M., Chijiwa, D., and Kumagai, A. Do we really need permutations? Impact of model width on linear mode connectivity. In *The Fourteenth International Conference on Learning Representations*, 2026.
- Lee, N., Ajanthan, T., and Torr, P. H. S. SNIP: Single-shot network pruning based on connection sensitivity. In *The Seventh International Conference on Learning Representations*, 2019.
- Li, H., Xu, Z., Taylor, G., Studer, C., and Goldstein, T. Visualizing the loss landscape of neural nets. *Advances in Neural Information Processing Systems*, 31, 2018.
- Li, Y., Yosinski, J., Clune, J., Lipson, H., and Hopcroft, J. Convergent learning: Do different neural networks learn the same representations? In Storcheus, D., Ros-tamizadeh, A., and Kumar, S. (eds.), *Proceedings of the 1st International Workshop on Feature Extraction: Modern Questions and Challenges at NIPS 2015*, volume 44 of *Proceedings of Machine Learning Research*, pp. 196–212, Montreal, Canada, 11 Dec 2015. PMLR.
- Lim, D., Maron, H., Law, M. T., Lorraine, J., and Lucas, J. Graph metanetworks for processing diverse neural architectures. In *The Twelfth International Conference on Learning Representations*, 2024a.
- Lim, D., Putterman, T., Walters, R., Maron, H., and Jegelka, S. The empirical impact of neural parameter symmetries, or lack thereof. *Advances in Neural Information Processing Systems*, 37:28322–28358, 2024b.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- Nguyen, Q., Mukkamala, M. C., and Hein, M. On the loss landscape of a class of deep neural networks with no bad local valleys. In *International Conference on Learning Representations*, 2019.
- Pittorino, F., Ferraro, A., Perugini, G., Feinauer, C., Baldassi, C., and Zecchina, R. Deep networks on toroids: Removing symmetries reveals the structure of flat regions in the landscape geometry. In *ICML*, pp. 17759–17781, 2022.
- Rudelson, M. and Vershynin, R. Small ball probabilities for linear images of high-dimensional distributions. *International Mathematics Research Notices*, 19:9594–9617, 2015.
- Savarese, P., Evron, I., Soudry, D., and Srebro, N. How do infinite width bounded norm networks look in function space? In *The Thirty Second Annual Conference on Learning Theory*. PMLR, 2019.
- Simsek, B., Ged, F., Jacot, A., Spadaro, F., Hongler, C., Gerstner, W., and Brea, J. Geometry of the loss landscape in overparameterized neural networks: Symmetries and invariances. In *ICML*, pp. 9722–9732, 2021.
- Singh, S. P. and Jaggi, M. Model fusion via optimal transport. *Advances in Neural Information Processing Systems*, 33:22045–22055, 2020.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Tanaka, H., Kunin, D., Yamins, D. L. K., and Ganguli, S. Pruning neural networks without any data by iteratively conserving synaptic flow. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Tatro, N., Chen, P.-Y., Das, P., Melnyk, I., Sattigeri, P., and Lai, R. Optimizing mode connectivity via neuron alignment. *Advances in Neural Information Processing Systems*, 33:15300–15311, 2020.
- Theus, A., Cabodi, A., Anagnostidis, S., Orvieto, A., Singh, S. P., and Boeva, V. Generalized linear mode connectivity for transformers. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
- Tropp, J. A. An introduction to matrix concentration inequalities. *Foundations and trends® in machine learning*, 8(1-2):1–230, 2015.
- Verma, N. and Elbayad, M. Merging text transformer models from different initializations. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856.

- Vershynin, R. *High-Dimensional Probability*. Cambridge Series in Statistical and Probabilistic Mathematics. 2018.
- Wager, S., Wang, S., and Liang, P. S. Dropout training as adaptive regularization. *Advances in Neural Information Processing Systems*, 26, 2013.
- Wan, L., Zeiler, M., Zhang, S., Le Cun, Y., and Fergus, R. Regularization of neural networks using DropConnect. In *International Conference on Machine Learning*, pp. 1058–1066. PMLR, 2013.
- Wang, C., Zhang, G., and Grosse, R. Picking winning tickets before training by preserving gradient flow. In *The Eighth International Conference on Learning Representations*, 2020.
- Xie, Y. and Smidt, T. A tale of two symmetries: Exploring the loss landscape of equivariant models. *Advances in Neural Information Processing Systems*, 38:53223–53255, 2025.
- Yehudai, G. and Shamir, O. On the power and limitations of random features for understanding neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.
- Zhao, B., Dehmamy, N., Walters, R., and Yu, R. Symmetry teleportation for accelerated optimization. *Advances in Neural Information Processing Systems*, 35:16679–16690, 2022.
- Zhao, B., Ganev, I., Walters, R., Yu, R., and Dehmamy, N. Symmetries, flat minima, and the conserved quantities of gradient flow. In *The Eleventh International Conference on Learning Representations*, 2023.
- Zhao, B., Gower, R. M., Walters, R., and Yu, R. Improving convergence and generalization using parameter symmetries. In *The Twelfth International Conference on Learning Representations*, 2024.
- Zhao, B., Dehmamy, N., Walters, R., and Yu, R. Understanding mode connectivity via parameter space symmetry. In *Forty-second International Conference on Machine Learning*, 2025.
- Zhao, B., Walters, R., and Yu, R. Symmetry in neural network parameter spaces. *Transactions on Machine Learning Research*, 2026. ISSN 2835-8856.
- Zhou, Z., Yang, Y., Yang, X., Yan, J., and Hu, W. Going beyond linear mode connectivity: The layerwise linear feature connectivity. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Ziyin, L., Xu, Y., and Chuang, I. L. Remove symmetries to control model expressivity and improve optimization. In *The Thirteenth International Conference on Learning Representations*, 2025.

Appendices

A	Notation	9
B	Related Work	10
	B.1 Parameter Symmetries and Mode Connectivity	10
	B.2 Fixed Masks, Pruning, and Sketching	10
	B.3 Representation Costs	11
C	Neuron Identifiability (§ 3)	12
	C.1 Realization Costs	12
	C.2 Neuron Reassignments	15
	C.3 Input Reindexings	20
	C.4 Identifiability vs. Feature Learning	23
D	Linear Mode Connectivity (§ 4)	25
E	Approximate Subspace Support and Relaxed Realization Costs	29
	E.1 Approximate Subspace Support	29
	E.2 Relaxed Realization Costs	30
F	Additional Figures	33
	F.1 Visualization of Effective Function Classes	33
	F.2 LMC Interpolation for MLPs and ResNets	33
G	Experimental Setup	34
	G.1 MLPs	34
	G.2 ResNets	34
	G.3 Datasets and Augmentations	34
	G.4 Activation Matching	34
	G.5 Coherence Experiment	35
	G.6 Gaussian Mixture Synthetic Dataset	35
	G.7 Empirical Neuron Realization and Swap Costs	35
H	Additional Experimental Results	37
	H.1 Layerwise Activation Matching Objectives	37
	H.2 Empirical Validation of Center Separation Rate	39
	H.3 Neuron Swap Costs via Ridge Regression and Additional Pairwise Neuron Swap Costs	39
	H.4 Transformers	43
	H.5 Subspace Coherence of MNIST Inputs and Hidden Representations	43

A. Notation

Table 2. We list the most important symbols used throughout this work.

$\mathbb{N}, \mathbb{N}_0, \mathbb{Q}, \mathbb{R}$	Natural, non-negative integer, rational, real numbers.
$[n]$	Set $\{1, \dots, n\}$ for $n \in \mathbb{N}$.
$\mathbb{1}_A$	Indicator function in a set A .
$\mathcal{O}(\cdot), \Theta(\cdot), \Omega(\cdot)$	Landau symbols for asymptotic growth of a function.
S_n	Symmetric group $\{\pi : [n] \rightarrow [n] \mid \pi \text{ bijective}\}$.
$\text{id}, (ij), \pi$	Identity, transposition $i \leftrightarrow j$, generic permutation $\in S_n$.
$\mathbf{P}, \mathbf{P}_\pi \in \{0, 1\}^{n \times n}$	Generic permutation matrix, permutation matrix of $\pi \in S_n$.
\odot	Elementwise multiplication of matrices.
\mathbf{I}_n	Identity matrix in n dimensions.
$\mathbf{1}_n$	Vector of all ones in n dimensions.
\mathbf{e}_i	i -th canonical basis vector.
$\text{Unif}(\cdot)$	Uniform distribution on a set.
$\text{Bernoulli}(p)$	Bernoulli distribution with success probability p .
$\mathcal{N}(\mu, \sigma^2)$	Normal distribution with mean μ , variance σ^2 .
$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	Multivariate normal distribution with mean vector $\boldsymbol{\mu}$, covariance matrix $\boldsymbol{\Sigma}$.
$\ \cdot\ _p, \ \cdot\ _{\text{op}}, \ \cdot\ _F, \ \cdot\ _{\mathbf{A}}$	ℓ^p norms ($p \in [1, \infty]$), spectral norm, Frobenius norm, Mahalanobis norm w.r.t. \mathbf{A} .
\mathbf{A}^\dagger	Moore-Penrose pseudoinverse of \mathbf{A} .
$\succ, \succeq, \prec, \preceq$	Löwner partial order for matrices.
$L^p(\mu)$	p -integrable functions w.r.t. a measure $\mu, p \in [1, \infty]$.
$\ \cdot\ _{L^p(\mu)}, \ \cdot\ _{L^p(\mu, \mathbb{R}^m)}$	L^p norms of functions mapping to $\mathbb{R}, \mathbb{R}^m, p \in [1, \infty]$.
$f * g$	Convolution of two functions f, g .
λ^n	n -dimensional Lebesgue measure.
B_t^n	n -dimensional ball in \mathbb{R}^n , i.e., $\{\mathbf{x} \in \mathbb{R}^n \mid \ \mathbf{x}\ _2 < t\}$.
V_n	Volume $\lambda^n(B_1^n)$ of the n -dimensional unit ball.
$\ \cdot\ _{\psi_2}, \ \cdot\ _{\psi_1}$	Subgaussian and subexponential norms.
\mathbb{S}^{n-1}	n -dimensional sphere $\{\mathbf{x} \in \mathbb{R}^n \mid \ \mathbf{x}\ _2 = 1\}$.
$\boldsymbol{\theta}, \Theta$	(Total) parameter vector and parameter space of a neural architecture.
\mathbf{W}, \mathbf{w}	Weight matrix of a layer, weight vector of a single neuron.
d, m	Ambient input dimension, output dimension (i.e., width) of a layer.
η	Activation function.
$\mathbf{F}, \mathbf{D} \in \mathbb{R}^{m \times d}$	Fixed offset and diagonal operator of symmetry breaking intervention.
$\mathbf{f}_i, \mathbf{d}_i, \mathbf{w}_i \in \mathbb{R}^d$	i -th rows of $\mathbf{F}, \mathbf{D}, \mathbf{W}, i \in [m]$.
$\sigma_{\mathbf{F}}$	Standard deviation (fixed weight scale) of entry distribution of \mathbf{F} .
$\mu_{\mathbf{D}}, \sigma_{\mathbf{D}}, b_{\mathbf{D}}$	Mean, standard deviation, a.s. deviation of entry distribution of \mathbf{d}_i^2 .
$\mathbf{x}, \mathcal{X}, \mathbb{P}$	Input, input support, input distribution of a layer ($\mathcal{X} = \text{supp}(\mathbb{P})$).
$\mathbf{H}(\mathbf{W}, \mathbf{x}), \mathbf{h}_i(\mathbf{w}_i, \mathbf{x})$	Layer map of a symmetry-broken layer, single neuron i .
$\mathcal{H}, \mathcal{H}_i, \mathcal{H}_i(\mathcal{X})$	Fct. class of a layer, fct. class of neuron i , fct. class of neuron i w/ explicit input supp. \mathcal{X} .
k	Intrinsic input dimension $\leq d$.
\mathcal{U}, \mathbf{U}	Input subspace, orthonormal basis of \mathcal{U} .
$\mathbf{v}_i \in \mathbb{R}^k, \mathbf{M}_i \in \mathbb{R}^{k \times d}$	$\mathbf{F}_i, \text{Diag}(\mathbf{d}_i)$ in projected \mathbf{U} coordinates.
$\mathbf{S}_i \in \mathbb{R}^{k \times k}$	Gram matrix $\mathbf{S}_i = \mathbf{M}_i \mathbf{M}_i^\top$.
$\ \mathbf{h}\ _{\mathcal{H}_i(\mathcal{X})}$	Realization cost of $\mathbf{h} : \mathcal{X} \rightarrow \mathbb{R}$ w.r.t. neuron i on input support \mathcal{X} .
$\nu(\mathcal{U})$	Subspace coherence.
$\Delta_{\pi, \tau}(\mathbf{H}; \mathcal{X})$	Realization cost sensitivity of $\mathbf{H} \in \mathcal{H}(\mathcal{X})$ under output-input permutation pair (π, τ) .
$\delta_\pi^{\text{out}}, \delta_\tau^{\text{in}} \in \mathbb{R}^{m \times k}$	Projected center differences under output / input permutation π / τ .
$\gamma_{\text{out}}, \gamma_{\text{in}}$	Minimum proj. center distances for output / input.
$\xi_{\mathbf{H}}(\lambda; \mathbf{x})$	Chord deviation of layer \mathbf{H} at interpolation parameter $\lambda \in [0, 1]$.
ℓ, \mathcal{L}	Loss on one data point, population loss.

B. Related Work

B.1. Parameter Symmetries and Mode Connectivity

Mode connectivity of neural networks was introduced by early works which observed that trained neural network solutions can be connected by paths in parameter space where the performance (measured by training or validation loss or accuracy) does not significantly worsen (Garipov et al., 2018; Draxler et al., 2018). The focus soon shifted to linear mode connectivity (LMC) (Frankle et al., 2020), which has been widely interpreted to occur when two solutions are in the same basin of the optimization landscape, which is a desirable property for applications such as model merging. Several settings which admit LMC have been identified: Frankle et al. (2020) found LMC for networks trained with different optimization randomness but starting from the same partially converged solution. Later studies obtained LMC by aligning networks with respect to symmetries (Ainsworth et al., 2023), or removing symmetries (Lim et al., 2024b). Variants of LMC focusing on individual layers have been investigated (Adilova et al., 2024; Zhou et al., 2023).

Exact structural neural network parameter symmetries, in particular *permutation symmetries*, have been investigated as early as Hecht-Nielsen (1990); Chen et al. (1993), where it was observed that the whole parameter space can be mapped to a small cone via function-preserving permutations. Recent work has also studied hidden, local, and data-dependent symmetries, where equivalences may depend on the activation region, parameter point, or data distribution rather than preserving the realized function globally (Zhao et al., 2023; Grigsby et al., 2023; Xie & Smidt, 2025; Zhao et al., 2026).

Entezari et al. (2022) connected parameter symmetries to LMC and stated the conjecture that pairs of SGD solutions are likely linearly mode connected after applying a suitable permutation. A number of works have proposed permutation alignment algorithms, with the goal of obtaining interpolating paths of low barrier, based either on optimal transport (OT) of activations or weights (*activation or weight matching*), or directly optimizing for a permutation that decreases the midpoint barrier (*straight-through estimator or learned matching*) (Singh & Jaggi, 2020; Tatro et al., 2020; Pittorino et al., 2022; Ainsworth et al., 2023), alongside non-OT variants (Li et al., 2015; Crisostomi et al., 2024). Alignment methods have been primarily applied to MLPs (permuting neurons) and convolutional architectures such as ResNets (permuting channels), but have recently been successfully applied to Transformers (Imfeld et al., 2024; Verma & Elbayad, 2024; Theus et al., 2025), both for vision and language (where vision task models seem to be easier to align). Theoretical results in Ferbach et al. (2024) show aligned LMC can be obtained for trained wide one-hidden-layer neural networks in the mean-field regime.

There has also been interest in *unaligned LMC*, i.e., conditions under which LMC is given without post-hoc alignment. Recently, Lim et al. (2024b) proposed architectural symmetry breaking methods and obtained unaligned LMC with *W-asymmetric models*, where parts of the weights of each neuron are fixed in a way that breaks the usual permutation symmetries. Not primarily motivated by LMC but in a similar fashion, Ziyin et al. (2025) add fixed noise to parameters. Ito et al. (2026) find that increasingly wide models tend to yield LMC without symmetry alignment or removal, and explain this by wide layers exhibiting reciprocally orthogonal null spaces, so that merged models increasingly behave like an ensemble.

B.2. Fixed Masks, Pruning, and Sketching

When $\mathbf{d}_i \in \{0, 1\}^d$, the trainable part $\mathbf{d}_i \odot \mathbf{w}_i$ corresponds to a fixed sparse connectivity pattern, which is similar to a pruning-at-initialization or fixed DropConnect-style mask on the incoming weights (Wan et al., 2013). In contrast, standard dropout resamples activation masks during training (Srivastava et al., 2014). Dropout and related feature-noising schemes have been analyzed as induced regularizers (Wager et al., 2013) and, under variational assumptions, as approximate Bayesian inference (Gal & Ghahramani, 2016). The masks in this work, however, are sampled once and kept fixed across the independent training runs.

Prior work on pruning has studied post-training sparsification, lottery-ticket subnetworks (Frankle & Carbin, 2019), dynamic sparse training (Evci et al., 2020), and pruning at initialization (Lee et al., 2019; Wang et al., 2020; Tanaka et al., 2020; Frankle et al., 2021). Recent work has also connected pruning at initialization to randomized sketching (Bar & Giryes, 2025). Our use of this viewpoint is different: we do not use masks primarily to preserve an approximation of a dense network, but to understand when fixed diagonal operators make neurons functionally distinguishable on the input support. The relevant quantity for this distinction is the subspace coherence $\nu(\mathcal{U})$, equivalently the maximum leverage score of the input subspace (Drineas et al., 2012). Coherence is standard in matrix completion, compressed sensing, and randomized numerical linear algebra, where it controls whether uniform coordinate sampling preserves a low-dimensional structure (Candes & Tao, 2005; Candes & Recht, 2009). In our setting, the same quantity controls the concentration of $\mathbf{S}_i = \mathbf{U}^\top \text{Diag}(\mathbf{d}_i)^2 \mathbf{U}$. I.e., we use coherence not for exact recovery, but to quantify when fixed masks are visible to the effective function classes of neurons.

B.3. Representation Costs

In (4), we define the *realization cost* of a neuron or layer as the minimum weight norm required to implement it. This coincides with the concept of *representation cost* used in prior work (Dai et al., 2021; Gunasekar et al., 2018; Savarese et al., 2019). We call it *realization cost* to avoid confusion with *representations* in a layer.

C. Neuron Identifiability (§ 3)

Assumption C.1 (Subspace support model). The distribution \mathbb{P} is supported on a k -dimensional subspace $\mathcal{U} \subset \mathbb{R}^d$, i.e., $\mathbb{P}_x(x \in \mathcal{U}) = 1$. Throughout, we fix an orthonormal basis $\mathbf{U} \in \mathbb{R}^{d \times k}$ of \mathcal{U} , i.e., $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_k$ and $\text{im}(\mathbf{U}) = \mathcal{U}$.

Assumption C.2 (Non-degeneracy). Under **Assumption C.1**, additionally assume the following. If η is injective (e.g., tanh, LeakyReLU), assume $\text{span}(\text{supp}(\mathbb{P})) = \mathcal{U}$. If $s \mapsto (\eta(s), \eta(-s))$ is injective (e.g., ReLU, GELU), assume there is $T \subseteq \text{supp}(\mathbb{P})$ s.t. $T = -T$ and $\text{span}(T) = \mathcal{U}$.

C.1. Realization Costs

Lemma C.3 (Neuron identifiability from preactivations). *Under Assumptions C.1 and C.2, any \mathbf{h}_i , $i \in [m]$, can be identified from its preactivation. Formally,*

$$\mathbf{h}_i(\mathbf{w}; \cdot) = \mathbf{h}_i(\mathbf{w}'; \cdot) \Leftrightarrow \mathbf{U}^\top \text{Diag}(\mathbf{d}_i)(\mathbf{w} - \mathbf{w}') = 0, \quad (17)$$

where the l.h.s. is understood almost surely w.r.t. \mathbb{P} .

Proof. Fix $i \in [m]$, $\mathbf{w}, \mathbf{w}' \in \mathbb{R}^d$, and write the (row-wise) preactivation vectors

$$\mathbf{w}_{\text{eff}} := \mathbf{f}_i + \mathbf{d}_i \odot \mathbf{w} = \mathbf{f}_i + \text{Diag}(\mathbf{d}_i)\mathbf{w}, \quad \mathbf{w}'_{\text{eff}} := \mathbf{f}_i + \text{Diag}(\mathbf{d}_i)\mathbf{w}'. \quad (18)$$

By **Assumption C.1**, every input \mathbf{x} satisfies $\mathbf{x} \in \mathcal{U} = \text{im}(\mathbf{U})$, hence we can write $\mathbf{x} = \mathbf{U}\mathbf{c}$ for some $\mathbf{c} \in \mathbb{R}^k$. Define the subspace coefficients $\mathbf{a} := \mathbf{U}^\top \mathbf{w}_{\text{eff}} \in \mathbb{R}^k$ and $\mathbf{a}' := \mathbf{U}^\top \mathbf{w}'_{\text{eff}} \in \mathbb{R}^k$. For any $\mathbf{x} = \mathbf{U}\mathbf{c} \in \mathcal{U}$, we have the identity

$$\mathbf{w}_{\text{eff}}^\top \mathbf{x} = \mathbf{w}_{\text{eff}}^\top \mathbf{U}\mathbf{c} = (\mathbf{U}^\top \mathbf{w}_{\text{eff}})^\top \mathbf{c} = \mathbf{a}^\top \mathbf{c}, \quad (19)$$

and likewise $(\mathbf{w}'_{\text{eff}})^\top \mathbf{x} = \mathbf{a}'^\top \mathbf{c}$. Therefore, the neuron functions restricted to \mathcal{U} can be written as

$$\mathbf{h}_i(\mathbf{w}; \mathbf{U}\mathbf{c}) = \eta(\mathbf{a}^\top \mathbf{c}), \quad \mathbf{h}_i(\mathbf{w}'; \mathbf{U}\mathbf{c}) = \eta(\mathbf{a}'^\top \mathbf{c}). \quad (20)$$

Note that

$$\mathbf{a} - \mathbf{a}' = \mathbf{U}^\top (\mathbf{w}_{\text{eff}} - \mathbf{w}'_{\text{eff}}) = \mathbf{U}^\top \text{Diag}(\mathbf{d}_i)(\mathbf{w} - \mathbf{w}'). \quad (21)$$

We can now prove the equivalence.

“ \Leftarrow ”: If $\mathbf{U}^\top \text{Diag}(\mathbf{d}_i)(\mathbf{w} - \mathbf{w}') = 0$, then (21) and (20) directly yield $\mathbf{h}_i(\mathbf{w}; \cdot) = \mathbf{h}_i(\mathbf{w}'; \cdot)$ on \mathcal{U} .

“ \Rightarrow ”: Assume $\mathbf{h}_i(\mathbf{w}; \cdot) = \mathbf{h}_i(\mathbf{w}'; \cdot)$ on $\text{supp}(\mathbb{P}) \subseteq \mathcal{U}$ (by **Assumption C.1**). Equivalently,

$$\eta(\mathbf{a}^\top \mathbf{c}) = \eta(\mathbf{a}'^\top \mathbf{c}) \quad \text{for all } \mathbf{c} \in \mathbf{U}^\top(\text{supp}(\mathbb{P})) \subseteq \mathbb{R}^k. \quad (22)$$

We show this forces $\mathbf{a} = \mathbf{a}'$ under **Assumption C.2**, distinguishing the two activation cases.

Case ①: η injective. Injectivity gives $\mathbf{a}^\top \mathbf{c} = \mathbf{a}'^\top \mathbf{c}$ for all $\mathbf{c} \in \mathbf{U}^\top(\text{supp}(\mathbb{P}))$, hence $(\mathbf{a} - \mathbf{a}')^\top \mathbf{c} = 0$ for all such \mathbf{c} . By **Assumption C.2**, we have $\text{span}(\text{supp}(\mathbb{P})) = \mathcal{U}$, hence $\text{span}(\mathbf{U}^\top \text{supp}(\mathbb{P})) = \mathbb{R}^k$. Therefore, $(\mathbf{a} - \mathbf{a}')^\top \mathbf{c} = 0$ on a spanning set of \mathbb{R}^k , which implies $\mathbf{a} - \mathbf{a}' = \mathbf{U}^\top \text{Diag}(\mathbf{d}_i)(\mathbf{w} - \mathbf{w}') = 0$.

Case ②: $s \mapsto (\eta(s), \eta(-s))$ injective. Let $T \subseteq \text{supp}(\mathbb{P})$ be as in **Assumption C.2**, so $T = -T$ and $\text{span}(T) = \mathcal{U}$. Set $T_U := \mathbf{U}^\top T \subseteq \mathbb{R}^k$. Then $T_U = -T_U$ and $\text{span}(T_U) = \mathbb{R}^k$. Since $\eta(\mathbf{a}^\top \mathbf{c}) = \eta(\mathbf{a}'^\top \mathbf{c})$ for all $\mathbf{c} \in \mathbf{U}^\top(\text{supp}(\mathbb{P}))$, in particular $\eta(\mathbf{a}^\top \mathbf{c}) = \eta(\mathbf{a}'^\top \mathbf{c})$ for all $\mathbf{c} \in T_U$. Moreover, for any $\mathbf{c} \in T_U$ we have $-\mathbf{c} \in T_U$, hence also $\eta(-\mathbf{a}^\top \mathbf{c}) = \eta(-\mathbf{a}'^\top \mathbf{c})$. Therefore, for all $\mathbf{c} \in T_U$, $(\eta(\mathbf{a}^\top \mathbf{c}), \eta(-\mathbf{a}^\top \mathbf{c})) = (\eta(\mathbf{a}'^\top \mathbf{c}), \eta(-\mathbf{a}'^\top \mathbf{c}))$, and injectivity yields $\mathbf{a}^\top \mathbf{c} = \mathbf{a}'^\top \mathbf{c}$ for all $\mathbf{c} \in T_U$, i.e. $(\mathbf{a} - \mathbf{a}')^\top \mathbf{c} = 0$ for all $\mathbf{c} \in T_U$. Since $\text{span}(T_U) = \mathbb{R}^k$, we get $\mathbf{a} = \mathbf{a}'$, which concludes the proof. \square

Proposition C.4. Define \mathbf{v}_i as the projection of the fixed center of the i -th asymmetric neuron to the input subspace, and \mathbf{M}_i as its projected diagonal operator, i.e.,

$$\mathbf{v}_i := \mathbf{U}^\top \mathbf{f}_i \in \mathbb{R}^k, \quad \mathbf{M}_i := \mathbf{U}^\top \text{Diag}(\mathbf{d}_i) \in \mathbb{R}^{k \times k}. \quad (23)$$

Under **Assumptions C.1** and **C.2**, $\mathbf{h} \in \mathcal{H}_i(\mathcal{X})$ belonging to weight \mathbf{w} can be written as $\mathbf{h}(\mathbf{x}) = \eta((\mathbf{U}\mathbf{a})^\top \mathbf{x})$, where

$$\mathbf{a} := \mathbf{U}^\top (\mathbf{f}_i + \text{Diag}(\mathbf{d}_i)\mathbf{w}) = \mathbf{v}_i + \mathbf{M}_i \mathbf{w} \in \mathbb{R}^k. \quad (24)$$

The mapping $\mathbf{h} \mapsto \mathbf{a}$ is a bijection that identifies $\mathcal{H}_i(\mathcal{X})$ with $\mathbf{v}_i + \text{im}(\mathbf{M}_i) \subseteq \mathbb{R}^k$. In particular, $\mathcal{H}_i(\mathcal{X})$ is an affine subspace of \mathbb{R}^k of dimension $\text{rk}(\mathbf{M}_i) \leq k$.

Proof. Fix $i \in [m]$. Under **Assumption C.1**, every $\mathbf{x} \in \mathcal{X}$ satisfies $\mathbf{x} \in \text{im}(\mathbf{U})$, so we can write $\mathbf{x} = \mathbf{U}\mathbf{c}$ with $\mathbf{c} = \mathbf{U}^\top \mathbf{x} \in \mathbb{R}^k$. For any $\mathbf{w} \in \mathbb{R}^d$, define the induced subspace preactivation coefficient

$$\mathbf{a}(\mathbf{w}) := \mathbf{U}^\top (\mathbf{f}_i + \text{Diag}(\mathbf{d}_i)\mathbf{w}) = \mathbf{v}_i + \mathbf{M}_i \mathbf{w} \in \mathbb{R}^k. \quad (25)$$

Then for every $\mathbf{x} = \mathbf{U}\mathbf{c} \in \mathcal{X}$, $\mathbf{h}_i(\mathbf{w}; \mathbf{x}) = \eta((\mathbf{U}\mathbf{a}(\mathbf{w}))^\top \mathbf{x})$. This yields

$$\mathcal{H}_i(\mathcal{X}) \subseteq \left\{ \eta((\mathbf{U}\mathbf{a})^\top \cdot) \mid \mathbf{a} \in \mathbf{v}_i + \text{im}(\mathbf{M}_i) \right\}. \quad (26)$$

Conversely, let $\mathbf{a} \in \mathbf{v}_i + \text{im}(\mathbf{M}_i)$. Then there exists $\mathbf{w} \in \mathbb{R}^d$ such that $\mathbf{M}_i \mathbf{w} = \mathbf{a} - \mathbf{v}_i$, i.e., $\mathbf{a}(\mathbf{w}) = \mathbf{a}$ and $\mathbf{h}_i(\mathbf{w}; \cdot) = \eta((\mathbf{U}\mathbf{a})^\top \cdot)$ on \mathcal{X} , so the reverse inclusion in (26) holds as well.

It remains to argue that $\mathbf{a} \mapsto \eta((\mathbf{U}\mathbf{a})^\top \cdot)$ is injective on \mathbb{R}^k under [Assumption C.2](#). Let $\mathbf{a}, \mathbf{a}' \in \mathbb{R}^k$ satisfy $\eta((\mathbf{U}\mathbf{a})^\top \mathbf{x}) = \eta((\mathbf{U}\mathbf{a}')^\top \mathbf{x})$ for all $\mathbf{x} \in \text{supp}(\mathbb{P})$. Set $\mathbf{d} := \mathbf{a} - \mathbf{a}'$. If η is injective, then $(\mathbf{U}\mathbf{d})^\top \mathbf{x} = 0$ for all $\mathbf{x} \in \text{supp}(\mathbb{P})$, hence $\mathbf{d}^\top (\mathbf{U}^\top \mathbf{x}) = 0$ for all $\mathbf{x} \in \text{supp}(\mathbb{P})$. Since $\text{span}(\text{supp}(\mathbb{P})) = \mathcal{U} = \text{im}(\mathbf{U})$, we obtain $\mathbf{d} = 0$, i.e. $\mathbf{a} = \mathbf{a}'$. If $s \mapsto (\eta(s), \eta(-s))$ is injective, then for all $\mathbf{x} \in T \subseteq \text{supp}(\mathbb{P})$, with $T = -T$, we have both

$$\eta((\mathbf{U}\mathbf{a})^\top \mathbf{x}) = \eta((\mathbf{U}\mathbf{a}')^\top \mathbf{x}) \quad \text{and} \quad \eta((\mathbf{U}\mathbf{a})^\top (-\mathbf{x})) = \eta((\mathbf{U}\mathbf{a}')^\top (-\mathbf{x})), \quad (27)$$

so the injectivity of $s \mapsto (\eta(s), \eta(-s))$ yields $(\mathbf{U}\mathbf{a})^\top \mathbf{x} = (\mathbf{U}\mathbf{a}')^\top \mathbf{x}$ for all $\mathbf{x} \in T$. Since $\text{span}(T) = \mathcal{U}$, the same argument gives $\mathbf{a} = \mathbf{a}'$.

Having obtained the injectivity of $\mathbf{a} \mapsto \eta((\mathbf{U}\mathbf{a})^\top \cdot)$, the final claim follows immediately. \square

Lemma C.5 (Realization cost in subspace coefficients). *Let $i \in [m]$ and $\mathbf{a} \in \mathbb{R}^k$. Then,*

$$\left\| \eta((\mathbf{U}\mathbf{a})^\top \cdot) \right\|_{\mathcal{H}_i(\mathcal{X})} = \sqrt{(\mathbf{a} - \mathbf{v}_i)^\top \mathbf{S}_i^\dagger (\mathbf{a} - \mathbf{v}_i)} \quad (28)$$

$$=: \|\mathbf{a} - \mathbf{v}_i\|_{\mathbf{S}_i^\dagger} \quad (29)$$

for $\mathbf{a} - \mathbf{v}_i \in \text{im}(\mathbf{S}_i) = \text{im}(\mathbf{M}_i)$ (and $+\infty$ otherwise). Here, \mathbf{S}_i^\dagger denotes the Moore-Penrose pseudoinverse of \mathbf{S}_i .

Proof. Fix $i \in [m]$ and $\mathbf{a} \in \mathbb{R}^k$. By [Prop. C.4](#), we have

$$\eta((\mathbf{U}\mathbf{a})^\top \cdot) \in \mathcal{H}_i(\mathcal{X}) \quad \Leftrightarrow \quad \mathbf{a} - \mathbf{v}_i \in \text{im}(\mathbf{M}_i), \quad (30)$$

and in that case $\eta((\mathbf{U}\mathbf{a})^\top \cdot) = \mathbf{h}_i(\mathbf{w}; \cdot)$ holds iff $\mathbf{M}_i \mathbf{w} = \mathbf{a} - \mathbf{v}_i$. Therefore, by definition of realization cost (4),

$$\left\| \eta((\mathbf{U}\mathbf{a})^\top \cdot) \right\|_{\mathcal{H}_i(\mathcal{X})} = \inf \{ \|\mathbf{w}\|_2 \mid \mathbf{M}_i \mathbf{w} = \mathbf{a} - \mathbf{v}_i \}, \quad (31)$$

with the convention $+\infty$ if $\mathbf{a} - \mathbf{v}_i \notin \text{im}(\mathbf{M}_i)$. Assume now that $\mathbf{a} - \mathbf{v}_i \in \text{im}(\mathbf{M}_i)$. Recalling $\mathbf{S}_i = \mathbf{M}_i \mathbf{M}_i^\top$, the (unique) minimum ℓ^2 norm solution is

$$\mathbf{w}^* = \mathbf{M}_i^\top (\mathbf{M}_i \mathbf{M}_i^\top)^\dagger (\mathbf{a} - \mathbf{v}_i) = \mathbf{M}_i^\top \mathbf{S}_i^\dagger (\mathbf{a} - \mathbf{v}_i), \quad (32)$$

and every solution can be written as $\mathbf{w} = \mathbf{w}^* + \mathbf{w}_\perp$ with $\mathbf{w}_\perp \in \ker(\mathbf{M}_i)$, which implies $\|\mathbf{w}\|_2^2 = \|\mathbf{w}^*\|_2^2 + \|\mathbf{w}_\perp\|_2^2 \geq \|\mathbf{w}^*\|_2^2$. Hence, $\inf \{ \|\mathbf{w}\|_2 \mid \mathbf{M}_i \mathbf{w} = \mathbf{a} - \mathbf{v}_i \} = \|\mathbf{w}^*\|_2$. Moreover,

$$\|\mathbf{w}^*\|_2^2 = (\mathbf{a} - \mathbf{v}_i)^\top \mathbf{S}_i^\dagger \mathbf{S}_i \mathbf{S}_i^\dagger (\mathbf{a} - \mathbf{v}_i) \quad (33)$$

$$= (\mathbf{a} - \mathbf{v}_i)^\top \mathbf{S}_i^\dagger (\mathbf{a} - \mathbf{v}_i), \quad (34)$$

where we used $\mathbf{S}_i^\dagger \mathbf{S}_i \mathbf{S}_i^\dagger = \mathbf{S}_i^\dagger$. This yields

$$\left\| \eta((\mathbf{U}\mathbf{a})^\top \cdot) \right\|_{\mathcal{H}_i(\mathcal{X})} = \sqrt{(\mathbf{a} - \mathbf{v}_i)^\top \mathbf{S}_i^\dagger (\mathbf{a} - \mathbf{v}_i)}, \quad (35)$$

which holds for $\mathbf{a} - \mathbf{v}_i \in \text{im}(\mathbf{S}_i) = \text{im}(\mathbf{M}_i)$ (and equals $+\infty$ otherwise). \square

Theorem C.6 (Spectral concentration of \mathbf{S}_i). *Let $i \in [m]$ and sample the entries of \mathbf{d}_i i.i.d. with $(\mathbf{d}_i)_\ell^2$ having mean $\mu_{\mathbf{D}}$, variance $\sigma_{\mathbf{D}}^2$, and a.s. deviation $b_{\mathbf{D}}$ from its mean (all $< \infty$). For any $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$\|\mathbf{S}_i - \mu_{\mathbf{D}} \mathbf{I}_k\|_{\text{op}} \leq C \left\{ \sqrt{\sigma_{\mathbf{D}}^2 \nu(\mathcal{U}) \Lambda} + b_{\mathbf{D}} \nu(\mathcal{U}) \Lambda \right\}, \quad (36)$$

where $\Lambda := \log(2k/\delta)$, $\nu(\mathcal{U}) := \max_{\ell \in [d]} (\mathbf{U}\mathbf{U}^\top)_{\ell\ell}$ is the subspace coherence, and C is a universal constant.³

The proof of this theorem relies on an application of the matrix Bernstein inequality.

Lemma C.7 (Matrix Bernstein Inequality; Theorem 1.4 of [Tropp \(2015\)](#)). *Let $\mathbf{A}_1, \dots, \mathbf{A}_d \in \mathbb{R}^{k \times k}$ be independent,*

³One can also control $\max_i \|\mathbf{S}_i - \mu_{\mathbf{D}} \mathbf{I}_k\|_{\text{op}}$ via a union bound at the cost of replacing $\log(2k/\delta)$ by $\log(2mk/\delta)$.

centered random symmetric matrices. Assume a uniform bound $\|\mathbf{A}_l\|_{\text{op}} \leq R$ almost surely for each addend $l \in [d]$. Let

$$\sigma^2 := \left\| \sum_{l=1}^d \mathbb{E}[\mathbf{A}_l^2] \right\|_{\text{op}}. \quad (37)$$

Then, for all $t \geq 0$,

$$\mathbb{P} \left(\left\| \sum_{l=1}^d \mathbf{A}_l \right\|_{\text{op}} \geq t \right) \leq 2k \exp \left(\frac{-t^2/2}{\sigma^2 + Rt/3} \right). \quad (38)$$

Proof of Thm. C.6. Denote the l -th row of \mathbf{U} by $\mathbf{u}_l \in \mathbb{R}^k$. Observe that

$$(\mathbf{U}\mathbf{U}^\top)_{ll} = \mathbf{e}_l^\top \mathbf{U}\mathbf{U}^\top \mathbf{e}_l = \|\mathbf{U}^\top \mathbf{e}_l\|_2^2 = \|\mathbf{u}_l\|_2^2, \quad (39)$$

where \mathbf{e}_l denotes the l -th canonical basis vector in \mathbb{R}^d . The fluctuation of the Gram matrix \mathbf{S}_i around its mean

$$\mathbb{E}[\mathbf{S}_i] = \mathbf{U}^\top \mathbb{E}[\text{diag}(\mathbf{d}_i)^2] \mathbf{U} = \mathbf{U}^\top (\mu_{\mathbf{D}} \mathbf{I}_d) \mathbf{U} = \mu_{\mathbf{D}} \mathbf{I}_k \quad (40)$$

can be written as a sum of independent rank 1 matrices

$$\mathbf{S}_i - \mu_{\mathbf{D}} \mathbf{I}_k = \sum_{l=1}^d \underbrace{((\mathbf{d}_i)_l^2 - \mu_{\mathbf{D}}) \mathbf{u}_l \mathbf{u}_l^\top}_{=: \mathbf{A}_l} = \sum_{l=1}^d \mathbf{A}_l. \quad (41)$$

Clearly, the addends \mathbf{A}_l are centered and symmetric. Furthermore, we can bound

$$\|\mathbf{A}_l\|_{\text{op}} = |(\mathbf{d}_i)_l^2 - \mu_{\mathbf{D}}| \cdot \|\mathbf{u}_l \mathbf{u}_l^\top\|_{\text{op}} \leq b_{\mathbf{D}} \|\mathbf{u}_l\|_2^2 \leq b_{\mathbf{D}} \nu(\mathcal{U}) =: R. \quad (42)$$

Next, we compute the variance parameter σ^2 . Since

$$\mathbf{A}_l^2 = ((\mathbf{d}_i)_l^2 - \mu_{\mathbf{D}})^2 (\mathbf{u}_l \mathbf{u}_l^\top)^2 = ((\mathbf{d}_i)_l^2 - \mu_{\mathbf{D}})^2 \|\mathbf{u}_l\|_2^2 (\mathbf{u}_l \mathbf{u}_l^\top) \quad (43)$$

and $\mathbb{E}[(\mathbf{d}_i)_l^2 - \mu_{\mathbf{D}}]^2 = \sigma_{\mathbf{D}}^2$ by definition, we have

$$\sigma^2 = \left\| \sum_{l=1}^d \sigma_{\mathbf{D}}^2 \|\mathbf{u}_l\|_2^2 (\mathbf{u}_l \mathbf{u}_l^\top) \right\|_{\text{op}} \quad (44)$$

$$\stackrel{(*)}{\leq} \sigma_{\mathbf{D}}^2 \underbrace{\left(\max_k \|\mathbf{u}_k\|_2^2 \right)}_{=: \nu(\mathcal{U})} \left\| \sum_{l=1}^d \mathbf{u}_l \mathbf{u}_l^\top \right\|_{\text{op}} \quad (45)$$

$$= \sigma_{\mathbf{D}}^2 \nu(\mathcal{U}) \underbrace{\|\mathbf{U}^\top \mathbf{U}\|_{\text{op}}}_{=: \mathbf{I}_k} \quad (46)$$

$$= \sigma_{\mathbf{D}}^2 \nu(\mathcal{U}). \quad (47)$$

In (*), we have utilized that for positive semidefinite matrices $\mathbf{B}_l \succeq 0$ and scalars $c_l \in [0, C]$, one has $\|\sum_l c_l \mathbf{B}_l\|_{\text{op}} \leq C \|\sum_l \mathbf{B}_l\|_{\text{op}}$. We are now ready to apply [Lem. C.7](#) to $\sum_{l=1}^d \mathbf{A}_l$. We set the target failure probability to $\delta \in (0, 1)$. We seek a threshold t such that the failure probability is at most δ . This holds if

$$2k \exp \left(\frac{-t^2/2}{\sigma^2 + Rt/3} \right) \leq \delta \Leftrightarrow \frac{t^2}{2} \geq \left(\sigma^2 + \frac{Rt}{3} \right) \log \left(\frac{2k}{\delta} \right). \quad (48)$$

Let $L := \log(2k/\delta)$. The inequality becomes $t^2 - \frac{2}{3}RLt - 2\sigma^2L \geq 0$. Solving for t yields

$$t \geq \frac{RL}{3} + \sqrt{\left(\frac{RL}{3} \right)^2 + 2\sigma^2L} \stackrel{(**)}{\leq} \frac{2}{3}RL + \sigma\sqrt{2L}, \quad (49)$$

where (**) uses $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b \geq 0$ (setting t to this upper bound suffices). Finally, we can plug in $R = b_{\mathbf{D}}\nu(\mathcal{U})$ ((42)) and $\sigma^2 \leq \sigma_{\mathbf{D}}^2\nu(\mathcal{U})$ ((47)), and absorbing absolute constants into C , we obtain

$$t \geq C \left\{ \sqrt{\sigma_{\mathbf{D}}^2\nu(\mathcal{U}) \log(2k/\delta)} + b_{\mathbf{D}}\nu(\mathcal{U}) \log(2k/\delta) \right\}, \quad (50)$$

which completes the proof. \square

Lemma C.8. Let $\mathbf{H} \in \mathcal{H}(\mathcal{X})$ with $\mathbf{h}_i = \eta((\mathbf{U}\mathbf{a}_i)^\top \cdot)$ and $\mathbf{r}_i := \mathbf{a}_i - \mathbf{v}_i$. Then,

$$\Delta_{(ij)}^{\text{out}}(\mathbf{H}) = \underbrace{\|\mathbf{v}_j - \mathbf{v}_i\|_{\mathbf{S}_i^{-1} + \mathbf{S}_j^{-1}}^2}_{(i) \text{ center term}} + \underbrace{2\langle \mathbf{v}_j - \mathbf{v}_i, \mathbf{S}_i^{-1}\mathbf{r}_j - \mathbf{S}_j^{-1}\mathbf{r}_i \rangle}_{(ii) \text{ cross term}} + \underbrace{\|\mathbf{r}_j\|_{\mathbf{S}_i^{-1} - \mathbf{S}_j^{-1}}^2 + \|\mathbf{r}_i\|_{\mathbf{S}_j^{-1} - \mathbf{S}_i^{-1}}^2}_{(iii) \text{ metric mismatch term}}. \quad (51)$$

Proof. Fix $i \neq j$ and write $\mathbf{v}_{ij} := \mathbf{v}_j - \mathbf{v}_i$. For a transposition (ij) , only the i - and j -addends change, hence

$$\begin{aligned} \Delta_{(ij)}^{\text{out}}(\mathbf{H}) &= \|\mathbf{v}_j - \mathbf{v}_i + \mathbf{r}_j\|_{\mathbf{S}_i^{-1}}^2 + \|\mathbf{v}_i - \mathbf{v}_j + \mathbf{r}_i\|_{\mathbf{S}_j^{-1}}^2 - \|\mathbf{r}_i\|_{\mathbf{S}_i^{-1}}^2 - \|\mathbf{r}_j\|_{\mathbf{S}_j^{-1}}^2 \\ &= \|\mathbf{v}_{ij} + \mathbf{r}_j\|_{\mathbf{S}_i^{-1}}^2 + \|\mathbf{v}_{ij} + \mathbf{r}_i\|_{\mathbf{S}_j^{-1}}^2 - \|\mathbf{r}_i\|_{\mathbf{S}_i^{-1}}^2 - \|\mathbf{r}_j\|_{\mathbf{S}_j^{-1}}^2. \end{aligned} \quad (52)$$

Using $\|\mathbf{x}\|_{\mathbf{A}}^2 = \mathbf{x}^\top \mathbf{A} \mathbf{x}$ and expanding the quadratic forms in (52) yields

$$\|\mathbf{v}_{ij} + \mathbf{r}_j\|_{\mathbf{S}_i^{-1}}^2 = \mathbf{v}_{ij}^\top \mathbf{S}_i^{-1} \mathbf{v}_{ij} + 2\mathbf{v}_{ij}^\top \mathbf{S}_i^{-1} \mathbf{r}_j + \mathbf{r}_j^\top \mathbf{S}_i^{-1} \mathbf{r}_j, \quad (53)$$

$$\|\mathbf{v}_{ij} + \mathbf{r}_i\|_{\mathbf{S}_j^{-1}}^2 = \mathbf{v}_{ij}^\top \mathbf{S}_j^{-1} \mathbf{v}_{ij} + 2\mathbf{v}_{ij}^\top \mathbf{S}_j^{-1} \mathbf{r}_i + \mathbf{r}_i^\top \mathbf{S}_j^{-1} \mathbf{r}_i. \quad (54)$$

Substituting (53) and (54) into (52) and regrouping gives

$$\begin{aligned} \Delta_{(ij)}^{\text{out}}(\mathbf{H}) &= \mathbf{v}_{ij}^\top (\mathbf{S}_i^{-1} + \mathbf{S}_j^{-1}) \mathbf{v}_{ij} + 2\mathbf{v}_{ij}^\top (\mathbf{S}_i^{-1} \mathbf{r}_j - \mathbf{S}_j^{-1} \mathbf{r}_i) \\ &\quad + \mathbf{r}_j^\top (\mathbf{S}_i^{-1} - \mathbf{S}_j^{-1}) \mathbf{r}_j + \mathbf{r}_i^\top (\mathbf{S}_j^{-1} - \mathbf{S}_i^{-1}) \mathbf{r}_i, \end{aligned} \quad (55)$$

which is precisely (51). \square

C.2. Neuron Reassignments

Assumption C.9 (Spectral concentration of Gram anisotropies). Fix $\varepsilon \in (0, \frac{\mu_{\mathbf{D}}}{2})$, and assume

$$\max_{i \in [m]} \|\mathbf{S}_i - \mu_{\mathbf{D}} \mathbf{I}_k\|_{\text{op}} \leq \varepsilon. \quad (56)$$

Theorem C.10 (Global output sensitivity). Consider a layer $\mathbf{H} = \eta((\mathbf{F} + \mathbf{D} \odot \mathbf{W}) \cdot)$, and assume [Assumption C.9](#) holds. Set $\mathbf{R} := (\mathbf{D} \odot \mathbf{W})\mathbf{U} \in \mathbb{R}^{m \times k}$ (proj. trainable parts). Let $\pi \in S_m$ and set $\mathbf{M}_\pi := \mathbf{I}_m - \text{Diag}(\text{diag}(\mathbf{P}_\pi))$ to be the projection onto the coordinates affected by π . Assume that $\delta_\pi^{\text{out}} \neq 0$, $\rho_\pi^{\text{out}} := \|\mathbf{M}_\pi \mathbf{R}\|_F / \|\delta_\pi^{\text{out}}\|_F \leq 1$. Then,

$$\Delta_\pi^{\text{out}}(\mathbf{H}) = \left(1 \pm \mathcal{O}(\rho_\pi^{\text{out}} + \varepsilon \mu_{\mathbf{D}}^{-1})\right) \cdot \mu_{\mathbf{D}}^{-1} \|\delta_\pi^{\text{out}}\|_F^2. \quad (57)$$

Proof. First, note that by definition,

$$\Delta_\pi^{\text{out}}(\mathbf{H}, \mathcal{X}) = \sum_{i=1}^m \left(\|\delta_\pi^{\text{out}}\|_{i,:}^2 + \|\mathbf{P}_\pi \mathbf{R}\|_{i,:}^2_{\mathbf{S}_i^{-1}} - \|\mathbf{R}_{i,:}\|_{\mathbf{S}_i^{-1}}^2 \right). \quad (58)$$

Now if $\pi(i) = i$, then $\delta_\pi^{\text{out}}\|_{i,:} = 0$ and $\mathbf{P}_\pi \mathbf{R}\|_{i,:} = \mathbf{R}_{i,:}$, so the i -th summand in (58) vanishes. Since \mathbf{M}_π is the (diagonal) projection onto the rows moved by π , we may write

$$\Delta_\pi^{\text{out}}(\mathbf{H}, \mathcal{X}) = \sum_{i=1}^m \left(\|\mathbf{M}_\pi (\delta_\pi^{\text{out}} + \mathbf{P}_\pi \mathbf{R})\|_{i,:}^2_{\mathbf{S}_i^{-1}} - \|\mathbf{M}_\pi \mathbf{R}\|_{i,:}^2_{\mathbf{S}_i^{-1}} \right). \quad (59)$$

Since $\varepsilon < \mu_{\mathbf{D}}$, the uniform bound $\|\mathbf{S}_i - \mu_{\mathbf{D}} \mathbf{I}_k\|_{\text{op}} \leq \varepsilon$ in i yields $(\mu_{\mathbf{D}} - \varepsilon) \mathbf{I}_k \preceq \mathbf{S}_i \preceq (\mu_{\mathbf{D}} + \varepsilon) \mathbf{I}_k$. Hence \mathbf{S}_i is invertible, and

$$\|\mathbf{S}_i^{-1} - \mu_{\mathbf{D}}^{-1} \mathbf{I}_k\|_{\text{op}} \leq \frac{\varepsilon}{\mu_{\mathbf{D}}(\mu_{\mathbf{D}} - \varepsilon)} =: \beta. \quad (60)$$

Thus, for every $\mathbf{a} \in \mathbb{R}^k$,

$$\left| \|\mathbf{a}\|_{\mathbf{S}_i^{-1}}^2 - \mu_{\mathbf{D}}^{-1} \|\mathbf{a}\|_2^2 \right| \leq \beta \|\mathbf{a}\|_2^2. \quad (61)$$

Applying (61) to each term in (59) yields

$$\Delta_\pi^{\text{out}}(\mathbf{H}, \mathcal{X}) = \mu_{\mathbf{D}}^{-1} \left(\|\mathbf{M}_\pi (\delta_\pi^{\text{out}} + \mathbf{P}_\pi \mathbf{R})\|_F^2 - \|\mathbf{M}_\pi \mathbf{R}\|_F^2 \right) + E_\pi, \quad (62)$$

with

$$|E_\pi| \leq \beta \left(\|\mathbf{M}_\pi (\delta_\pi^{\text{out}} + \mathbf{P}_\pi \mathbf{R})\|_F^2 + \|\mathbf{M}_\pi \mathbf{R}\|_F^2 \right). \quad (63)$$

Because $\mathbf{M}_\pi \delta_\pi^{\text{out}} = \delta_\pi^{\text{out}}$ and $\mathbf{M}_\pi \mathbf{P}_\pi = \mathbf{P}_\pi \mathbf{M}_\pi$ implies $\|\mathbf{M}_\pi \mathbf{P}_\pi \mathbf{R}\|_F = \|\mathbf{P}_\pi \mathbf{M}_\pi \mathbf{R}\|_F = \|\mathbf{M}_\pi \mathbf{R}\|_F$, we can apply Cauchy-Schwarz to bound (63) as

$$|E_\pi| \leq \beta \left(\|\delta_\pi^{\text{out}}\|_F^2 + 2\|\delta_\pi^{\text{out}}\|_F \|\mathbf{M}_\pi \mathbf{R}\|_F + 2\|\mathbf{M}_\pi \mathbf{R}\|_F^2 \right). \quad (64)$$

It remains to simplify the isotropic term in (62), for which one can obtain

$$\|M_\pi(\delta_\pi^{\text{out}} + P_\pi \mathbf{R})\|_F^2 - \|M_\pi \mathbf{R}\|_F^2 = \|\delta_\pi^{\text{out}}\|_F^2 + 2\langle \delta_\pi^{\text{out}}, M_\pi P_\pi \mathbf{R} \rangle_F = \|\delta_\pi^{\text{out}}\|_F^2 + 2\langle \delta_\pi^{\text{out}}, P_\pi \mathbf{R} \rangle_F, \quad (65)$$

where the last equality holds since $M_\pi^\top = M_\pi$ and $M_\pi \delta_\pi^{\text{out}} = \delta_\pi^{\text{out}}$. Define

$$\mathbf{Q}_\pi := \mathbf{I}_m - \text{ord}(\pi)^{-1} \sum_{t=0}^{\text{ord}(\pi)-1} P_\pi^t. \quad (66)$$

Since $P_\pi^{\text{ord}(\pi)} = \mathbf{I}_m$,

$$(\mathbf{I}_m - \mathbf{Q}_\pi)P_\pi = P_\pi(\mathbf{I}_m - \mathbf{Q}_\pi) = \mathbf{I}_m - \mathbf{Q}_\pi. \quad (67)$$

Thus, $\mathbf{I}_m - \mathbf{Q}_\pi$ is the orthogonal projection onto the fixed subspace of P_π , and \mathbf{Q}_π is the orthogonal projection onto its orthogonal complement. Moreover,

$$(\mathbf{I}_m - \mathbf{Q}_\pi)\delta_\pi^{\text{out}} = (\mathbf{I}_m - \mathbf{Q}_\pi)(P_\pi - \mathbf{I}_m)\mathbf{F}\mathbf{U} = 0. \quad (68)$$

Hence, $\mathbf{Q}_\pi \delta_\pi^{\text{out}} = \delta_\pi^{\text{out}}$ and, using $P_\pi(\mathbf{I}_m - \mathbf{Q}_\pi) = \mathbf{I}_m - \mathbf{Q}_\pi$,

$$\langle \delta_\pi^{\text{out}}, P_\pi(\mathbf{I}_m - \mathbf{Q}_\pi)\mathbf{R} \rangle_F = \langle \delta_\pi^{\text{out}}, (\mathbf{I}_m - \mathbf{Q}_\pi)\mathbf{R} \rangle_F = 0. \quad (69)$$

Therefore,

$$\langle \delta_\pi^{\text{out}}, P_\pi \mathbf{R} \rangle_F = \langle \delta_\pi^{\text{out}}, P_\pi \mathbf{Q}_\pi \mathbf{R} \rangle_F. \quad (70)$$

By Cauchy–Schwarz and the orthogonality of P_π ,

$$|\langle \delta_\pi^{\text{out}}, P_\pi \mathbf{R} \rangle_F| \leq \|\delta_\pi^{\text{out}}\|_F \|\mathbf{Q}_\pi \mathbf{R}\|_F. \quad (71)$$

Combining (62), (64), (65), and (71) gives

$$\begin{aligned} |\Delta_\pi^{\text{out}}(\mathbf{H}, \mathcal{X}) - \mu_{\mathbf{D}}^{-1} \|\delta_\pi^{\text{out}}\|_F^2| &\leq \frac{2}{\mu_{\mathbf{D}}} \|\delta_\pi^{\text{out}}\|_F \|\mathbf{Q}_\pi \mathbf{R}\|_F \\ &\quad + \frac{\varepsilon}{\mu_{\mathbf{D}}(\mu_{\mathbf{D}} - \varepsilon)} \left(\|\delta_\pi^{\text{out}}\|_F^2 + 2\|\delta_\pi^{\text{out}}\|_F \|M_\pi \mathbf{R}\|_F + 2\|M_\pi \mathbf{R}\|_F^2 \right). \end{aligned} \quad (72)$$

For the final statement, note that $\mathbf{Q}_\pi = \mathbf{Q}_\pi M_\pi$, and as \mathbf{Q}_π is an orthogonal projection,

$$\|\mathbf{Q}_\pi \mathbf{R}\|_F = \|\mathbf{Q}_\pi M_\pi \mathbf{R}\|_F \leq \|M_\pi \mathbf{R}\|_F. \quad (73)$$

Dividing (72) by $\mu_{\mathbf{D}}^{-1} \|\delta_\pi^{\text{out}}\|_F^2$ therefore yields

$$\left| \frac{\mu_{\mathbf{D}} \Delta_\pi^{\text{out}}(\mathbf{H}, \mathcal{X})}{\|\delta_\pi^{\text{out}}\|_F^2} - 1 \right| \leq 2\rho_\pi^{\text{out}} + \frac{\varepsilon}{\mu_{\mathbf{D}} - \varepsilon} (1 + 2\rho_\pi^{\text{out}} + 2(\rho_\pi^{\text{out}})^2). \quad (74)$$

If $\rho_\pi^{\text{out}} \leq 1$ and $\varepsilon \leq \mu_{\mathbf{D}}/2$, then

$$\frac{\varepsilon}{\mu_{\mathbf{D}} - \varepsilon} \leq \frac{2\varepsilon}{\mu_{\mathbf{D}}}, \quad 1 + 2\rho_\pi^{\text{out}} + 2(\rho_\pi^{\text{out}})^2 \leq 5. \quad (75)$$

Hence, the r.h.s. is $\mathcal{O}(\rho_\pi^{\text{out}} + \varepsilon\mu_{\mathbf{D}}^{-1})$, and therefore

$$\Delta_\pi^{\text{out}}(\mathbf{H}, \mathcal{X}) = \left(1 \pm \mathcal{O}(\rho_\pi^{\text{out}} + \varepsilon\mu_{\mathbf{D}}^{-1})\right) \cdot \mu_{\mathbf{D}}^{-1} \|\delta_\pi^{\text{out}}\|_F^2. \quad (76)$$

Corollary C.11 (Minimum output reassignment cost). *In the setup of Thm. C.10, suppose further that $\rho_{(ij)}^{\text{out}} \leq \rho^{\text{out}} \leq 1$ for some ρ^{out} , uniformly in $i \neq j \in [m]$. Then,*

$$\min_{\text{id} \neq \pi \in S_m} \Delta_\pi^{\text{out}}(\mathbf{H}, \mathcal{X}) = \left(1 \pm \mathcal{O}(\rho^{\text{out}} + \varepsilon\mu_{\mathbf{D}}^{-1})\right) \cdot \mu_{\mathbf{D}}^{-1} \gamma_{\text{out}}^2. \quad (77)$$

Proof. Fix $\text{id} \neq \pi \in S_m$. Then,

$$\|\delta_\pi^{\text{out}}\|_F^2 = \|(P_\pi - \mathbf{I}_m)\mathbf{F}\mathbf{U}\|_F^2 = \sum_{i=1}^m \|\mathbf{v}_{\pi(i)} - \mathbf{v}_i\|_2^2. \quad (78)$$

Write π as a product of disjoint cycles, and consider one nontrivial cycle $(i_1 \dots i_\ell)$ of length $\ell \geq 2$, and use the convention $i_{\ell+1} := i_1$. Its contribution to the above sum is

$$\sum_{r=1}^{\ell} \|\mathbf{v}_{i_{r+1}} - \mathbf{v}_{i_r}\|_2^2 \geq \ell \min_{i \neq j} \|\mathbf{v}_j - \mathbf{v}_i\|_2^2 \geq 2 \min_{i \neq j} \|\mathbf{v}_j - \mathbf{v}_i\|_2^2. \quad (79)$$

On the other hand, for a transposition (ij) one has $\|\delta_{(ij)}^{\text{out}}\|_F^2 = 2\|\mathbf{v}_j - \mathbf{v}_i\|_2^2$. Therefore,

$$\min_{\text{id} \neq \pi \in S_m} \|\delta_\pi^{\text{out}}\|_F = \min_{i \neq j} \|\delta_{(ij)}^{\text{out}}\|_F = \gamma_{\text{out}}. \quad (80)$$

Next, for every $i \neq j$,

$$\|\mathbf{R}_{i,:}\|_2^2 + \|\mathbf{R}_{j,:}\|_2^2 = \|\mathbf{M}_{(ij)}\mathbf{R}\|_F^2 = (\rho_{(ij)}^{\text{out}})^2 \|\delta_{(ij)}^{\text{out}}\|_F^2 \leq 2(\rho^{\text{out}})^2 \|\mathbf{v}_j - \mathbf{v}_i\|_2^2. \quad (81)$$

Now let $\pi \neq \text{id}$ and decompose its moved indices into s nontrivial disjoint cycles $(i_1^{(s)} \dots i_{\ell_s}^{(s)})$. Then,

$$2\|\mathbf{M}_\pi\mathbf{R}\|_F^2 = \sum_s \sum_{r=1}^{\ell_s} \left(\|\mathbf{R}_{i_r^{(s)}}\|_2^2 + \|\mathbf{R}_{i_{r+1}^{(s)}}\|_2^2 \right) \leq 2(\rho^{\text{out}})^2 \sum_s \sum_{r=1}^{\ell_s} \|\mathbf{v}_{i_{r+1}^{(s)}} - \mathbf{v}_{i_r^{(s)}}\|_2^2 = 2(\rho^{\text{out}})^2 \|\delta_\pi^{\text{out}}\|_F^2. \quad (82)$$

Hence, $\rho_\pi^{\text{out}} = \|\mathbf{M}_\pi\mathbf{R}\|_F / \|\delta_\pi^{\text{out}}\|_F \leq \rho^{\text{out}}$ for all $\pi \neq \text{id}$. By Theorem C.10, it follows that

$$\Delta_\pi^{\text{out}}(\mathbf{H}, \mathcal{X}) = \left(1 \pm \mathcal{O}(\rho^{\text{out}} + \varepsilon\mu_{\mathbf{D}}^{-1})\right) \cdot \mu_{\mathbf{D}}^{-1} \|\delta_\pi^{\text{out}}\|_F^2 \quad (83)$$

uniformly in $\pi \neq \text{id}$. Using (80) yields

$$\min_{\text{id} \neq \pi \in S_m} \Delta_\pi^{\text{out}}(\mathbf{H}, \mathcal{X}) = \left(1 \pm \mathcal{O}(\rho^{\text{out}} + \varepsilon\mu_{\mathbf{D}}^{-1})\right) \cdot \mu_{\mathbf{D}}^{-1} \gamma_{\text{out}}^2. \quad (84)$$

Theorem C.12 (Asymptotics of projected center gap γ_{out}). *Let \mathbf{F} have i.i.d. entries of the form $B \cdot \sigma_{\mathbf{F}} Y_0$, where $B \sim \text{Bernoulli}(p)$, $p \in (0, 1]$, and Y_0 is independent of B with a Lebesgue density bounded by $M_0 < \infty$. Set $q := 2p - p^2$ and fix $\delta \in (0, 1)$. Let $m \geq 2$ and assume*

$$q \geq C \left\{ \sqrt{q\nu(\mathcal{U})\Lambda} + \nu(\mathcal{U})\Lambda \right\}, \quad \Lambda := \log(4km^2/\delta). \quad (85)$$

Then, with probability at least $1 - \delta$,

$$\gamma_{\text{out}} \geq c \frac{\sqrt{q}}{M_0} \sigma_{\mathbf{F}} V_k^{-1/k} \left(\frac{\delta}{m^2} \right)^{1/k} = \Theta(\sigma_{\mathbf{F}} \sqrt{km}^{-2/k}), \quad (86)$$

where V_k denotes the volume of the k -dim. unit ball, $V_k^{-1/k} = \Theta(\sqrt{k})$, and $c, C > 0$ are absolute constants.

Proof. Recall that $\mathbf{f}_i \in \mathbb{R}^d$ is the i -th row of \mathbf{F} and $\mathbf{v}_i = \mathbf{U}^\top \mathbf{f}_i \in \mathbb{R}^k$, further $\gamma_{\text{out}} = \sqrt{2} \min_{i \neq j} \|\mathbf{v}_i - \mathbf{v}_j\|_2$ and $\mathbf{v}_{ij} = \mathbf{v}_i - \mathbf{v}_j = \mathbf{U}^\top (\mathbf{f}_i - \mathbf{f}_j) \in \mathbb{R}^k$. Moreover, define $Y := \sigma_{\mathbf{F}} Y_0$, which has a Lebesgue density bounded by $M := \sigma_{\mathbf{F}}^{-1} M_0$.

Step ①: Random mask concentration. Fix a pair (i, j) . For each $\ell \in [d]$, let

$$\xi_\ell := \mathbb{1}\{(B_{i\ell}, B_{j\ell}) \neq (0, 0)\} \in \{0, 1\}, \quad \mathbf{P} := \mathbf{U}^\top \text{Diag}(\xi) \mathbf{U} \in \mathbb{R}^{k \times k}. \quad (87)$$

Then, ξ_ℓ are i.i.d. Bernoulli variables with

$$\mathbb{E}[\xi_\ell] = \mathbb{P}(\xi_\ell = 1) = 1 - (1 - p)^2 = 2p - p^2 =: q. \quad (88)$$

Let $\mathbf{u}_\ell := \mathbf{U}^\top \mathbf{e}_\ell \in \mathbb{R}^k$, so that $\|\mathbf{u}_\ell\|_2^2 = (\mathbf{U}\mathbf{U}^\top)_{\ell\ell} \leq \nu(\mathcal{U})$. Then,

$$\mathbf{P} - q\mathbf{I}_k = \sum_{\ell=1}^d (\xi_\ell - q) \mathbf{u}_\ell \mathbf{u}_\ell^\top. \quad (89)$$

Set $\mathbf{A}_\ell := (\xi_\ell - q) \mathbf{u}_\ell \mathbf{u}_\ell^\top$. Then, $\mathbb{E}[\mathbf{A}_\ell] = 0$ and

$$\|\mathbf{A}_\ell\|_{\text{op}} = |\xi_\ell - q| \|\mathbf{u}_\ell\|_2^2 \leq \nu(\mathcal{U}) =: R. \quad (90)$$

Further,

$$\mathbf{A}_\ell^2 = (\xi_\ell - q)^2 (\mathbf{u}_\ell \mathbf{u}_\ell^\top)^2 = (\xi_\ell - q)^2 \|\mathbf{u}_\ell\|_2^2 \mathbf{u}_\ell \mathbf{u}_\ell^\top, \quad (91)$$

so, since $\mathbb{E}[(\xi_\ell - q)^2] = q(1 - q) \leq q$,

$$\mathbb{E}[\mathbf{A}_\ell^2] \preceq q \|\mathbf{u}_\ell\|_2^2 \mathbf{u}_\ell \mathbf{u}_\ell^\top \preceq q\nu(\mathcal{U}) \mathbf{u}_\ell \mathbf{u}_\ell^\top. \quad (92)$$

Because $\sum_{\ell=1}^d \mathbf{u}_\ell \mathbf{u}_\ell^\top = \mathbf{U}^\top \mathbf{U} = \mathbf{I}_k$, it follows that

$$\sum_{\ell=1}^d \mathbb{E}[\mathbf{A}_\ell^2] \preceq q\nu(\mathcal{U}) \mathbf{I}_k, \quad \text{hence} \quad \sigma^2 := \left\| \sum_{\ell=1}^d \mathbb{E}[\mathbf{A}_\ell^2] \right\|_{\text{op}} \leq q\nu(\mathcal{U}). \quad (93)$$

By the matrix Bernstein inequality (Tropp (2015) and Lem. C.7), there exists an absolute constant $C_0 > 0$ such that for all $\eta \in (0, 1)$, with probability at least $1 - \eta$,

$$\|\mathbf{P} - q\mathbf{I}_k\|_{\text{op}} \leq C_0 \left(\sqrt{\sigma^2 \log(2k/\eta)} + R \log(2k/\eta) \right) \leq C_0 \left(\sqrt{q\nu(\mathcal{U}) \log(2k/\eta)} + \nu(\mathcal{U}) \log(2k/\eta) \right). \quad (94)$$

Now set $\eta := \delta/(2m^2)$, let

$$\Lambda := \log(2k/\eta) = \log(4km^2/\delta), \quad (95)$$

and define the event

$$\mathcal{G} := \left\{ \|\mathbf{P} - q\mathbf{I}_k\|_{\text{op}} \leq \frac{q}{2} \right\}. \quad (96)$$

By (94) and the assumption $q \geq C(\sqrt{q\nu(\mathcal{U})\Lambda} + \nu(\mathcal{U})\Lambda)$, taken with $C := 2C_0$, we have

$$\mathbb{P}(\mathcal{G}) \geq 1 - \frac{\delta}{2m^2}. \quad (97)$$

On \mathcal{G} , we have $\mathbf{P} \succeq (q/2)\mathbf{I}_k$, and therefore

$$\det(\mathbf{P}) \geq (q/2)^k. \quad (98)$$

Step ②: Uniform conditional density bound for \mathbf{v}_{ij} . Fix $\xi^* \in \{0, 1\}^d$ such that $\mathbb{P}(\xi = \xi^*) > 0$ and

$$\left\| \mathbf{U}^\top \text{Diag}(\xi^*) \mathbf{U} - q\mathbf{I}_k \right\|_{\text{op}} \leq \frac{q}{2}. \quad (99)$$

Let

$$S := \{\ell \in [d] \mid \xi_\ell^* = 1\} =: \{\ell_1, \dots, \ell_{|S|}\}, \quad (100)$$

define

$$\zeta := ((\mathbf{f}_i - \mathbf{f}_j)_{\ell_1}, \dots, (\mathbf{f}_i - \mathbf{f}_j)_{\ell_{|S|}}) \in \mathbb{R}^{|S|}, \quad (101)$$

and let $\mathbf{U}_S \in \mathbb{R}^{|S| \times k}$ be the submatrix of \mathbf{U} with rows indexed by S . Then,

$$\mathbf{P}^* := \mathbf{U}^\top \text{Diag}(\xi^*) \mathbf{U} = \mathbf{U}_S^\top \mathbf{U}_S \succeq \frac{q}{2} \mathbf{I}_k. \quad (102)$$

In particular, \mathbf{P}^* is positive definite and $\text{rk}(\mathbf{U}_S^\top) = k$. For each $\ell \in [d]$, write $X_\ell := (\mathbf{f}_i - \mathbf{f}_j)_\ell = B_{i\ell}Y_{i\ell} - B_{j\ell}Y_{j\ell}$ and let $\mathcal{E}_\ell := \{\xi_\ell = \xi_\ell^*\}$. Then, $\{\xi = \xi^*\} = \bigcap_{\ell=1}^d \mathcal{E}_\ell$, and each \mathcal{E}_ℓ is measurable w.r.t. $(B_{i\ell}, B_{j\ell})$. Since the families $(B_{i\ell}, B_{j\ell}, Y_{i\ell}, Y_{j\ell})$ are independent across $\ell \in [d]$, it follows that for arbitrary Borel sets $A_1, \dots, A_{|S|} \subseteq \mathbb{R}$,

$$\mathbb{P} \left(\bigcap_{r=1}^{|S|} \{X_{\ell_r} \in A_r\} \mid \xi = \xi^* \right) = \frac{\mathbb{P} \left(\bigcap_{r=1}^{|S|} \{X_{\ell_r} \in A_r\} \cap \bigcap_{\ell=1}^d \mathcal{E}_\ell \right)}{\mathbb{P}(\xi = \xi^*)} \quad (103)$$

$$= \frac{\prod_{r=1}^{|S|} \mathbb{P}(X_{\ell_r} \in A_r, \mathcal{E}_{\ell_r}) \prod_{\ell \notin S} \mathbb{P}(\mathcal{E}_\ell)}{\prod_{r=1}^{|S|} \mathbb{P}(\mathcal{E}_{\ell_r}) \prod_{\ell \notin S} \mathbb{P}(\mathcal{E}_\ell)} \quad (104)$$

$$= \prod_{r=1}^{|S|} \mathbb{P}(X_{\ell_r} \in A_r \mid \mathcal{E}_{\ell_r}). \quad (105)$$

Thus, the coordinates of ζ are independent under $\mathbb{P}(\cdot \mid \xi = \xi^*)$. Now fix $\ell \in S$ for which $\xi_\ell^* = 1$ and $\mathcal{E}_\ell = \{\xi_\ell = 1\}$. Let f_Y denote the Lebesgue density of Y , with $\|f_Y\|_{L^\infty(\lambda)} \leq M$, and let Y' be an independent copy of Y . For every Borel set $A \subseteq \mathbb{R}$,

$$\mathbb{P}(X_\ell \in A \mid \mathcal{E}_\ell) = \frac{p(1-p)}{q} \mathbb{P}(Y \in A) + \frac{p(1-p)}{q} \mathbb{P}(-Y \in A) + \frac{p^2}{q} \mathbb{P}(Y - Y' \in A). \quad (106)$$

Hence, under $\mathbb{P}(\cdot \mid \mathcal{E}_\ell)$, X_ℓ has a Lebesgue density which is a convex combination of the densities of Y , $-Y$, and $Y - Y'$. $-Y$ has density $x \mapsto f_{-Y}(x) = f_Y(-x)$ and $Y - Y'$ has density $f_{Y-Y'} = f_Y * f_{-Y}$. Thus,

$$\|f_{Y-Y'}\|_{L^\infty(\lambda)} = \|f_Y * f_{-Y}\|_{L^\infty(\lambda)} \leq \underbrace{\|f_Y\|_{L^\infty(\lambda)}}_{\leq M} \underbrace{\|f_{-Y}\|_{L^1(\lambda)}}_{=1} \leq M. \quad (107)$$

Therefore, for every $\ell \in S$, the conditional distribution of X_ℓ given \mathcal{E}_ℓ admits a Lebesgue density bounded by M . Together with (105), this shows that under $\mathbb{P}(\cdot \mid \xi = \xi^*)$, $\zeta \in \mathbb{R}^{|S|}$ has independent coordinates, each with Lebesgue density bounded by M . Now set $\mathbf{A} := \mathbf{U}_S^\top \in \mathbb{R}^{k \times |S|}$, so that $\mathbf{A}\mathbf{A}^\top = \mathbf{P}^*$. Since $\xi_\ell^* = 0$ implies $B_{i\ell} = B_{j\ell} = 0$, we have $X_\ell = 0$ a.s. under $\mathbb{P}(\cdot \mid \xi = \xi^*)$ for every $\ell \notin S$. Consequently,

$$\mathbf{v}_{ij} = \mathbf{U}^\top (\mathbf{f}_i - \mathbf{f}_j) = \mathbf{U}_S^\top \zeta = \mathbf{A}\zeta \quad \text{a.s. under } \mathbb{P}(\cdot \mid \xi = \xi^*). \quad (108)$$

Let $\Sigma := (\mathbf{A}\mathbf{A}^\top)^{1/2} = (\mathbf{P}^*)^{1/2}$ and $\mathbf{Q} := \Sigma^{-1}\mathbf{A} \in \mathbb{R}^{k \times |S|}$. Then, $\mathbf{Q}\mathbf{Q}^\top = \mathbf{I}_k$. Writing $\mathbf{\Pi} := \mathbf{Q}^\top \mathbf{Q}$, $\mathbf{\Pi}$ is the orthogonal projection onto the k -dim. subspace $\text{im}(\mathbf{Q}^\top) \subseteq \mathbb{R}^{|S|}$, and $\mathbf{Q}\zeta = \mathbf{Q}\mathbf{\Pi}\zeta$. Let $\lambda_{\text{im}(\mathbf{Q}^\top)}$ denote the k -dimensional Lebesgue measure on the subspace $\text{im}(\mathbf{Q}^\top)$. By [Rudelson & Vershynin \(2015, Theorem 1.1\)](#), applied under the conditional probability measure $\mathbb{P}(\cdot | \xi = \xi^*)$ to ζ and $\mathbf{\Pi}$, the random vector $\mathbf{\Pi}\zeta$ admits a density $f_{\mathbf{\Pi}\zeta | \xi = \xi^*}$ w.r.t. $\lambda_{\text{im}(\mathbf{Q}^\top)}$ satisfying

$$\|f_{\mathbf{\Pi}\zeta | \xi = \xi^*}\|_{L^\infty(\lambda_{\text{im}(\mathbf{Q}^\top)})} \leq (C_1 M)^k \quad (109)$$

for an absolute constant $C_1 > 0$. Since the restriction $\mathbf{Q}|_{\text{im}(\mathbf{Q}^\top)} : \text{im}(\mathbf{Q}^\top) \rightarrow \mathbb{R}^k$ is an isometry, it preserves the k -dim. Lebesgue measure λ^k , and thus, $\mathbf{Q}\zeta$ admits a Lebesgue density $f_{\mathbf{Q}\zeta | \xi = \xi^*}$ on \mathbb{R}^k with

$$\|f_{\mathbf{Q}\zeta | \xi = \xi^*}\|_{L^\infty(\lambda^k)} \leq (C_1 M)^k. \quad (110)$$

Since $\mathbf{v}_{ij} = \mathbf{A}\zeta = \Sigma(\mathbf{Q}\zeta)$ under $\mathbb{P}(\cdot | \xi = \xi^*)$, a change of variables for the invertible linear map $\Sigma : \mathbb{R}^k \rightarrow \mathbb{R}^k$ shows that \mathbf{v}_{ij} admits a Lebesgue density $f_{\mathbf{v}_{ij} | \xi = \xi^*}$ on \mathbb{R}^k given by

$$f_{\mathbf{v}_{ij} | \xi = \xi^*}(\mathbf{x}) = \frac{f_{\mathbf{Q}\zeta | \xi = \xi^*}(\Sigma^{-1}\mathbf{x})}{|\det \Sigma|}, \quad \mathbf{x} \in \mathbb{R}^k. \quad (111)$$

Therefore,

$$\|f_{\mathbf{v}_{ij} | \xi = \xi^*}\|_{L^\infty(\lambda^k)} \leq \frac{\|f_{\mathbf{Q}\zeta | \xi = \xi^*}\|_{L^\infty(\lambda^k)}}{|\det \Sigma|} \leq \frac{(C_1 M)^k}{\sqrt{\det(\mathbf{A}\mathbf{A}^\top)}} = \frac{(C_1 M)^k}{\sqrt{\det(\mathbf{P}^*)}} \leq (C_1 M)^k \left(\frac{2}{q}\right)^{k/2}. \quad (112)$$

Consequently, for every $t \geq 0$,

$$\mathbb{P}(\|\mathbf{v}_{ij}\|_2 \leq t | \xi = \xi^*) = \int_{B_t^k} f_{\mathbf{v}_{ij} | \xi = \xi^*} d\lambda^k \quad (113)$$

$$\leq \lambda^k(B_t^k) \|f_{\mathbf{v}_{ij} | \xi = \xi^*}\|_{L^\infty(\lambda^k)} \quad (114)$$

$$\leq V_k t^k (C_1 M)^k \left(\frac{2}{q}\right)^{k/2} = \left(C_2 V_k^{1/k} \frac{Mt}{\sqrt{q}}\right)^k, \quad (115)$$

where we set $C_2 := \sqrt{2} C_1$.

The bound (115) is uniform over all $\xi^* \in \{0, 1\}^d$ with $\mathbb{P}(\xi = \xi^*) > 0$ and $\|\mathbf{U}^\top \text{Diag}(\xi^*) \mathbf{U} - q \mathbf{I}_k\|_{\text{op}} \leq q/2$. Since the event \mathcal{G} is measurable w.r.t. ξ , and ξ takes values in the finite set $\{0, 1\}^d$, conditioning on \mathcal{G} yields

$$\mathbb{P}(\|\mathbf{v}_{ij}\|_2 \leq t | \mathcal{G}) = \sum_{\substack{\xi^* \in \{0, 1\}^d \\ \mathbb{P}(\xi = \xi^* | \mathcal{G}) > 0}} \mathbb{P}(\xi = \xi^* | \mathcal{G}) \mathbb{P}(\|\mathbf{v}_{ij}\|_2 \leq t | \xi = \xi^*) \leq \left(C_2 V_k^{1/k} \frac{Mt}{\sqrt{q}}\right)^k. \quad (116)$$

Combining (116) with $\mathbb{P}(\mathcal{G}^c) \leq \delta/(2m^2)$ from Step ①, we obtain

$$\mathbb{P}(\|\mathbf{v}_{ij}\|_2 \leq t) \leq \mathbb{P}(\mathcal{G}^c) + \mathbb{P}(\|\mathbf{v}_{ij}\|_2 \leq t | \mathcal{G}) \leq \frac{\delta}{2m^2} + \left(C_2 V_k^{1/k} \frac{Mt}{\sqrt{q}}\right)^k. \quad (117)$$

Step ③: Union bound over all pairs. Choose

$$t := c \frac{\sqrt{q}}{M} V_k^{-1/k} \left(\frac{\delta}{m^2}\right)^{1/k} \quad (118)$$

with $c := (2C_2)^{-1}$. Then,

$$\left(C_2 V_k^{1/k} \frac{Mt}{\sqrt{q}}\right)^k = (C_2 c)^k \frac{\delta}{m^2} = \left(\frac{1}{2}\right)^k \frac{\delta}{m^2} \leq \frac{1}{2} \frac{\delta}{m^2}, \quad (119)$$

so (117) gives for every fixed pair (i, j)

$$\mathbb{P}(\|\mathbf{v}_{ij}\|_2 \leq t) \leq \frac{\delta}{2m^2} + \frac{1}{2} \frac{\delta}{m^2} = \frac{\delta}{m^2}. \quad (120)$$

By a union bound over the $\binom{m}{2} \leq m^2/2$ pairs,

$$\mathbb{P}\left(\min_{i \neq j} \|\mathbf{v}_{ij}\| \leq t\right) \leq \sum_{i < j} \mathbb{P}(\|\mathbf{v}_{ij}\|_2 \leq t) \leq \frac{m^2}{2} \cdot \frac{\delta}{m^2} \leq \delta. \quad (121)$$

Hence, recalling that $M = \sigma_{\mathbf{F}}^{-1} M_0$,

$$\gamma_{\text{out}} = \sqrt{2} \min_{i \neq j} \|\mathbf{v}_{ij}\| \geq \sqrt{2}t = \sqrt{2}c \frac{\sqrt{q}}{M_0} \sigma_{\mathbf{F}} V_k^{-1/k} \left(\frac{\delta}{m^2} \right)^{1/k} \quad (122)$$

with probability at least $1 - \delta$, which is precisely the first inequality of (86), absorbing $\sqrt{2}$ into c .

Step ④: Stirling approximation. Fixing δ, q, M_0 and using the Stirling formula

$$\Gamma(x+1) \sim \sqrt{2\pi x} \left(\frac{x}{e} \right)^x \quad \text{as } x \rightarrow \infty, \quad (123)$$

we obtain

$$V_k^{-1/k} = \left(\frac{\Gamma(k/2+1)}{\pi^{k/2}} \right)^{1/k} \sim \left(\frac{\sqrt{\pi k} \left(\frac{k}{2e} \right)^{k/2}}{\pi^{k/2}} \right)^{1/k} = \underbrace{(\pi k)^{1/(2k)}}_{\rightarrow 1, k \rightarrow \infty} \sqrt{\frac{k}{2\pi e}} \sim \sqrt{\frac{k}{2\pi e}}, \quad (124)$$

where $\Gamma(x) := \int_0^\infty t^{x-1} e^{-t} dt$ denotes the gamma function. Hence, $V_k^{-1/k} = \Theta(\sqrt{k})$, and absorbing the constant $(2\pi e)^{-1/2}$ into c yields that the r.h.s. of (122) is $\Theta(\sigma_{\mathbf{F}} \sqrt{k} m^{-2/k})$ in $m, k, \sigma_{\mathbf{F}}$. \square

C.3. Input Reindexings

Theorem C.13 (Global input sensitivity). *In the setup of Thm. C.10, fix an input reindexing $\tau \in S_d$. Assume that the concentration event Assumption C.9 holds both for \mathbf{S}_i and $\mathbf{S}_i^\tau := \mathbf{U}^\top \mathbf{P}_\tau^\top \text{Diag}(\mathbf{d}_i)^2 \mathbf{P}_\tau \mathbf{U}$.⁴ Also assume that $\delta_\tau^{\text{in}} \neq 0$ and $\rho_\tau^{\text{in}} := \|\mathbf{R}\|_F / \|\delta_\tau^{\text{in}}\|_F \leq 1$. Then,*

$$\Delta_\tau^{\text{in}}(\mathbf{H}) = \left(1 \pm \mathcal{O}(\rho_\tau^{\text{in}} + \varepsilon \mu_{\mathbf{D}}^{-1}) \right) \cdot \mu_{\mathbf{D}}^{-1} \|\delta_\tau^{\text{in}}\|_F^2. \quad (125)$$

Proof. The argument is partly analogous to the proof of Thm. C.10. For each $i \in [m]$, let $\mathbf{a}_i := \mathbf{v}_i + \mathbf{R}_{i,\cdot} \in \mathbb{R}^k$ denote the coefficient vector realized by the i -th neuron on \mathcal{X} . Under the input permutation τ , the transformed function is evaluated on the reindexed support $\mathbf{P}_\tau \mathcal{X} \subseteq \mathbf{P}_\tau \mathcal{U}$, which has $\mathbf{P}_\tau \mathbf{U}$ as an orthonormal basis. W.r.t. this basis, the same transformed neuron is still represented by the coefficient vector \mathbf{a}_i , while the corresponding projected center becomes

$$\mathbf{v}_i^\tau := (\mathbf{F} \mathbf{P}_\tau \mathbf{U})_{i,\cdot} = \mathbf{v}_i + (\delta_\tau^{\text{in}})_{i,\cdot}. \quad (126)$$

Hence,

$$\Delta_\tau^{\text{in}}(\mathbf{H}, \mathcal{X}) = \sum_{i=1}^m \left(\|\mathbf{R}_{i,\cdot} - (\delta_\tau^{\text{in}})_{i,\cdot}\|_{(\mathbf{S}_i^\tau)^{-1}}^2 - \|\mathbf{R}_{i,\cdot}\|_{\mathbf{S}_i^{-1}}^2 \right). \quad (127)$$

Since $\varepsilon < \mu_{\mathbf{D}}$, the assumptions imply

$$(\mu_{\mathbf{D}} - \varepsilon) \mathbf{I}_k \preceq \mathbf{S}_i, \mathbf{S}_i^\tau \preceq (\mu_{\mathbf{D}} + \varepsilon) \mathbf{I}_k \quad \text{for all } i \in [m]. \quad (128)$$

Thus, both \mathbf{S}_i and \mathbf{S}_i^τ are invertible, and

$$\|\mathbf{S}_i^{-1} - \mu_{\mathbf{D}}^{-1} \mathbf{I}_k\|_{\text{op}}, \|(\mathbf{S}_i^\tau)^{-1} - \mu_{\mathbf{D}}^{-1} \mathbf{I}_k\|_{\text{op}} \leq \frac{\varepsilon}{\mu_{\mathbf{D}}(\mu_{\mathbf{D}} - \varepsilon)} =: \beta. \quad (129)$$

Therefore, for every $\mathbf{a} \in \mathbb{R}^k$,

$$\left| \|\mathbf{a}\|_{\mathbf{S}_i^{-1}}^2 - \mu_{\mathbf{D}}^{-1} \|\mathbf{a}\|_2^2 \right| \leq \beta \|\mathbf{a}\|_2^2, \quad \left| \|\mathbf{a}\|_{(\mathbf{S}_i^\tau)^{-1}}^2 - \mu_{\mathbf{D}}^{-1} \|\mathbf{a}\|_2^2 \right| \leq \beta \|\mathbf{a}\|_2^2. \quad (130)$$

Applying this to each term in (127) yields

$$\Delta_\tau^{\text{in}}(\mathbf{H}, \mathcal{X}) = \mu_{\mathbf{D}}^{-1} \left(\|\mathbf{R} - \delta_\tau^{\text{in}}\|_F^2 - \|\mathbf{R}\|_F^2 \right) + E_\tau, \quad (131)$$

with

$$|E_\tau| \leq \beta \left(\|\mathbf{R} - \delta_\tau^{\text{in}}\|_F^2 + \|\mathbf{R}\|_F^2 \right). \quad (132)$$

By Cauchy–Schwarz,

$$|E_\tau| \leq \beta \left(\|\delta_\tau^{\text{in}}\|_F^2 + 2\|\delta_\tau^{\text{in}}\|_F \|\mathbf{R}\|_F + 2\|\mathbf{R}\|_F^2 \right). \quad (133)$$

⁴When \mathbf{D} has i.i.d. entries, $\mathbf{S}_i^\tau = \mathbf{U}^\top \mathbf{P}_\tau^\top \text{Diag}(\mathbf{d}_i)^2 \mathbf{P}_\tau \mathbf{U} \stackrel{d}{=} \mathbf{S}_i$.

Moreover,

$$\|\mathbf{R} - \delta_\tau^{\text{in}}\|_F^2 - \|\mathbf{R}\|_F^2 = \|\delta_\tau^{\text{in}}\|_F^2 - 2\langle \delta_\tau^{\text{in}}, \mathbf{R} \rangle_F, \quad (134)$$

and by Cauchy–Schwarz,

$$|\langle \delta_\tau^{\text{in}}, \mathbf{R} \rangle_F| \leq \|\delta_\tau^{\text{in}}\|_F \|\mathbf{R}\|_F. \quad (135)$$

Combining (131), (133), (134), and (135) gives

$$|\Delta_\tau^{\text{in}}(\mathbf{H}, \mathcal{X}) - \mu_{\mathbf{D}}^{-1} \|\delta_\tau^{\text{in}}\|_F^2| \leq \frac{2}{\mu_{\mathbf{D}}} \|\delta_\tau^{\text{in}}\|_F \|\mathbf{R}\|_F + \frac{\varepsilon}{\mu_{\mathbf{D}}(\mu_{\mathbf{D}} - \varepsilon)} \left(\|\delta_\tau^{\text{in}}\|_F^2 + 2\|\delta_\tau^{\text{in}}\|_F \|\mathbf{R}\|_F + 2\|\mathbf{R}\|_F^2 \right). \quad (136)$$

Finally, divide (136) by $\mu_{\mathbf{D}}^{-1} \|\delta_\tau^{\text{in}}\|_F^2$ to obtain

$$\left| \frac{\mu_{\mathbf{D}} \Delta_\tau^{\text{in}}(\mathbf{H}, \mathcal{X})}{\|\delta_\tau^{\text{in}}\|_F^2} - 1 \right| \leq 2\rho_\tau^{\text{in}} + \frac{\varepsilon}{\mu_{\mathbf{D}} - \varepsilon} (1 + 2\rho_\tau^{\text{in}} + 2(\rho_\tau^{\text{in}})^2). \quad (137)$$

If $\rho_\tau^{\text{in}} \leq 1$ and $\varepsilon \leq \mu_{\mathbf{D}}/2$, then

$$\frac{\varepsilon}{\mu_{\mathbf{D}} - \varepsilon} \leq \frac{2\varepsilon}{\mu_{\mathbf{D}}}, \quad 1 + 2\rho_\tau^{\text{in}} + 2(\rho_\tau^{\text{in}})^2 \leq 5. \quad (138)$$

Thus the r.h.s. is $\mathcal{O}(\rho_\tau^{\text{in}} + \varepsilon\mu_{\mathbf{D}}^{-1})$, and therefore

$$\Delta_\tau^{\text{in}}(\mathbf{H}, \mathcal{X}) = \left(1 \pm \mathcal{O}(\rho_\tau^{\text{in}} + \varepsilon\mu_{\mathbf{D}}^{-1})\right) \cdot \mu_{\mathbf{D}}^{-1} \|\delta_\tau^{\text{in}}\|_F^2. \quad (139)$$

Corollary C.14 (Minimum input transposition cost). *Assume that the assumptions of Thm. C.13 hold for every transposition $(ab) \in S_d$, and that $\rho_{(ab)}^{\text{in}} \leq \rho^{\text{in}} \leq 1$ for some ρ^{in} uniformly in $a \neq b$. Then,*

$$\min_{a \neq b} \Delta_{(ab)}^{\text{in}}(\mathbf{H}, \mathcal{X}) = \left(1 \pm \mathcal{O}(\rho^{\text{in}} + \varepsilon\mu_{\mathbf{D}}^{-1})\right) \cdot \mu_{\mathbf{D}}^{-1} \gamma_{\text{in}}^2. \quad (140)$$

Proof. Since $\rho_{(ab)}^{\text{in}} \leq \rho^{\text{in}}$ for every transposition (ab) , by Thm. C.13,

$$\Delta_{(ab)}^{\text{in}}(\mathbf{H}, \mathcal{X}) = \left(1 \pm \mathcal{O}(\rho^{\text{in}} + \varepsilon\mu_{\mathbf{D}}^{-1})\right) \cdot \mu_{\mathbf{D}}^{-1} \|\delta_{(ab)}^{\text{in}}\|_F^2 \quad (141)$$

uniformly in $a \neq b$. Taking the minimum over all transpositions gives

$$\min_{a \neq b} \Delta_{(ab)}^{\text{in}}(\mathbf{H}, \mathcal{X}) = \left(1 \pm \mathcal{O}(\rho^{\text{in}} + \varepsilon\mu_{\mathbf{D}}^{-1})\right) \cdot \mu_{\mathbf{D}}^{-1} \gamma_{\text{in}}^2, \quad (142)$$

which is exactly the claim. \square

Theorem C.15 (Asymptotics of γ_{in}). *Let $\mathbf{F}_{:, \ell} \in \mathbb{R}^m$ denote the ℓ -th column of \mathbf{F} . Let $a \neq b \in [d]$. Then,*

$$\|\delta_{(ab)}^{\text{in}}\|_F = \|\mathbf{F}_{:, b} - \mathbf{F}_{:, a}\|_2 \|\mathbf{U}_{b,:} - \mathbf{U}_{a,:}\|_2. \quad (143)$$

If further \mathbf{F} has i.i.d. centered subgaussian entries with variance $\sigma_{\mathbf{F}}^2$, subgaussian norm uniformly bounded by $C\sigma_{\mathbf{F}}$, and $m \gtrsim C^4 \log(d^2/\delta)$, then with probability at least $1 - \delta$,

$$\gamma_{\text{in}} = \Theta(\sigma_{\mathbf{F}} \sqrt{m}) \cdot \min_{a \neq b} \|\mathbf{U}_{b,:} - \mathbf{U}_{a,:}\|_2. \quad (144)$$

Proof. **Step ①: Proof of (143).** Fix $a \neq b$. For a transposition (ab) , we have for any $\mathbf{x} \in \mathbb{R}^d$

$$(\mathbf{P}_{(ab)} - \mathbf{I}_d)\mathbf{x} = (x_b - x_a)(\mathbf{e}_a - \mathbf{e}_b). \quad (145)$$

Thus,

$$(\delta_{(ab)}^{\text{in}}) = \mathbf{F}(\mathbf{P}_{(ab)} - \mathbf{I}_d)\mathbf{U} = \mathbf{F}(\mathbf{e}_a - \mathbf{e}_b)(\mathbf{U}_{b,:} - \mathbf{U}_{a,:}) = (\mathbf{F}_{:, a} - \mathbf{F}_{:, b})(\mathbf{U}_{b,:} - \mathbf{U}_{a,:}), \quad (146)$$

which is a matrix of rank ≤ 1 . Using $\|\mathbf{xy}^\top\|_F = \|\mathbf{x}\|_2 \|\mathbf{y}\|_2$ gives (143).

Step ②: Proof of (144). We assume that the entries of \mathbf{F} are i.i.d., centered, have variance $\sigma_{\mathbf{F}}^2$, and subgaussian norm $K := \|\mathbf{F}_{ij}\|_{\psi_2}$. Fix $a < b$ and define $Z_i := \mathbf{F}_{ib} - \mathbf{F}_{ia}$, for $i \in [m]$. Then, Z_1, \dots, Z_m are independent and centered, with

$$\mathbb{E}Z_i^2 = \mathbb{E}\mathbf{F}_{ib}^2 + \mathbb{E}\mathbf{F}_{ia}^2 = 2\sigma_{\mathbf{F}}^2. \quad (147)$$

Moreover, by the triangle inequality for the subgaussian norm,

$$\|Z_i\|_{\psi_2} \leq \|\mathbf{F}_{ib}\|_{\psi_2} + \|\mathbf{F}_{ia}\|_{\psi_2} = 2K. \quad (148)$$

Therefore, for $Y_i := Z_i^2 - \mathbb{E}Z_i^2$, the random variables Y_1, \dots, Y_m are independent, centered, and subexponential. In particular, there exists a universal constant $C_0 > 0$ such that

$$\|Y_i\|_{\psi_1} \leq C_0 \|Z_i\|_{\psi_2}^2 \leq 4C_0 K^2. \quad (149)$$

Hence, by the Bernstein inequality for sums of independent subexponential random variables (e.g., [Vershynin \(2018, Thm. 2.8.2\)](#)), and absorbing the universal constant in the bound on $\|Y_i\|_{\psi_1}$, there exists a universal constant $c_1 > 0$ such that for every $t \geq 0$,

$$\mathbb{P}\left(\left|\sum_{i=1}^m Y_i\right| \geq t\right) \leq 2 \exp\left(-c_1 \min\left\{\frac{t^2}{mK^4}, \frac{t}{K^2}\right\}\right). \quad (150)$$

Set $L := \log(d^2/\delta)$ and choose

$$t := C_1 K^2 (\sqrt{mL} + L), \quad (151)$$

where $C_1 > 0$ is a universal constant yet to be determined. Then,

$$\frac{t}{K^2} = C_1 (\sqrt{mL} + L) \geq C_1 L, \quad \frac{t^2}{mK^4} = C_1^2 \frac{(\sqrt{mL} + L)^2}{m} \geq C_1^2 L. \quad (152)$$

Thus,

$$\min\left\{\frac{t^2}{mK^4}, \frac{t}{K^2}\right\} \geq \min\{C_1^2, C_1\}L. \quad (153)$$

Choosing C_1 large enough so that $c_1 \min\{C_1^2, C_1\} \geq 1$, we obtain

$$\mathbb{P}\left(\left|\sum_{i=1}^m Y_i\right| \geq t\right) \leq 2 \exp(-c_1 \min\{C_1^2, C_1\} \cdot L) \leq 2e^{-L} = \frac{2\delta}{d^2}. \quad (154)$$

Since

$$\sum_{i=1}^m Y_i = \sum_{i=1}^m Z_i^2 - 2m\sigma_{\mathbf{F}}^2 = \|\mathbf{F}_{:,b} - \mathbf{F}_{:,a}\|_2^2 - 2m\sigma_{\mathbf{F}}^2, \quad (155)$$

we have shown that, for this fixed pair $a < b$,

$$\mathbb{P}\left(\left|\|\mathbf{F}_{:,b} - \mathbf{F}_{:,a}\|_2^2 - 2m\sigma_{\mathbf{F}}^2\right| \geq t\right) \leq \frac{2\delta}{d^2}. \quad (156)$$

Taking a union bound over the $\binom{d}{2} \leq d^2/2$ pairs, we conclude that, with probability at least $1 - \delta$, simultaneously for all $a < b$,

$$\left|\|\mathbf{F}_{:,b} - \mathbf{F}_{:,a}\|_2^2 - 2m\sigma_{\mathbf{F}}^2\right| \leq C_1 K^2 \left(\sqrt{m \log(d^2/\delta)} + \log(d^2/\delta)\right). \quad (157)$$

On this event, if $t < 2m\sigma_{\mathbf{F}}^2$, then for every $a \neq b$,

$$\sqrt{2m\sigma_{\mathbf{F}}^2 - t} \leq \|\mathbf{F}_{:,b} - \mathbf{F}_{:,a}\|_2 \leq \sqrt{2m\sigma_{\mathbf{F}}^2 + t}. \quad (158)$$

Combining this with (143) gives

$$\sqrt{2m\sigma_{\mathbf{F}}^2 - t} \min_{a \neq b} \|\mathbf{U}_{a,:} - \mathbf{U}_{b,:}\|_2 \leq \gamma_{\text{in}} \leq \sqrt{2m\sigma_{\mathbf{F}}^2 + t} \min_{a \neq b} \|\mathbf{U}_{a,:} - \mathbf{U}_{b,:}\|_2. \quad (159)$$

It remains to translate (159) into the asymptotic statement. By assumption, $K = \|\mathbf{F}_{ij}\|_{\psi_2} \leq C\sigma_{\mathbf{F}}$. Increasing C if necessary, we may assume $C \geq 1$. We obtain

$$t \leq C_1 C^2 \sigma_{\mathbf{F}}^2 (\sqrt{mL} + L). \quad (160)$$

We now choose a universal constant $C_2 \geq \max\{4C_1^2, 2C_1\}$. If $m \geq C_2 C^4 L$, then

$$C_1 C^2 \sqrt{mL} \leq m/2, \quad C_1 C^2 L \leq m/2. \quad (161)$$

Indeed, the first inequality follows from $C_1 C^2 \sqrt{mL} \leq m/2 \Leftrightarrow m \geq 4C_1^2 C^4 L$; the second follows since $C \geq 1$ and $C_2 \geq 2C_1$. Therefore, $t \leq m\sigma_{\mathbf{F}}^2$, and substituting this into (159) gives

$$\sigma_{\mathbf{F}} \sqrt{m} \min_{a \neq b} \|\mathbf{U}_{b,:} - \mathbf{U}_{a,:}\|_2 \leq \gamma_{\text{in}} \leq \sqrt{3} \sigma_{\mathbf{F}} \sqrt{m} \min_{a \neq b} \|\mathbf{U}_{b,:} - \mathbf{U}_{a,:}\|_2. \quad (162)$$

Equivalently,

$$\gamma_{\text{in}} = \Theta(\sigma_{\mathbf{F}} \sqrt{m}) \cdot \min_{a \neq b} \|\mathbf{U}_{b,:} - \mathbf{U}_{a,:}\|_2 \quad (163)$$

with probability at least $1 - \delta$. \square

C.4. Identifiability vs. Feature Learning

Let $\mathbf{x} \sim \mathbb{P} := \mathcal{N}(0, \mathbf{I}_d)$ and let $\eta = \text{ReLU}$. For any measurable $h : \mathbb{R}^d \rightarrow \mathbb{R}$, write $\|h\|_{L^2(\mathbb{P})}^2 := \mathbb{E}_{\mathbf{x}}[h(\mathbf{x})^2]$. For the symmetric full feature learning regime, consider a two-layer network of width m as

$$f_{\text{sym}}(\mathbf{x}) := \boldsymbol{\alpha}^\top \eta(\mathbf{W}\mathbf{x} + \mathbf{b}), \quad (164)$$

where $\boldsymbol{\alpha} \in \mathbb{R}^m$, $\mathbf{W} \in \mathbb{R}^{m \times d}$, and $\mathbf{b} \in \mathbb{R}^m$ are trainable and η acts elementwise. For an asymmetric model, fix centers $\mathbf{F} \in \mathbb{R}^{m \times d}$ and consider

$$f_{\text{asym}}(\mathbf{x}) := \boldsymbol{\alpha}^\top \eta((\mathbf{F} + \mathbf{W})\mathbf{x}), \quad (165)$$

where $\boldsymbol{\alpha} \in \mathbb{R}^m$ and $\mathbf{W} \in \mathbb{R}^{m \times d}$ are trainable.

Lemma C.16 (Lipschitz stability of ridge features). *Let $\mathbf{x} \sim \mathbb{P} = \mathcal{N}(0, \mathbf{I}_d)$ and $\eta = \text{ReLU}$. Then for any $\mathbf{f} \in \mathbb{R}^d$ and any perturbation $\mathbf{w} \in \mathbb{R}^d$,*

$$\|\eta((\mathbf{f} + \mathbf{w})^\top \mathbf{x}) - \eta(\mathbf{f}^\top \mathbf{x})\|_{L^2(\mathbb{P})} \leq \|\mathbf{w}\|_2. \quad (166)$$

Proof. Since $\eta = \text{ReLU}$ is 1-Lipschitz, we have pointwise for all $\mathbf{x} \in \mathbb{R}^d$,

$$|\eta((\mathbf{f} + \mathbf{w})^\top \mathbf{x}) - \eta(\mathbf{f}^\top \mathbf{x})| \leq |(\mathbf{f} + \mathbf{w})^\top \mathbf{x} - \mathbf{f}^\top \mathbf{x}| = |\mathbf{w}^\top \mathbf{x}|. \quad (167)$$

Taking the expectation under \mathbf{x} yields

$$\|\eta((\mathbf{f} + \mathbf{w})^\top \mathbf{x}) - \eta(\mathbf{f}^\top \mathbf{x})\|_{L^2(\mathbb{P})}^2 \leq \mathbb{E}_{\mathbf{x}}[(\mathbf{w}^\top \mathbf{x})^2] = \mathbf{w}^\top \mathbb{E}_{\mathbf{x}}[\mathbf{x}\mathbf{x}^\top] \mathbf{w} = \|\mathbf{w}\|_2^2, \quad (168)$$

since $\mathbb{E}_{\mathbf{x}}[\mathbf{x}\mathbf{x}^\top] = \mathbf{I}_d$. \square

Lemma C.17 (Asymmetric network stability under bounded perturbations). *Let $\mathbf{x} \sim \mathbb{P} = \mathcal{N}(0, \mathbf{I}_d)$ and $\eta = \text{ReLU}$. Fix $\mathbf{F} \in \mathbb{R}^{m \times d}$ and consider the following two networks with the same output weights $\boldsymbol{\alpha} \in \mathbb{R}^m$*

$$f_{\text{asym}}(\mathbf{x}) := \boldsymbol{\alpha}^\top \eta((\mathbf{F} + \mathbf{W})\mathbf{x}), \quad f_0(\mathbf{x}) := \boldsymbol{\alpha}^\top \eta(\mathbf{F}\mathbf{x}), \quad (169)$$

where $\mathbf{W} \in \mathbb{R}^{m \times d}$ is arbitrary. Then,

$$\|f_{\text{asym}} - f_0\|_{L^2(\mathbb{P})} \leq \|\boldsymbol{\alpha}\|_2 \|\mathbf{W}\|_F. \quad (170)$$

Proof. Write $\mathbf{f}_i \in \mathbb{R}^d$ and $\mathbf{w}_i \in \mathbb{R}^d$ for the i -th rows of \mathbf{F} and \mathbf{W} , respectively, so that

$$f_{\text{asym}}(\mathbf{x}) - f_0(\mathbf{x}) = \sum_{i=1}^m \alpha_i \left(\eta((\mathbf{f}_i + \mathbf{w}_i)^\top \mathbf{x}) - \eta(\mathbf{f}_i^\top \mathbf{x}) \right). \quad (171)$$

Let $\delta_i(\mathbf{x}) := \eta((\mathbf{f}_i + \mathbf{w}_i)^\top \mathbf{x}) - \eta(\mathbf{f}_i^\top \mathbf{x})$. By Cauchy–Schwarz applied pointwise in \mathbf{x} ,

$$|f_{\text{asym}}(\mathbf{x}) - f_0(\mathbf{x})| \leq \|\boldsymbol{\alpha}\|_2 \|\boldsymbol{\delta}(\mathbf{x})\|_2, \quad \boldsymbol{\delta}(\mathbf{x}) := (\delta_1(\mathbf{x}), \dots, \delta_m(\mathbf{x}))^\top. \quad (172)$$

Squaring and taking expectation gives

$$\|f_{\text{asym}} - f_0\|_{L^2(\mathbb{P})}^2 \leq \|\boldsymbol{\alpha}\|_2^2 \mathbb{E}_{\mathbf{x}}[\|\boldsymbol{\delta}(\mathbf{x})\|_2^2] = \|\boldsymbol{\alpha}\|_2^2 \sum_{i=1}^m \|\delta_i\|_{L^2(\mathbb{P})}^2. \quad (173)$$

Applying [Lem. C.16](#) to each $(\mathbf{f}_i, \mathbf{w}_i)$ yields $\|\delta_i\|_{L^2(\mathbb{P})} \leq \|\mathbf{w}_i\|_2$, hence

$$\|f_{\text{asym}} - f_0\|_{L^2(\mathbb{P})}^2 \leq \|\boldsymbol{\alpha}\|_2^2 \sum_{i=1}^m \|\mathbf{w}_i\|_2^2 = \|\boldsymbol{\alpha}\|_2^2 \|\mathbf{W}\|_F^2. \quad (174) \quad \square$$

Theorem C.18 (Random feature hardness for single ReLU neuron ([Yehudai & Shamir, 2019](#)), specialization of Theorem 4.2). *There exists a universal constant $c > 0$ such that for all $d > 40$, the following holds: For every $\mathbf{w}_\star \in \mathbb{R}^d$ with $\|\mathbf{w}_\star\|_2 = d^2$, there exists a bias $b_\star \in \mathbb{R}$ with $|b_\star| \leq 6d^3 + 1$ such that for any coefficients $\alpha_1, \dots, \alpha_m \in \mathbb{R}$, with probability at least $1 - \exp(-cd)$ over i.i.d. $\mathbf{b}_1, \dots, \mathbf{b}_m \sim \text{Unif}(\mathbb{S}^{d-1})$,*

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_d)} \left[\left(\sum_{i=1}^m \alpha_i \eta(\mathbf{b}_i^\top \mathbf{x}) - \eta(\mathbf{w}_\star^\top \mathbf{x} + b_\star) \right)^2 \right] \leq \frac{1}{50} \implies m \cdot \max_{i \in [m]} |\alpha_i| \geq \frac{1}{48d^2} \exp(cd). \quad (175)$$

Theorem C.19 (Hardness in center-dominated regime). *Let $\mathbf{x} \sim \mathbb{P} = \mathcal{N}(0, \mathbf{I}_d)$, $\eta = \text{ReLU}$, and fix $d > 40$, $m \in \mathbb{N}$. Draw $\mathbf{F} \in \mathbb{R}^{m \times d}$ with i.i.d. $\mathbf{f}_i \sim \text{Unif}(\mathbb{S}^{d-1})$. The following holds for every $\mathbf{w}_\star \in \mathbb{R}^d$ with $\|\mathbf{w}_\star\|_2 = d^2$: There exists a bias $b_\star \in \mathbb{R}$ with $|b_\star| \leq 6d^3 + 1$ such that for all $\boldsymbol{\alpha} \in \mathbb{R}^m$ and $\mathbf{W} \in \mathbb{R}^{m \times d}$ satisfying $\|\boldsymbol{\alpha}\|_\infty < \frac{1}{48md^2} \exp(cd)$ and*

$\|\boldsymbol{\alpha}\|_2\|\mathbf{W}\|_F \leq \frac{1}{10\sqrt{2}}$, with probability at least $1 - \exp(-cd)$ over sampling \mathbf{F} , the asymmetric network $f_{\text{asym}}(\mathbf{x}) := \boldsymbol{\alpha}^\top \eta((\mathbf{F} + \mathbf{W})\mathbf{x})$ incurs the population error lower bound

$$\|f_{\text{asym}} - \eta(\mathbf{w}_*^\top(\cdot) + b_*)\|_{L^2(\mathbb{P})} > 1/(10\sqrt{2}). \quad (176)$$

Proof. Work on the event \mathcal{E} (of probability at least $1 - \exp(-cd)$) on which [Thm. C.18](#) holds for $\{\mathbf{f}_i\}_{i=1}^m$. Fix any $\mathbf{w}_* \in \mathbb{R}^d$ with $\|\mathbf{w}_*\|_2 = d^2$. On \mathcal{E} , [Thm. C.18](#) provides a bias $b_* \in \mathbb{R}$ with $|b_*| \leq 6d^3 + 1$ such that for every choice of coefficients $\alpha_1, \dots, \alpha_m \in \mathbb{R}$,

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} \left[\left(\sum_{i=1}^m \alpha_i \eta(\mathbf{f}_i^\top \mathbf{x}) - \eta(\mathbf{w}_*^\top \mathbf{x} + b_*) \right)^2 \right] \leq \frac{1}{50} \implies m \max_{i \in [m]} |\alpha_i| \geq \frac{1}{48d^2} \exp(cd). \quad (177)$$

Now fix arbitrary $\boldsymbol{\alpha} \in \mathbb{R}^m$ and $\mathbf{W} \in \mathbb{R}^{m \times d}$ satisfying

$$\|\boldsymbol{\alpha}\|_\infty < \frac{1}{48md^2} \exp(cd), \quad \|\boldsymbol{\alpha}\|_2\|\mathbf{W}\|_F \leq \frac{1}{10\sqrt{2}}. \quad (178)$$

Define

$$f_{\text{asym}}(\mathbf{x}) := \boldsymbol{\alpha}^\top \eta((\mathbf{F} + \mathbf{W})\mathbf{x}), \quad f_0(\mathbf{x}) := \boldsymbol{\alpha}^\top \eta(\mathbf{F}\mathbf{x}) = \sum_{i=1}^m \alpha_i \eta(\mathbf{f}_i^\top \mathbf{x}). \quad (179)$$

By [Lem. C.17](#),

$$\|f_{\text{asym}} - f_0\|_{L^2(\mathbb{P})} \leq \|\boldsymbol{\alpha}\|_2\|\mathbf{W}\|_F \leq \frac{1}{10\sqrt{2}}. \quad (180)$$

Next,

$$m \max_{i \in [m]} |\alpha_i| < \frac{1}{48d^2} \exp(cd) \quad (181)$$

as required in the theorem. Therefore, by the contrapositive of [\(177\)](#), we must have

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} [(f_0(\mathbf{x}) - \eta(\mathbf{w}_*^\top \mathbf{x} + b_*))^2] > \frac{1}{50}, \quad \text{i.e.} \quad \|f_0 - \eta(\mathbf{w}_*^\top \mathbf{x} + b_*)\|_{L^2(\mathbb{P})} > \frac{1}{\sqrt{50}} = \frac{1}{5\sqrt{2}}. \quad (182)$$

Finally, by the triangle inequality and [\(180\)](#),

$$\|f_{\text{asym}} - \eta(\mathbf{w}_*^\top \mathbf{x} + b_*)\|_{L^2(\mathbb{P})} \geq \|f_0 - \eta(\mathbf{w}_*^\top \mathbf{x} + b_*)\|_{L^2(\mathbb{P})} - \|f_{\text{asym}} - f_0\|_{L^2(\mathbb{P})} > \frac{1}{5\sqrt{2}} - \frac{1}{10\sqrt{2}} = \frac{1}{10\sqrt{2}}. \quad (183) \quad \square$$

D. Linear Mode Connectivity (§ 4)

Theorem D.1 (Chord dev. in center-dom. regime, ReLU). *Let $\eta = \text{ReLU}$ and $\mathbf{x} \sim \mathcal{N}(0, \Sigma)$ with $\Sigma \succeq 0$. Fix $\mathbf{W}^A, \mathbf{W}^B \in \mathbb{R}^{m \times d}$. Suppose there exists $\beta \in [0, 1)$ such that for all $i \in [m]$,*

$$\|\mathbf{d}_i \odot \mathbf{w}_i^A\|_{\Sigma} \leq \beta \|\mathbf{f}_i\|_{\Sigma}, \quad \|\mathbf{d}_i \odot \mathbf{w}_i^B\|_{\Sigma} \leq \beta \|\mathbf{f}_i\|_{\Sigma}, \quad (184)$$

where $\|u\|_{\Sigma} := \sqrt{u^{\top} \Sigma u}$. Then,

$$\sup_{\lambda \in [0, 1]} \|\xi_{\mathbf{H}}(\lambda; \cdot)\|_{L^2(\mathbb{P}; \mathbb{R}^m)} = \mathcal{O}(\beta^{3/2}) \|\mathbf{F} \Sigma^{1/2}\|_F. \quad (185)$$

Proof. Fix $i \in [m]$. Write

$$\mathbf{w}_{\text{eff},i}^A := \mathbf{f}_i + \mathbf{d}_i \odot \mathbf{w}_i^A, \quad \mathbf{w}_{\text{eff},i}^B := \mathbf{f}_i + \mathbf{d}_i \odot \mathbf{w}_i^B \quad (186)$$

for the effective weights and

$$Z_i^A := (\mathbf{w}_{\text{eff},i}^A)^{\top} \mathbf{x}, \quad Z_i^B := (\mathbf{w}_{\text{eff},i}^B)^{\top} \mathbf{x} \quad (187)$$

for the preactivation random variables under the input distribution. If $\|\mathbf{f}_i\|_{\Sigma} = 0$, then (184) forces $\|\mathbf{d}_i \odot \mathbf{w}_i^A\|_{\Sigma} = \|\mathbf{d}_i \odot \mathbf{w}_i^B\|_{\Sigma} = 0$, hence $Z_i^A = Z_i^B = 0$ a.s. and $\xi_i(\lambda; \mathbf{x}) \equiv 0$; in this case the desired bound holds trivially for neuron i . Assume from now on that $\|\mathbf{f}_i\|_{\Sigma} > 0$. Then, $Z_i^A \sim \mathcal{N}(0, (\sigma_i^A)^2)$, $Z_i^B \sim \mathcal{N}(0, (\sigma_i^B)^2)$ with $\sigma_i^A := \|\mathbf{w}_{\text{eff},i}^A\|_{\Sigma}$, $\sigma_i^B := \|\mathbf{w}_{\text{eff},i}^B\|_{\Sigma}$. Since $\beta < 1$, we have $\sigma_i^A, \sigma_i^B > 0$. Further, the correlation of Z_i^A and Z_i^B is

$$\rho_i := \frac{\langle \mathbf{w}_{\text{eff},i}^A, \mathbf{w}_{\text{eff},i}^B \rangle_{\Sigma}}{\sigma_i^A \sigma_i^B}, \quad \theta_i := \arccos(\rho_i). \quad (188)$$

For ReLU, one obtains

$$\xi_i(\lambda; \mathbf{x})^2 = \left(\text{ReLU}((1-\lambda)Z_i^A + \lambda Z_i^B) - ((1-\lambda)\text{ReLU}(Z_i^A) + \lambda \text{ReLU}(Z_i^B)) \right)^2 \quad (189)$$

$$= \min\{(1-\lambda)^2(Z_i^A)^2, \lambda^2(Z_i^B)^2\} \mathbf{1}\{Z_i^A Z_i^B \leq 0\} \quad (190)$$

$$\leq (1-\lambda)^2(Z_i^A)^2 \mathbf{1}\{Z_i^A Z_i^B \leq 0\} + \lambda^2(Z_i^B)^2 \mathbf{1}\{Z_i^A Z_i^B \leq 0\}. \quad (191)$$

For centered bivariate Gaussian random variables with variances $(\sigma_i^A)^2, (\sigma_i^B)^2$ and correlation $\rho_i = \cos(\theta_i)$,

$$\mathbb{E}[(Z_i^A)^2 \mathbf{1}\{Z_i^A Z_i^B \leq 0\}] = (\sigma_i^A)^2 \frac{\theta_i - \sin(\theta_i) \cos(\theta_i)}{\pi}, \quad (192)$$

and likewise

$$\mathbb{E}[(Z_i^B)^2 \mathbf{1}\{Z_i^A Z_i^B \leq 0\}] = (\sigma_i^B)^2 \frac{\theta_i - \sin(\theta_i) \cos(\theta_i)}{\pi}. \quad (193)$$

Indeed, if $(G_1, G_2) \sim \mathcal{N}(0, \mathbf{I}_2)$, then

$$(Z_i^A/\sigma_i^A, Z_i^B/\sigma_i^B) \stackrel{d}{=} (G_1, \cos(\theta_i)G_1 + \sin(\theta_i)G_2), \quad (194)$$

where $\theta_i = \arccos(\rho_i) \in [0, \pi]$. Since (192) only depends on the joint distribution, we may compute it under this representation. Write $(G_1, G_2) = R(\cos \varphi, \sin \varphi)$ in polar coordinates. Since the standard Gaussian distribution is invariant under rotations, $\varphi \sim \text{Unif}([0, 2\pi])$ is independent of R , and $\mathbb{E}[R^2] = \mathbb{E}[G_1^2 + G_2^2] = 2$. In polar coordinates, we have

$$Z_i^A/\sigma_i^A = R \cos(\varphi), \quad (195)$$

$$Z_i^B/\sigma_i^B = \cos(\theta_i)R \cos(\varphi) + \sin(\theta_i)R \sin(\varphi) = R \cos(\varphi - \theta_i), \quad (196)$$

where the last equality uses $\cos(\varphi - \theta_i) = \cos(\varphi) \cos(\theta_i) + \sin(\varphi) \sin(\theta_i)$. Thus, since $\sigma_i^A, \sigma_i^B > 0$ and $R > 0$ a.s.,

$$Z_i^A Z_i^B \leq 0 \Leftrightarrow \sigma_i^A \sigma_i^B R^2 \cos(\varphi) \cos(\varphi - \theta_i) \leq 0 \quad (197)$$

$$\Leftrightarrow \cos(\varphi) \cos(\varphi - \theta_i) \leq 0. \quad (198)$$

For $\theta_i \in [0, \pi]$, this holds exactly for

$$\varphi \in \left[\frac{\pi}{2}, \frac{\pi}{2} + \theta_i \right] \cup \left[\frac{3\pi}{2}, \frac{3\pi}{2} + \theta_i \right], \quad (199)$$

with endpoints modulo 2π . Therefore,

$$\mathbb{E}[(Z_i^A)^2 \mathbb{1}\{Z_i^A Z_i^B \leq 0\}] = (\sigma_i^A)^2 \mathbb{E}[R^2 \cos^2(\varphi) \mathbb{1}\{\cos(\varphi) \cos(\varphi - \theta_i) \leq 0\}] \quad (200)$$

$$= (\sigma_i^A)^2 \underbrace{\mathbb{E}[R^2]}_{=2} \mathbb{E}[\cos^2(\varphi) \mathbb{1}\{\cos(\varphi) \cos(\varphi - \theta_i) \leq 0\}] \quad (201)$$

$$= (\sigma_i^A)^2 \frac{2}{2\pi} \int_{\{\cos(\varphi) \cos(\varphi - \theta_i) \leq 0\}} \cos^2(\varphi) d\varphi \quad (202)$$

$$= (\sigma_i^A)^2 \frac{2}{\pi} \int_{\pi/2}^{\pi/2 + \theta_i} \cos^2(\varphi) d\varphi \quad (203)$$

$$= (\sigma_i^A)^2 \frac{\theta_i - \sin(\theta_i) \cos(\theta_i)}{\pi}, \quad (204)$$

where the last equality follows from $\int \cos^2(\varphi) d\varphi = \varphi/2 + \sin(2\varphi)/4$ and $\sin(2\theta_i) = 2 \sin(\theta_i) \cos(\theta_i)$. Hence,

$$\|\xi_i(\lambda; \cdot)\|_{L^2(\mathbb{P})}^2 \leq \left((1-\lambda)^2 (\sigma_i^A)^2 + \lambda^2 (\sigma_i^B)^2 \right) \frac{\theta_i - \sin(\theta_i) \cos(\theta_i)}{\pi}. \quad (205)$$

It remains to bound the size and angular terms in (205) using (184). For the size term, write $\mathbf{w}_{\text{eff},i}^A = \mathbf{f}_i + \mathbf{r}_i^A$ and $\mathbf{w}_{\text{eff},i}^B = \mathbf{f}_i + \mathbf{r}_i^B$, where $\mathbf{r}_i^A := \mathbf{d}_i \odot \mathbf{w}_i^A$ and $\mathbf{r}_i^B := \mathbf{d}_i \odot \mathbf{w}_i^B$. Then

$$\sigma_i^A = \|\mathbf{f}_i + \mathbf{r}_i^A\|_{\Sigma} \leq \|\mathbf{f}_i\|_{\Sigma} + \|\mathbf{r}_i^A\|_{\Sigma} \leq (1 + \beta) \|\mathbf{f}_i\|_{\Sigma}, \quad (206)$$

and similarly $\sigma_i^B \leq (1 + \beta) \|\mathbf{f}_i\|_{\Sigma}$. Hence, since $(1 - \lambda)^2 + \lambda^2 \leq 1$,

$$(1 - \lambda)^2 (\sigma_i^A)^2 + \lambda^2 (\sigma_i^B)^2 \leq (1 + \beta)^2 \|\mathbf{f}_i\|_{\Sigma}^2. \quad (207)$$

To bound θ_i , define

$$\mathbf{u}_i^0 := \frac{\mathbf{f}_i}{\|\mathbf{f}_i\|_{\Sigma}}, \quad \mathbf{u}_i^A := \frac{\mathbf{f}_i + \mathbf{r}_i^A}{\|\mathbf{f}_i + \mathbf{r}_i^A\|_{\Sigma}}, \quad \mathbf{u}_i^B := \frac{\mathbf{f}_i + \mathbf{r}_i^B}{\|\mathbf{f}_i + \mathbf{r}_i^B\|_{\Sigma}}. \quad (208)$$

These vectors are normalized in $\|\cdot\|_{\Sigma}$, and $\rho_i = \langle \mathbf{u}_i^A, \mathbf{u}_i^B \rangle_{\Sigma}$. Set

$$\eta_i^A := \frac{\|\mathbf{r}_i^A\|_{\Sigma}}{\|\mathbf{f}_i\|_{\Sigma}}, \quad s_i^A := \frac{\langle \mathbf{r}_i^A, \mathbf{f}_i \rangle_{\Sigma}}{\|\mathbf{f}_i\|_{\Sigma}^2}. \quad (209)$$

Then $\eta_i^A \leq \beta$ and $s_i^A \geq -\eta_i^A$ by Cauchy-Schwarz. Since $\beta < 1$, this gives $1 + s_i^A > 0$, and

$$\langle \mathbf{u}_i^A, \mathbf{u}_i^0 \rangle_{\Sigma} = \frac{1 + s_i^A}{\sqrt{1 + 2s_i^A + (\eta_i^A)^2}}. \quad (210)$$

Moreover,

$$\frac{(1 + s_i^A)^2}{1 + 2s_i^A + (\eta_i^A)^2} - (1 - (\eta_i^A)^2) = \frac{(s_i^A + (\eta_i^A)^2)^2}{1 + 2s_i^A + (\eta_i^A)^2} \geq 0. \quad (211)$$

Thus,

$$\langle \mathbf{u}_i^A, \mathbf{u}_i^0 \rangle_{\Sigma} \geq \sqrt{1 - (\eta_i^A)^2} \geq \sqrt{1 - \beta^2}. \quad (212)$$

The same argument gives $\langle \mathbf{u}_i^B, \mathbf{u}_i^0 \rangle_{\Sigma} \geq \sqrt{1 - \beta^2}$. Now write $a_i^A := \langle \mathbf{u}_i^A, \mathbf{u}_i^0 \rangle_{\Sigma}$ and $a_i^B := \langle \mathbf{u}_i^B, \mathbf{u}_i^0 \rangle_{\Sigma}$. Since $\|\mathbf{u}_i^A - a_i^A \mathbf{u}_i^0\|_{\Sigma}^2 = 1 - (a_i^A)^2$ and similarly for B , Cauchy-Schwarz gives

$$\rho_i = \langle \mathbf{u}_i^A, \mathbf{u}_i^B \rangle_{\Sigma} = a_i^A a_i^B + \langle \mathbf{u}_i^A - a_i^A \mathbf{u}_i^0, \mathbf{u}_i^B - a_i^B \mathbf{u}_i^0 \rangle_{\Sigma} \quad (213)$$

$$\geq a_i^A a_i^B - \sqrt{1 - (a_i^A)^2} \sqrt{1 - (a_i^B)^2}. \quad (214)$$

The function

$$(a, b) \mapsto ab - \sqrt{1 - a^2} \sqrt{1 - b^2} \quad (215)$$

is increasing in each argument on $[0, 1]^2$. Since $a_i^A, a_i^B \geq \sqrt{1 - \beta^2}$, we obtain

$$\rho_i \geq (1 - \beta^2) - \beta^2 = 1 - 2\beta^2. \quad (216)$$

Therefore,

$$\theta_i = \arccos(\rho_i) \leq \arccos(1 - 2\beta^2) = 2 \arcsin(\beta) \leq \pi\beta. \quad (217)$$

Thus,

$$\frac{\theta_i - \sin(\theta_i) \cos(\theta_i)}{\pi} = \frac{1}{\pi} \int_0^{\theta_i} \underbrace{2 \sin^2(s)}_{\leq s^2} ds \leq \frac{2}{3\pi} \theta_i^3 \leq \frac{2\pi^2}{3} \beta^3. \quad (218)$$

Combining (205), (207), and (218) yields

$$\|\xi_i(\lambda; \cdot)\|_{L^2(\mathbb{P})}^2 \leq \frac{2\pi^2}{3} (1 + \beta)^2 \beta^3 \|\mathbf{f}_i\|_{\Sigma}^2. \quad (219)$$

Summing over i and using $\|\xi_{\mathbf{H}}(\lambda; \cdot)\|_{L^2(\mathbb{P}; \mathbb{R}^m)}^2 = \sum_{i=1}^m \|\xi_i(\lambda; \cdot)\|_{L^2(\mathbb{P})}^2$ gives

$$\|\xi_{\mathbf{H}}(\lambda; \cdot)\|_{L^2(\mathbb{P}; \mathbb{R}^m)}^2 \leq \frac{2\pi^2}{3} (1 + \beta)^2 \beta^3 \|\mathbf{F}\Sigma^{1/2}\|_F^2. \quad (220)$$

Finally, taking $\sqrt{\cdot}$ and the supremum over $\lambda \in [0, 1]$ on the l.h.s. proves the claim. \square

We also record a version of [Thm. D.1](#) for smooth activation functions.

Theorem D.2 (Chord deviation in center-dominated regime, C^2 activations). *Let $\mathbf{x} \sim \mathcal{N}(0, \Sigma)$ with $\Sigma \succeq 0$, and let $\eta \in C^2(\mathbb{R})$ with $L_\eta := \sup_{z \in \mathbb{R}} |\eta''(z)| < \infty$. Fix $\mathbf{W}^A, \mathbf{W}^B \in \mathbb{R}^{m \times d}$. Suppose there exists $\beta \in [0, 1)$ such that, for all $i \in [m]$,*

$$\|\mathbf{d}_i \odot \mathbf{w}_i^A\|_{\Sigma} \leq \beta \|\mathbf{f}_i\|_{\Sigma}, \quad \|\mathbf{d}_i \odot \mathbf{w}_i^B\|_{\Sigma} \leq \beta \|\mathbf{f}_i\|_{\Sigma}. \quad (221)$$

Then

$$\sup_{\lambda \in [0, 1]} \|\xi_{\mathbf{H}}(\lambda; \cdot)\|_{L^2(\mathbb{P}; \mathbb{R}^m)} \leq \frac{\sqrt{3}}{2} L_\eta \beta^2 \|\mathbf{F}\Sigma^{1/2}\|_F^2. \quad (222)$$

Proof. Fix $i \in [m]$ and write $\mathbf{r}_i^A := \mathbf{d}_i \odot \mathbf{w}_i^A$ and $\mathbf{r}_i^B := \mathbf{d}_i \odot \mathbf{w}_i^B$. Set

$$Z_i^A := (\mathbf{f}_i + \mathbf{r}_i^A)^\top \mathbf{x}, \quad Z_i^B := (\mathbf{f}_i + \mathbf{r}_i^B)^\top \mathbf{x}, \quad \Delta Z_i := Z_i^B - Z_i^A. \quad (223)$$

For fixed \mathbf{x} , define $h_i(s) := \eta(Z_i^A + s\Delta Z_i)$ for $s \in [0, 1]$. Then

$$\xi_i(\lambda; \mathbf{x}) = h_i(\lambda) - ((1 - \lambda)h_i(0) + \lambda h_i(1)). \quad (224)$$

We obtain

$$\xi_i(\lambda; \mathbf{x}) = (1 - \lambda) \int_0^\lambda h_i'(s) ds - \lambda \int_\lambda^1 h_i'(s) ds \quad (225)$$

$$= (1 - \lambda) \int_0^\lambda (h_i'(s) - h_i'(\lambda)) ds - \lambda \int_\lambda^1 (h_i'(s) - h_i'(\lambda)) ds. \quad (226)$$

Moreover, $h_i''(s) = \eta''(Z_i^A + s\Delta Z_i)(\Delta Z_i)^2$, and hence $|h_i''(s)| \leq L_\eta |\Delta Z_i|^2$ for all $s \in [0, 1]$. Therefore,

$$|\xi_i(\lambda; \mathbf{x})| \leq L_\eta |\Delta Z_i|^2 \left((1 - \lambda) \int_0^\lambda (\lambda - s) ds + \lambda \int_\lambda^1 (s - \lambda) ds \right) \quad (227)$$

$$= \frac{L_\eta}{2} \lambda(1 - \lambda) |\Delta Z_i|^2. \quad (228)$$

Thus

$$\|\xi_i(\lambda; \cdot)\|_{L^2(\mathbb{P})}^2 \leq \frac{L_\eta^2}{4} \lambda^2 (1 - \lambda)^2 \mathbb{E}[(\Delta Z_i)^4]. \quad (229)$$

Since $\Delta Z_i = (\mathbf{r}_i^B - \mathbf{r}_i^A)^\top \mathbf{x}$ is centered Gaussian,

$$\mathbb{E}[(\Delta Z_i)^4] = 3 \|\mathbf{r}_i^B - \mathbf{r}_i^A\|_{\Sigma}^4. \quad (230)$$

By (221),

$$\|\mathbf{r}_i^B - \mathbf{r}_i^A\|_{\Sigma} \leq \|\mathbf{r}_i^B\|_{\Sigma} + \|\mathbf{r}_i^A\|_{\Sigma} \leq 2\beta \|\mathbf{f}_i\|_{\Sigma}. \quad (231)$$

Hence

$$\|\xi_i(\lambda; \cdot)\|_{L^2(\mathbb{P})}^2 \leq 12 L_\eta^2 \lambda^2 (1 - \lambda)^2 \beta^4 \|\mathbf{f}_i\|_{\Sigma}^4. \quad (232)$$

Summing over i gives

$$\|\xi_{\mathbf{H}}(\lambda; \cdot)\|_{L^2(\mathbb{P}; \mathbb{R}^m)}^2 \leq 12 L_\eta^2 \lambda^2 (1 - \lambda)^2 \beta^4 \sum_{i=1}^m \|\mathbf{f}_i\|_{\Sigma}^4. \quad (233)$$

Taking $\sqrt{\cdot}$ and using $\lambda(1 - \lambda) \leq 1/4$ yields

$$\sup_{\lambda \in [0,1]} \|\xi_{\mathbf{H}}(\lambda; \cdot)\|_{L^2(\mathbb{P}; \mathbb{R}^m)} \leq \frac{\sqrt{3}}{2} L_{\eta} \beta^2 \left(\sum_{i=1}^m \|\mathbf{f}_i\|_{\Sigma}^4 \right)^{1/2}. \quad (234)$$

Finally,

$$\left(\sum_{i=1}^m \|\mathbf{f}_i\|_{\Sigma}^4 \right)^{1/2} \leq \sum_{i=1}^m \|\mathbf{f}_i\|_{\Sigma}^2 = \|\mathbf{F}\Sigma^{1/2}\|_F^2, \quad (235)$$

which proves the claim. \square

Proposition D.3 (LMC bound by chord deviation). *Consider a dataset $\mathcal{X} \times \mathcal{Y} \ni (\mathbf{x}, \mathbf{y}) \sim \mathbb{P}$, a neural network $f_{\theta} : \mathcal{X} \rightarrow \mathcal{Y} \subseteq \mathbb{R}^{d_{\text{out}}}$, and define the population loss $\mathcal{L}(\theta) := \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathbb{P}}[\ell(f_{\theta}(\mathbf{x}), \mathbf{y})] < \infty$. Assume that for every $\mathbf{y}, \mathbf{z} \mapsto \ell(\mathbf{z}, \mathbf{y})$ is convex and L_{ℓ} -Lipschitz w.r.t. $\|\cdot\|_2$. Fix two parameter vectors θ^A, θ^B and define the linear path $\theta(\lambda) := (1 - \lambda)\theta^A + \lambda\theta^B$, $\lambda \in [0, 1]$. Define further*

$$\xi_f(\lambda; \mathbf{x}) := f_{\theta(\lambda)}(\mathbf{x}) - ((1 - \lambda)f_{\theta^A}(\mathbf{x}) + \lambda f_{\theta^B}(\mathbf{x})). \quad (236)$$

Then, for all $\lambda \in [0, 1]$, we obtain

$$\begin{aligned} \mathcal{L}(\theta(\lambda)) - ((1 - \lambda)\mathcal{L}(\theta^A) + \lambda\mathcal{L}(\theta^B)) \\ \leq L_{\ell} \|\xi_f(\lambda; \cdot)\|_{L^2(\mathbb{P}_{\mathbf{x}}, \mathbb{R}^{d_{\text{out}}})}. \end{aligned} \quad (237)$$

Proof. Fix $\lambda \in [0, 1]$. By convexity of $\mathbf{z} \mapsto \ell(\mathbf{z}, \mathbf{y})$, for every (\mathbf{x}, \mathbf{y}) ,

$$\ell((1 - \lambda)f_{\theta^A}(\mathbf{x}) + \lambda f_{\theta^B}(\mathbf{x}), \mathbf{y}) \leq (1 - \lambda)\ell(f_{\theta^A}(\mathbf{x}), \mathbf{y}) + \lambda\ell(f_{\theta^B}(\mathbf{x}), \mathbf{y}). \quad (238)$$

Moreover, by the definition of $\xi_f(\lambda, \mathbf{x})$ and since $\ell(\cdot, \mathbf{y})$ is L_{ℓ} -Lipschitz w.r.t. $\|\cdot\|_2$,

$$\ell(f_{\theta(\lambda)}(\mathbf{x}), \mathbf{y}) \leq \ell((1 - \lambda)f_{\theta^A}(\mathbf{x}) + \lambda f_{\theta^B}(\mathbf{x}), \mathbf{y}) + L_{\ell} \|\xi_f(\lambda, \mathbf{x})\|_2. \quad (239)$$

Combining (238) and (239) and taking expectations w.r.t. $(\mathbf{x}, \mathbf{y}) \sim \mathbb{P}$ gives

$$\mathcal{L}(\theta(\lambda)) \leq (1 - \lambda)\mathcal{L}(\theta^A) + \lambda\mathcal{L}(\theta^B) + L_{\ell} \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_{\mathbf{x}}} [\|\xi_f(\lambda, \mathbf{x})\|_2]. \quad (240)$$

Finally, by Jensen's inequality,

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}_{\mathbf{x}}} [\|\xi_f(\lambda, \mathbf{x})\|_2] \leq \left(\mathbb{E}_{\mathbf{x} \sim \mathbb{P}_{\mathbf{x}}} [\|\xi_f(\lambda, \mathbf{x})\|_2^2] \right)^{1/2} = \|\xi_f(\lambda, \cdot)\|_{L^2(\mathbb{P}_{\mathbf{x}}, \mathbb{R}^{d_{\text{out}}})}. \quad (241)$$

Substituting this into (240) and rearranging yields the claim. \square

E. Approximate Subspace Support and Relaxed Realization Costs

In this section, we make rigorous how the exact subspace support model (Assumptions C.1 and C.2) can be interpreted as a low-rank approximation of a general data distribution: Namely, after projection on the top k principal components, exact realizability in the projected model implies approximate realizability on the original distribution, with the error controlled by the discarded second-moment tail energy. In this sense, the results from § 3 should not only be read as applicable to this idealized model, but as a tractable approximation to the observable realizability structure when the input distribution to a layer is only approximately low-rank. We also explain a natural relaxed variant of the realization cost framework from § 3, which does not require exact realizability on the data support, and show that the hard realization cost from § 3 arises as its ridgeless limit.

E.1. Approximate Subspace Support

Let the input \mathbf{x} to a layer be drawn from a distribution \mathbb{P} on \mathbb{R}^d with finite second moment $\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} \|\mathbf{x}\|_2^2 < \infty$, and define

$$\Sigma := \mathbb{E}_{\mathbf{x} \sim \mathbb{P}} [\mathbf{x} \mathbf{x}^\top] \succeq 0. \quad (242)$$

Let $\Sigma = \mathbf{U} \Lambda \mathbf{U}^\top$ be an eigendecomposition, i.e., $\Lambda = \text{Diag}(\lambda_1, \dots, \lambda_d)$ with $\lambda_1 \geq \dots \geq \lambda_d \geq 0$ and $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_d$. For a target rank k , let $\mathbf{U}_k \in \mathbb{R}^{d \times k}$ be the top k eigenvectors and define the tail energy

$$\tau_k := \sum_{j>k} \lambda_j. \quad (243)$$

In this context, Assumption C.1 for a given k means precisely that $\tau_k = 0$ with $\mathcal{U} := \text{span } \mathbf{U}_k$, and the non-degeneracy assumption (Assumption C.2) is equivalent to $\tau_{k-1} > 0$ for injective η . More generally, one can show the following:

Lemma E.1 (Subspace support and second-moment tail energy). *Assume $\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} \|\mathbf{x}\|_2^2 < \infty$. Then, $\tau_k = 0$ iff $\mathbb{P}(\mathbf{x} \in \text{im}(\mathbf{U}_k)) = 1$. Moreover, setting $\mathbf{P}_k := \mathbf{U}_k \mathbf{U}_k^\top$, the mean-squared residual outside the top- k subspace is $\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} [\|\mathbf{x} - \mathbf{P}_k \mathbf{x}\|_2^2] = \tau_k$.*

Proof. One can directly calculate

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} [\|\mathbf{x} - \mathbf{P}_k \mathbf{x}\|_2^2] = \mathbb{E}_{\mathbf{x} \sim \mathbb{P}} [\mathbf{x}^\top (\mathbf{I} - \mathbf{P}_k) \mathbf{x}] = \text{tr}((\mathbf{I} - \mathbf{P}_k) \Sigma) = \sum_{j>k} \lambda_j = \tau_k. \quad (244)$$

If $\tau_k = 0$, then $\|\mathbf{x} - \mathbf{P}_k \mathbf{x}\|_2^2 \geq 0$ has expectation 0, hence $\mathbf{x} = \mathbf{P}_k \mathbf{x}$ a.s., which means $\mathbf{x} \in \text{im}(\mathbf{U}_k)$ a.s. Conversely, if $\mathbf{x} \in \text{im}(\mathbf{U}_k)$ a.s., then $\mathbf{x} = \mathbf{P}_k \mathbf{x}$ a.s., so $\tau_k = 0$. \square

Lem. E.1 shows that τ_k is the residual incurred by projecting \mathbb{P} onto its top- k principal subspace. This makes the projected distribution $\mathbb{P}_k := (\mathbf{P}_k)_* \mathbb{P}$ a natural low-rank surrogate for the original input distribution. One may therefore ask whether coincidence of neuron functions on this projected model still implies approximate coincidence on the original distribution. The next proposition answers this affirmatively: for Lipschitz neuron functions, the resulting $L^2(\mathbb{P})$ error can be controlled by the second-moment tail energy τ_k . In this sense, the subspace support model from § 3 can also be seen as a tractable approximation whenever the input distribution is only approximately low-rank.

Proposition E.2 (Projected coincidence implies approximate coincidence). *Assume $\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} \|\mathbf{x}\|_2^2 < \infty$. Let $\mathbf{h}_1, \mathbf{h}_2 : \mathbb{R}^d \rightarrow \mathbb{R}$ be L_1 - and L_2 -Lipschitz, respectively, i.e., $|\mathbf{h}_i(\mathbf{x}) - \mathbf{h}_i(\mathbf{y})| \leq L_i \|\mathbf{x} - \mathbf{y}\|_2$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, $i \in \{1, 2\}$. Assume that \mathbf{h}_1 and \mathbf{h}_2 coincide on the projected distribution \mathbb{P}_k , i.e. $\mathbf{h}_1(\mathbf{P}_k \mathbf{x}) = \mathbf{h}_2(\mathbf{P}_k \mathbf{x})$ for \mathbb{P} -a.e. \mathbf{x} . Then,*

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} [(\mathbf{h}_1(\mathbf{x}) - \mathbf{h}_2(\mathbf{x}))^2] \leq 2(L_1^2 + L_2^2) \tau_k. \quad (245)$$

This means that if τ_k is small, any two neuron functions that agree on the top- k projected model are close in $L^2(\mathbb{P})$ on the original distribution.

Proof. Since $\mathbf{h}_1(\mathbf{P}_k \mathbf{x}) = \mathbf{h}_2(\mathbf{P}_k \mathbf{x})$ for \mathbb{P} -a.e. \mathbf{x} ,

$$\mathbf{h}_1(\mathbf{x}) - \mathbf{h}_2(\mathbf{x}) = (\mathbf{h}_1(\mathbf{x}) - \mathbf{h}_1(\mathbf{P}_k \mathbf{x})) + (\mathbf{h}_2(\mathbf{P}_k \mathbf{x}) - \mathbf{h}_2(\mathbf{x})). \quad (246)$$

Using $(a + b)^2 \leq 2a^2 + 2b^2$, we obtain

$$\begin{aligned} (\mathbf{h}_1(\mathbf{x}) - \mathbf{h}_2(\mathbf{x}))^2 &\leq 2(\mathbf{h}_1(\mathbf{x}) - \mathbf{h}_1(\mathbf{P}_k \mathbf{x}))^2 + 2(\mathbf{h}_2(\mathbf{x}) - \mathbf{h}_2(\mathbf{P}_k \mathbf{x}))^2 \\ &\leq 2(L_1^2 + L_2^2) \|\mathbf{x} - \mathbf{P}_k \mathbf{x}\|_2^2. \end{aligned} \quad (247)$$

Taking expectations and using [Lem. E.1](#),

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}}[(\mathbf{h}_1(\mathbf{x}) - \mathbf{h}_2(\mathbf{x}))^2] \leq 2(L_1^2 + L_2^2) \mathbb{E}_{\mathbf{x} \sim \mathbb{P}} \|\mathbf{x} - \mathbf{P}_k \mathbf{x}\|_2^2 = 2(L_1^2 + L_2^2) \tau_k, \quad (248)$$

which proves [\(245\)](#). \square

In particular, if η is L -Lipschitz and $\mathbf{h}_i(\mathbf{x}) = \eta(\mathbf{w}_i^\top \mathbf{x})$ for $\mathbf{w}_i \in \mathbb{R}^d$, $i \in \{1, 2\}$, then each \mathbf{h}_i is $L\|\mathbf{w}_i\|_2$ -Lipschitz (by Cauchy-Schwarz). Hence, [Prop. E.2](#) yields

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}}[(\mathbf{h}_1(\mathbf{x}) - \mathbf{h}_2(\mathbf{x}))^2] \leq 2L^2(\|\mathbf{w}_1\|_2^2 + \|\mathbf{w}_2\|_2^2) \tau_k \quad (249)$$

whenever $\mathbf{h}_1(\mathbf{P}_k \mathbf{x}) = \mathbf{h}_2(\mathbf{P}_k \mathbf{x})$ for \mathbb{P} -a.e. \mathbf{x} .

E.2. Relaxed Realization Costs

In the previous subsection, we have seen in what sense the subspace support model can be viewed as an approximation to a general input distribution. For realization costs, a complementary question is how one should formulate the cost when exact realizability is not imposed. A natural choice is to relax the hard constraint by penalizing squared $L^2(\mathbb{P})$ approximation error and weight norm. We record it here to show that the realization cost from [§ 3](#) arises as its ridgeless limit.

Just as in [§ 3](#), consider a single asymmetric layer in isolation, and define $\mathbf{H}(\mathbf{W}; \mathbf{x}) := \eta((\mathbf{F} + \mathbf{D} \odot \mathbf{W})\mathbf{x}) \in \mathbb{R}^m$, where $\mathbf{W} \in \mathbb{R}^{m \times d}$, $\mathbf{x} \in \mathbb{R}^d$, $\mathbf{F}, \mathbf{D} \in \mathbb{R}^{m \times d}$ are fixed architectural components, and $\eta: \mathbb{R} \rightarrow \mathbb{R}$ acts elementwise.

Definition E.3 (Relaxed L^2 realization cost). Fix neuron i and $\beta > 0$. For any $\mathbf{h} \in L^2(\mathbb{P})$, define

$$\|\mathbf{h}\|_{\mathcal{H}_i(\mathbb{P}), \beta}^2 = \|\mathbf{h}\|_{\mathcal{H}_i, \beta}^2 := \inf_{\mathbf{w} \in \mathbb{R}^d} \left\{ \beta^{-1} \mathbb{E}_{\mathbf{x} \sim \mathbb{P}} [(\mathbf{h}(\mathbf{x}) - \mathbf{h}_i(\mathbf{w}; \mathbf{x}))^2] + \|\mathbf{w}\|_2^2 \right\}. \quad (250)$$

In cases where the distribution \mathbb{P} might vary, we will make the dependence explicit through writing $\mathcal{H}_i(\mathbb{P})$. It is straightforward to see that for Lipschitz η and any $\mathbf{h} \in L^2(\mathbb{P})$, $\|\mathbf{h}\|_{\mathcal{H}_i, \beta}^2 < \infty$ is finite. A canonical family of targets we will look at are functions of the form $\mathbf{h}_\mathbf{u}(\mathbf{x}) := \eta(\mathbf{u}^\top \mathbf{x})$, $\mathbf{u} \in \mathbb{R}^d$, i.e., precisely *realizable* target features, since this is precisely the setting in [§ 3](#).

Theorem E.4 (Explicit upper bound via a quadratic surrogate). Assume η is L -Lipschitz. Fix neuron i and define $\mathbf{D}_i := \text{Diag}(\mathbf{d}_i) \in \mathbb{R}^{d \times d}$. For any $\mathbf{u} \in \mathbb{R}^d$,

$$\|\mathbf{h}_\mathbf{u}\|_{\mathcal{H}_i, \beta}^2 \leq \inf_{\mathbf{w} \in \mathbb{R}^d} \left\{ \beta^{-1} L^2 \|\mathbf{u} - \mathbf{f}_i - \mathbf{D}_i \mathbf{w}\|_\Sigma^2 + \|\mathbf{w}\|_2^2 \right\}. \quad (251)$$

Moreover, the infimum in [\(251\)](#) has a unique minimizer

$$\mathbf{w}_i^*(\mathbf{u}) = \left(\beta \mathbf{I} + L^2 \mathbf{D}_i \Sigma \mathbf{D}_i \right)^{-1} L^2 \mathbf{D}_i \Sigma (\mathbf{u} - \mathbf{f}_i), \quad (252)$$

and the corresponding minimized value equals

$$\frac{L^2}{\beta} \|\mathbf{u} - \mathbf{f}_i\|_\Sigma^2 - \frac{L^4}{\beta^2} (\mathbf{u} - \mathbf{f}_i)^\top \Sigma \mathbf{D}_i \left(\mathbf{I} + \frac{L^2}{\beta} \mathbf{D}_i \Sigma \mathbf{D}_i \right)^{-1} \mathbf{D}_i \Sigma (\mathbf{u} - \mathbf{f}_i). \quad (253)$$

Proof. Fix $\mathbf{w} \in \mathbb{R}^d$ and abbreviate

$$\ell_\mathbf{u}(\mathbf{x}) := \mathbf{u}^\top \mathbf{x}, \quad \ell_i(\mathbf{w}; \mathbf{x}) := (\mathbf{f}_i + \mathbf{D}_i \mathbf{w})^\top \mathbf{x}. \quad (254)$$

Since η is L -Lipschitz,

$$(\mathbf{h}_\mathbf{u}(\mathbf{x}) - \mathbf{h}_i(\mathbf{w}; \mathbf{x}))^2 \leq L^2 (\ell_\mathbf{u}(\mathbf{x}) - \ell_i(\mathbf{w}; \mathbf{x}))^2. \quad (255)$$

Taking expectations yields

$$\mathbb{E}[(\mathbf{h}_\mathbf{u}(\mathbf{x}) - \mathbf{h}_i(\mathbf{w}; \mathbf{x}))^2] \leq L^2 \|\mathbf{u} - \mathbf{f}_i - \mathbf{D}_i \mathbf{w}\|_\Sigma^2. \quad (256)$$

Plugging into [Def. E.3](#) and then taking the infimum over \mathbf{w} proves [\(251\)](#). To compute the infimum, define $\alpha := L^2/\beta$ and consider the strictly convex quadratic

$$J(\mathbf{w}) := \|\mathbf{w}\|_2^2 + \alpha (\mathbf{u} - \mathbf{f}_i - \mathbf{D}_i \mathbf{w})^\top \Sigma (\mathbf{u} - \mathbf{f}_i - \mathbf{D}_i \mathbf{w}). \quad (257)$$

Differentiating gives

$$\nabla J(\mathbf{w}) = 2\mathbf{w} - 2\alpha \mathbf{D}_i \Sigma (\mathbf{u} - \mathbf{f}_i - \mathbf{D}_i \mathbf{w}) = 2 \left(\mathbf{I} + \alpha \mathbf{D}_i \Sigma \mathbf{D}_i \right) \mathbf{w} - 2\alpha \mathbf{D}_i \Sigma (\mathbf{u} - \mathbf{f}_i). \quad (258)$$

Since $\mathbf{I} + \alpha \mathbf{D}_i \boldsymbol{\Sigma} \mathbf{D}_i \succ 0$, setting $\nabla J(\mathbf{w}) = 0$ yields the unique minimizer (252). It remains to compute the minimum value. Write $\mathbf{r} := \mathbf{u} - \mathbf{f}_i$ and $\mathbf{A} := \mathbf{I} + \alpha \mathbf{D}_i \boldsymbol{\Sigma} \mathbf{D}_i \succ 0$. Expanding J yields

$$J(\mathbf{w}) = \mathbf{w}^\top \mathbf{A} \mathbf{w} - 2\alpha \mathbf{w}^\top \mathbf{D}_i \boldsymbol{\Sigma} \mathbf{r} + \alpha \mathbf{r}^\top \boldsymbol{\Sigma} \mathbf{r}. \quad (259)$$

Let $\mathbf{c} := \alpha \mathbf{D}_i \boldsymbol{\Sigma} \mathbf{r}$. Then, $J(\mathbf{w}) = \mathbf{w}^\top \mathbf{A} \mathbf{w} - 2\mathbf{w}^\top \mathbf{c} + \alpha \mathbf{r}^\top \boldsymbol{\Sigma} \mathbf{r}$. Since $\mathbf{A} \succ 0$, we have

$$\mathbf{w}^\top \mathbf{A} \mathbf{w} - 2\mathbf{w}^\top \mathbf{c} = (\mathbf{w} - \mathbf{A}^{-1} \mathbf{c})^\top \mathbf{A} (\mathbf{w} - \mathbf{A}^{-1} \mathbf{c}) - \mathbf{c}^\top \mathbf{A}^{-1} \mathbf{c}, \quad (260)$$

which implies that the minimum is attained at $\mathbf{w}^* := \mathbf{A}^{-1} \mathbf{c}$ (which coincides with (252)) and equals

$$\min_{\mathbf{w} \in \mathbb{R}^d} J(\mathbf{w}) = \alpha \mathbf{r}^\top \boldsymbol{\Sigma} \mathbf{r} - \mathbf{c}^\top \mathbf{A}^{-1} \mathbf{c} = \alpha \mathbf{r}^\top \boldsymbol{\Sigma} \mathbf{r} - \alpha^2 \mathbf{r}^\top \boldsymbol{\Sigma} \mathbf{D}_i (\mathbf{I} + \alpha \mathbf{D}_i \boldsymbol{\Sigma} \mathbf{D}_i)^{-1} \mathbf{D}_i \boldsymbol{\Sigma} \mathbf{r}, \quad (261)$$

which is exactly (253). \square

If η is bi-Lipschitz, one can prove a lower bound in a similar way. We will now show that the $\beta \rightarrow 0$ (ridgeless) limit recovers the hard realization cost from § 3, where non-realizable targets have infinite cost.

Theorem E.5 (Ridgeless limit). *Assume η is Lipschitz and $\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} \|\mathbf{x}\|_2^2 < \infty$. Fix neuron i and let*

$$\|\mathbf{h}\|_{\mathcal{H}_i(\mathbb{P})}^2 = \|\mathbf{h}\|_{\mathcal{H}_i}^2 := \inf \{ \|\mathbf{w}\|_2^2 \mid \mathbb{E}_{\mathbf{x} \sim \mathbb{P}} [(\mathbf{h}(\mathbf{x}) - \mathbf{h}_i(\mathbf{w}; \mathbf{x}))^2] = 0 \} \in [0, \infty] \quad (262)$$

with the convention $\inf \emptyset := +\infty$. Then for any $\mathbf{h} \in L^2(\mathbb{P})$,

$$\lim_{\beta \rightarrow 0} \|\mathbf{h}\|_{\mathcal{H}_i, \beta}^2 = \|\mathbf{h}\|_{\mathcal{H}_i}^2. \quad (263)$$

Proof. Write $f(\mathbf{w}) := \mathbb{E}[(\mathbf{h}(\mathbf{x}) - \mathbf{h}_i(\mathbf{w}; \mathbf{x}))^2] \geq 0$ and $\phi(\beta) := \inf_{\mathbf{w}} \{f(\mathbf{w}) + \beta \|\mathbf{w}\|_2^2\}$ so that $\|\mathbf{h}\|_{\mathcal{H}_i, \beta}^2 = \beta^{-1} \phi(\beta)$ by Def. E.3.

First, suppose $\|\mathbf{h}\|_{\mathcal{H}_i}^2 = +\infty$. We show $\|\mathbf{h}\|_{\mathcal{H}_i, \beta}^2 \rightarrow +\infty$. Assume for contradiction that there exists $0 < \beta_n \rightarrow 0$ and $C < \infty$ with $\|\mathbf{h}\|_{\mathcal{H}_i, \beta_n}^2 \leq C$ for all n . Then there exist \mathbf{w}_n such that

$$f(\mathbf{w}_n) + \beta_n \|\mathbf{w}_n\|_2^2 \leq \beta_n (C + 1). \quad (264)$$

In particular, this implies that $\|\mathbf{w}_n\|_2^2 \leq C + 1$ and $f(\mathbf{w}_n) \rightarrow 0$ as $n \rightarrow \infty$. Since $(\mathbf{w}_n)_n$ is bounded, we can extract a convergent subsequence $(\mathbf{w}_{n_k})_k$ with $\mathbf{w}_{n_k} \rightarrow \mathbf{w}^*$ as $k \rightarrow \infty$. By Lipschitzness of η and $\mathbb{E} \|\mathbf{x}\|_2^2 < \infty$, the map $\mathbf{w} \mapsto \mathbf{h}_i(\mathbf{w}; \cdot)$ is continuous from $(\mathbb{R}^d, \|\cdot\|_2)$ to $L^2(\mathbb{P})$, hence $f(\mathbf{w}_{n_k}) \rightarrow f(\mathbf{w}^*)$ and therefore $f(\mathbf{w}^*) = 0$. This contradicts $\|\mathbf{h}\|_{\mathcal{H}_i}^2 = +\infty$.

Assume now $\|\mathbf{h}\|_{\mathcal{H}_i}^2 < \infty$. Then, for every $\beta > 0$,

$$\phi(\beta) = \inf_{\mathbf{w}} \{f(\mathbf{w}) + \beta \|\mathbf{w}\|_2^2\} \leq \inf_{\mathbf{w}: f(\mathbf{w})=0} \{\beta \|\mathbf{w}\|_2^2\} = \beta \|\mathbf{h}\|_{\mathcal{H}_i}^2, \quad (265)$$

and

$$\limsup_{\beta \rightarrow 0} \beta^{-1} \phi(\beta) \leq \|\mathbf{h}\|_{\mathcal{H}_i}^2. \quad (266)$$

For the reverse inequality, let $0 < \beta_n \rightarrow 0$ and pick \mathbf{w}_n such that $f(\mathbf{w}_n) + \beta_n \|\mathbf{w}_n\|_2^2 \leq \phi(\beta_n) + \beta_n/n$. By the previous upper bound, $\phi(\beta_n) \leq \beta_n (\|\mathbf{h}\|_{\mathcal{H}_i}^2 + 1)$ for all large enough n , hence $(\|\mathbf{w}_n\|_2)_n$ is bounded and $f(\mathbf{w}_n) \rightarrow 0$. Similar as above, any limit point \mathbf{w}^* of a convergent subsequence satisfies $f(\mathbf{w}^*) = 0$ as $n \rightarrow \infty$. Therefore $\|\mathbf{w}^*\|_2^2 \geq \|\mathbf{h}\|_{\mathcal{H}_i}^2$ and thus $\liminf_{n \rightarrow \infty} \|\mathbf{w}_n\|_2^2 \geq \|\mathbf{h}\|_{\mathcal{H}_i}^2$. Since $\beta_n^{-1} \phi(\beta_n) \geq \|\mathbf{w}_n\|_2^2 - 1/n$, we can conclude

$$\liminf_{\beta \rightarrow 0} \beta^{-1} \phi(\beta) \geq \|\mathbf{h}\|_{\mathcal{H}_i}^2. \quad (267)$$

Combining (266) and (267) gives the claim. \square

When $\eta = \text{id}$, the relaxed cost in Def. E.3 reduces to a strictly convex quadratic ridge objective in \mathbf{w} , so the minimizer and optimum admit explicit closed forms for every $\beta > 0$. For general nonlinear $\eta \neq \text{id}$, the same objective is typically nonconvex in \mathbf{w} and no simple closed-form solution is available, so apart from the ridgeless limit, one would have to work with tractable surrogates such as Thm. E.4. By contrast, in the ridgeless limit the problem simplifies conceptually: Thm. E.5 shows that as $\beta \rightarrow 0$, one recovers the realization cost from § 3, which admits the explicit closed form (5). At the same time, because the limit enforces exact realizability, it is singular under perturbations of the *support*, i.e., even a small amount of mass outside the top- k subspace can destroy feasibility on the original distribution and make the hard cost jump to $+\infty$.

For fixed $\beta > 0$, however, the relaxed objective remains stable under such perturbations. The next corollary records the corresponding quantitative upper bound.

Corollary E.6 (Projected realizability yields small relaxed cost). *Fix neuron i and $\beta > 0$, and let $\mathbf{h} \in L^2(\mathbb{P})$. Assume that the hard realization cost of \mathbf{h} with respect to \mathbb{P}_k is finite and attained at some $\mathbf{w}^* \in \mathbb{R}^d$, i.e.*

$$\mathbf{h}(\mathbf{P}_k \mathbf{x}) = \mathbf{h}_i(\mathbf{w}^*; \mathbf{P}_k \mathbf{x}) \quad \text{for } \mathbb{P}\text{-a.e. } \mathbf{x}, \quad \|\mathbf{w}^*\|_2 = \|\mathbf{h}\|_{\mathcal{H}_i(\mathbb{P}_k)}. \quad (268)$$

Assume also that \mathbf{h} , $\mathbf{h}_i(\mathbf{w}^*; \cdot)$ are $L_{\mathbf{h}}$ - and $L_i(\mathbf{w}^*)$ -Lipschitz, respectively. Then,

$$\|\mathbf{h}\|_{\mathcal{H}_i(\mathbb{P}), \beta}^2 \leq \|\mathbf{h}\|_{\mathcal{H}_i(\mathbb{P}_k)}^2 + 2(L_{\mathbf{h}}^2 + L_i(\mathbf{w}^*)^2) \beta^{-1} \tau_k. \quad (269)$$

Proof. By [Def. E.3](#), evaluating the infimum at \mathbf{w}^* and applying [Prop. E.2](#) gives

$$\|\mathbf{h}\|_{\mathcal{H}_i(\mathbb{P}), \beta}^2 \leq \beta^{-1} \mathbb{E}_{\mathbf{x} \sim \mathbb{P}} [(\mathbf{h}(\mathbf{x}) - \mathbf{h}_i(\mathbf{w}^*; \mathbf{x}))^2] + \|\mathbf{w}^*\|_2^2 \leq 2(L_{\mathbf{h}}^2 + L_i(\mathbf{w}^*)^2) \beta^{-1} \tau_k + \|\mathbf{w}^*\|_2^2, \quad (270)$$

which yields the claim since $\|\mathbf{w}^*\|_2 = \|\mathbf{h}\|_{\mathcal{H}_i(\mathbb{P}_k)}$. □

[Corollary E.6](#) links the approximate subspace support from [§ E.1](#) and the relaxed realization costs we introduced in [§ E.2](#). Namely, the first term on the r.h.s. of [\(269\)](#) is the hard realization cost on the projected model \mathbb{P}_k . The second term is the approximation error incurred when passing back to the original distribution. Thus, for fixed $\beta > 0$, the projected subspace support model yields a quantitative upper bound on the relaxed realization cost under the true distribution.

F. Additional Figures

F.1. Visualization of Effective Function Classes

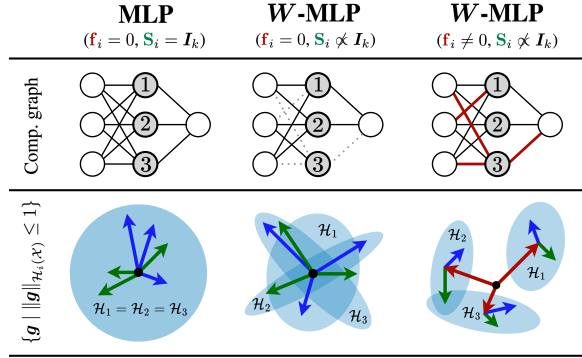


Figure 6. Eff. fct. classes and learned features (run A/run B) w/ varying levels of symmetry. Col. 1: In a fully symmetric MLP, each neuron can implement the same functions on \mathcal{X} . Col. 2: Pruning weights affects \mathbf{S}_i and can introduce anisotropy. Col. 3: Fixed weights via \mathbf{F} introduce functional biases.

F.2. LMC Interpolation for MLPs and ResNets

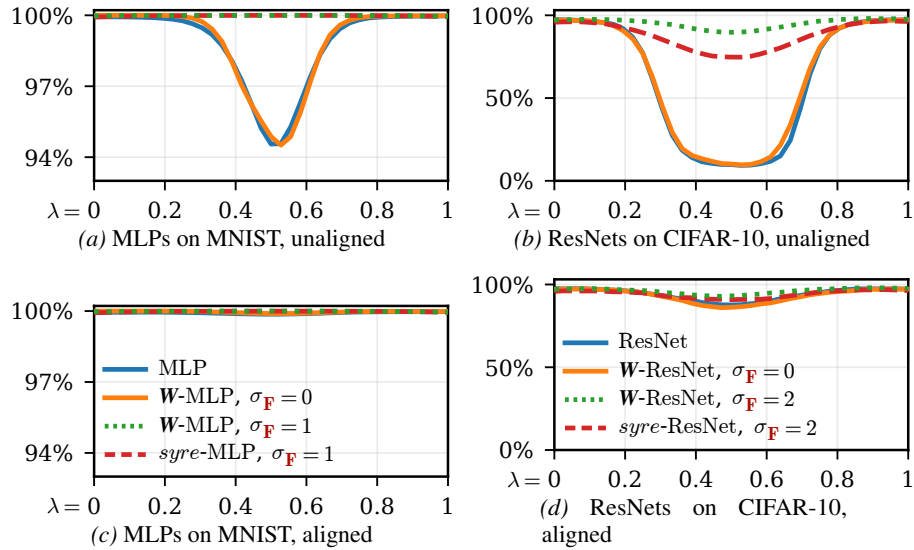


Figure 7. Aligned and unaligned training accuracy LMC interpolation for standard models and asymmetric models (average of 8 model pairs each). Alignment using activation matching.

G. Experimental Setup

G.1. MLPs

For MLP experiments, we use MLPs with 4 layers, input dimension 784 (MNIST input dimension), hidden dimensions 512, layer norms, and GELU activation functions. For \mathcal{W} -asymmetric MLPs (\mathcal{W} -MLPs), following Lim et al. (2024b), we fix 64 entries per linear layer row (= neuron) for the first three layers, and 256 entries per row for the final layer. Fixed entries are randomly sampled, making use of a mask seed that ensures the same mask in different model instances. Whenever we perform parameter interpolation between models, the models have the same mask (which is considered a fixed part of the architecture).

For the fixed weight scale $\sigma_{\mathbf{F}}$, we use $\sigma_{\mathbf{F}} = 1$ as the default choice for effective symmetry breaking \mathcal{W} -MLPs, and $\sigma_{\mathbf{F}} = 0$ for the variant that does not exhibit effective symmetry breaking, as well as a range of values for $\sigma_{\mathbf{F}}$ between 0.0 and 5.0 for one $\sigma_{\mathbf{F}}$ sweep experiment. In contrast to Lim et al. (2024b), we always apply the same value of $\sigma_{\mathbf{F}}$ at all layers.

We train the MLPs, \mathcal{W} -MLPs, and *syre*-MLPs on MNIST for 100 epochs with AdamW (Loshchilov & Hutter, 2019), a learning rate of 0.001, and a weight decay of 0.01.

G.2. ResNets

For our ResNet experiments, we use ResNet20 networks with $8\times$ width multiplier (following the setup of Lim et al. (2024b)), ReLU activation functions, and group norms. The architecture contains an initial convolution ($3 \rightarrow 128$ channels), three block groups ($128, 128 \rightarrow 256, 256 \rightarrow 512$ channels), and a final linear layer. For \mathcal{W} -asymmetric ResNets, in each convolution, a number of fixed weights are introduced per input channel (spread randomly over the output channels and kernel dimensions). The number of fixed parameters varies per block group, with the shortcut convolution in block groups 2 and 3 having a separate value for the number of fixed values. Table 3 gives the values across the network.

Table 3. Number of fixed parameters per layer.

Layer	# fixed
Initial convolution	12
Block group 1	108
Block group 2	162
Shortcut	18
Block group 3	216
Shortcut	24
Final linear	24

The default fixed weight scale in our \mathcal{W} -ResNet experiments is $\sigma_{\mathbf{F}} = 2$ for the effective symmetry breaking setting, and $\sigma_{\mathbf{F}} = 0$ for the ineffective symmetry breaking setting. We train the ResNets on CIFAR-10 for 100 epochs using AdamW (Loshchilov & Hutter, 2019) with weight decay of 0.01, and learning rate warm up from 0.0001 to 0.01 in the first 20 epochs ($\gamma = 1.259$).

G.3. Datasets and Augmentations

We use MNIST and CIFAR-10 datasets with a random 80%/10%/10% train/val/test split. On MNIST, we only use a normalization transform. On CIFAR-10, we additionally use a RandomHorizontalFlip transform, and a RandomCrop transform with 4 pixels of padding and a 32x32 output size.

G.4. Activation Matching

We implement activation matching as described in § 5. For MLPs, this is straightforward: we forward the full dataset through the models and record intermediate activations $\mathbf{Z}^A, \mathbf{Z}^B$ at each layer, which we use to compute the similarity matrix \mathbf{S}^{act} . For ResNets, some subtleties are involved. The permutable dimension is the convolution channel dimension, which takes the role of the hidden representation dimension in \mathbf{Z} . We could average over the pixel dimensions to have $\mathbf{Z} \in \mathbb{R}^{n \times d}$. However we find it beneficial for LMC to treat each pixel as an independent data sample, producing $\mathbf{Z} \in \mathbb{R}^{(nk^2) \times d}$. Furthermore, in ResNets, the residual connections cause certain activation permutations to affect several layer inputs: For example, the input

to block 1 in block group 1 is also indirectly fed into block 2 via the additive residual connection, without anything yielding an independent permutation symmetry in between. Each block group exhibits one such a global permutation that affects several layer inputs, and hence should be estimated jointly, yielding $\mathbf{Z} \in \mathbb{R}^{(nk^2c) \times d}$ (for $c = 4$ in our case). Besides, each block admits an inner permutation symmetry between its two convolution layers, which can be estimated independently. Due to memory constraints, we perform activation matching on 10% of our CIFAR-10 training set, which is sufficient to yield an alignment of low barrier. For our plots in § 5, we average the activation matching objective across all layers.

G.5. Coherence Experiment

We record details for the experiment isolating dependence on the coherence of the data subspace (see Fig. 4) in a teacher-student setup. The latent input variable is $(a, b) \sim \mathcal{N}(0, \mathbf{I}_2)$ and we aim to learn the teacher $f(a, b) := \text{ReLU}(a + b)$. We sample 1024 training examples and 2048 test examples from this latent distribution, which are used across all coherence levels. The latent (a, b) is embedded into the ambient space \mathbb{R}^d , $d = 20$, as

$$\mathbf{x} = \mathbf{U}_k(a, b)^\top, \quad \mathbf{U}_k \in \mathbb{R}^{d \times 2}, \quad \mathbf{U}_k^\top \mathbf{U}_k = \mathbf{I}_2, \quad (271)$$

where $k \in \{1, \dots, 10\}$ controls the active ambient support. For the “exact-copy” family, all rows of \mathbf{U}_k are zero except

$$U_{2r-1,1} = \frac{1}{\sqrt{k}}, \quad U_{2r,2} = \frac{1}{\sqrt{k}}, \quad r = 1, \dots, k. \quad (272)$$

I.e., each latent coordinate is copied uniformly over k ambient coordinates. For the “random-frame” family, we sample a permutation π and angles $\theta_1, \dots, \theta_{d/2} \sim \text{Unif}(0, \pi)$. These sampled values are reused for all k . For a given k , all rows of \mathbf{U}_k are zero except, for $r = 1, \dots, k$,

$$U_{\pi(2r-1),:} = \frac{1}{\sqrt{k}}(\cos \theta_r, \sin \theta_r), \quad U_{\pi(2r),:} = \frac{1}{\sqrt{k}}(-\sin \theta_r, \cos \theta_r). \quad (273)$$

The student is a one-hidden-layer asymmetric ReLU network of width 128 with $\mathbf{F} = 0$ and $\mathbf{D} \sim \text{Bernoulli}(0.15)$ i.i.d. Further, the output weights are fixed to 1. Hidden weights are initialized independently from $\mathcal{N}(0, 0.05^2)$ and hidden biases are initialized to zero. Each model is trained for 100 epochs on MSE on a batch size of 512 using AdamW (Loshchilov & Hutter, 2019) (learning rate 10^{-2} , weight decay 10^{-2}). For each k and subspace embedding, we sample 5 outer seeds for \mathbf{D} , and 10 initializations of student pairs for interpolation each, yielding in total 50 interpolation pairs per k and subspace embedding. For each k and subspace embedding, the reported midpoint barrier is the mean over the 50 pair-level midpoint barriers \pm standard deviation over these 50 values.

G.6. Gaussian Mixture Synthetic Dataset

In our experiments (Q3 in § 5, and § H.2, § H.3 in the appendix) we use a synthetic Gaussian mixture model dataset that allows us to control the intrinsic data dimension parameter k . We sample class centers $\boldsymbol{\mu}_c \in \mathbb{R}^k$, $(\boldsymbol{\mu}_c)_i \sim \mathcal{N}(0, \sigma_{\text{sep}}^2)$, where σ_{sep} is a class separation parameter. We generate k -dimensional class k latent samples as

$$\mathbf{z}_i = \mathbf{Q}_c(\boldsymbol{\mu}_c + \boldsymbol{\varepsilon}_i) \in \mathbb{R}^k, \quad (274)$$

where $\boldsymbol{\varepsilon}_i \sim \mathcal{N}(0, \text{Diag}(s_1^2, \dots, s_k^2))$ is anisotropic Gaussian noise (with direction-wise scaling vector $\mathbf{s} \in \mathbb{R}^k$) and \mathbf{Q}_c is a class-specific orthogonal transformation, and embed them into the ambient space as $\mathbf{x}_i = \mathbf{P}\mathbf{z}_i \in \mathbb{R}^d$, where $\mathbf{P} \in \mathbb{R}^{d \times k}$. We generate 50 000 training samples using five classes in an ambient space of dimension $d = 128$, varying the intrinsic dimension k depending on the experiment.

G.7. Empirical Neuron Realization and Swap Costs

We estimate neuron realization costs in trained models, and pairwise neuron swap costs derived from them (Q3 in § 5, and § H.3 in the appendix) using two methods: Using the Mahalanobis distance of (5), and using a ridge regression setup (see Def. E.3).

For the Mahalanobis distance estimation, we record activations during a forward pass on the whole dataset, and use them to estimate an input subspace for each layer as a PCA subspace that explains at least 90% of the variance. We then follow (5) to compute per neuron the induced subspace preactivation coefficient \mathbf{a} , the Gram matrix \mathbf{S}_i , and finally the realization cost (Mahalanobis distance) under the subspace support model.

For the ridge regression realization cost estimation, we optimize (250) for 15 000 iterations using AdamW (learning rate 0.005, no weight decay). We sweep $\beta \in \{0.001, 0.01, 0.1, 1.0\}$, and report results for $\beta = 0.01$. For the ridge regression

experiments, we re-train models using batch norm instead of layer norm, consider the neuron function to include the normalization and activation function, but leave the batch norm parameters fixed. The optimization for neuron i 's realization cost with respect to neuron j 's class starts at neuron i 's trained weights.

H. Additional Experimental Results

H.1. Layerwise Activation Matching Objectives

In addition to Fig. 2, we also report activation matching objectives per layer. In Fig. 8, we plot layerwise activation matching objectives by epoch for MLPs on MNIST. This reveals that for all observed cases, the patterns are consistent across layers and, therefore, the aggregate over all layers from Fig. 2 is informative.

We also report activation matching objectives for σ -asymmetric networks (Lim et al., 2024b). This different symmetry breaking intervention operates by replacing the nonlinearity with one that does not act elementwise. This specific approach has been found to not break symmetries as effectively as \mathcal{W} -asym. in the context of unaligned LMC. While σ -asym. networks do not fit into our framework of (1), one can still perform activation matching to qualitatively analyze neuron identifiability. Fig. 8 reveals that, in fact, the identity does not produce a higher objective than random permutations, and therefore, similar to \mathcal{W} -asym. networks with $\sigma_{\mathbf{F}} = 0$, symmetries are not broken effectively.

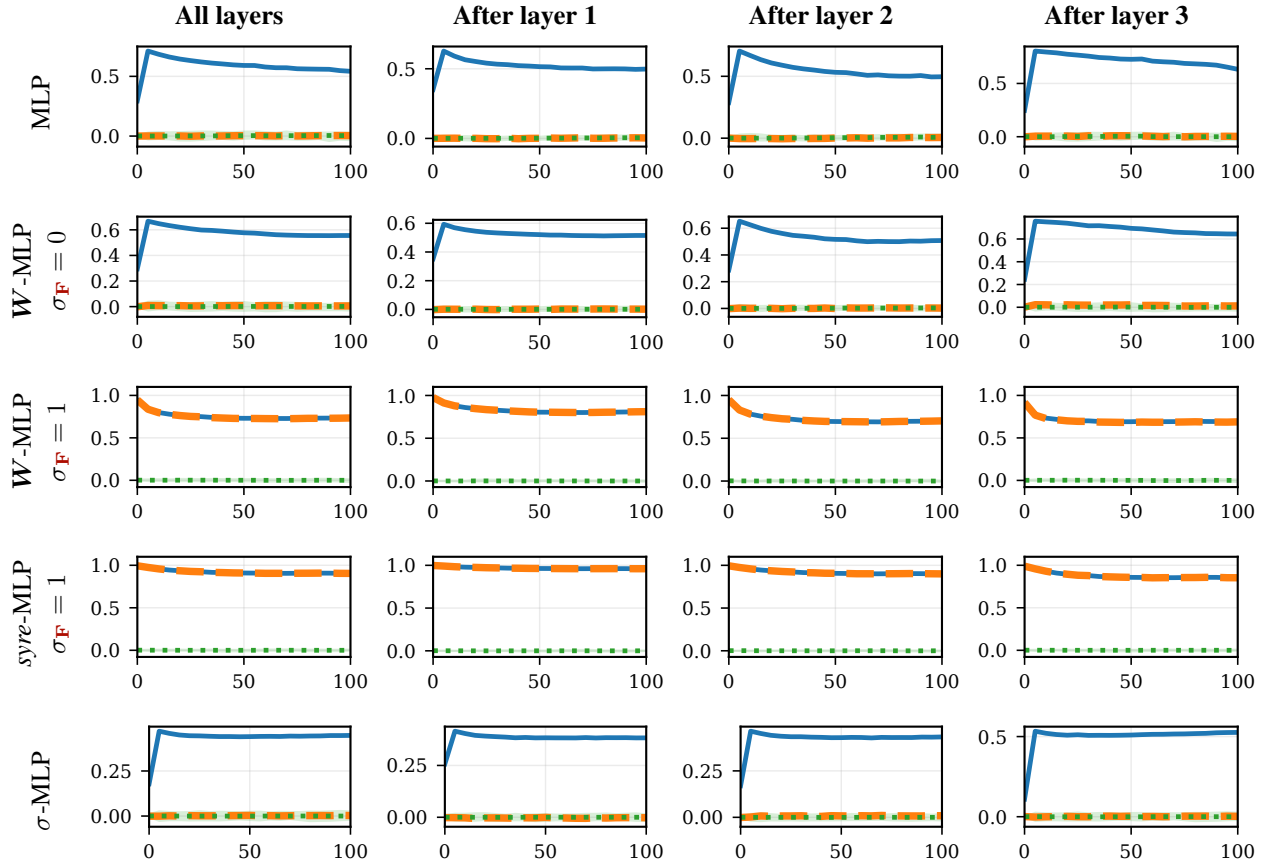


Figure 8. Layerwise activation matching objectives over epochs, separated by layer and by different symmetry-breaking architectures (including the σ -MLP variant, see Lim et al. (2024b)). We plot optimal, identity, and random permutations.

Further, in Figs. 9 and 10, we plot layerwise versions of the $\sigma_{\mathbf{F}}$ sweep for both MLPs on MNIST and ResNets on CIFAR-10 respectively. We also sweep the sparsity (proportion of fixed weights) in Fig. 11 to measure its effect on symmetry breaking, and find that sparsity alone without fixed weights does not yield strong neuron identifiability.

Beyond Structural Symmetries: Linear Mode Connectivity via Neuron Identifiability




Figure 9. Activation matching objectives per layer (W -MLP on MNIST), sweeping over fixed weight scale σ_F .




Figure 10. Activation matching objectives separated by layer/activation matching point (W -ResNet on CIFAR-10). In our W -ResNets, activation matching points lie within each of the three block groups’ (BG) residual streams (used to estimate a global permutation for the residual stream that multiple layers write into), and between the two lin. layers of the three inner two-layer MLPs within each block group.

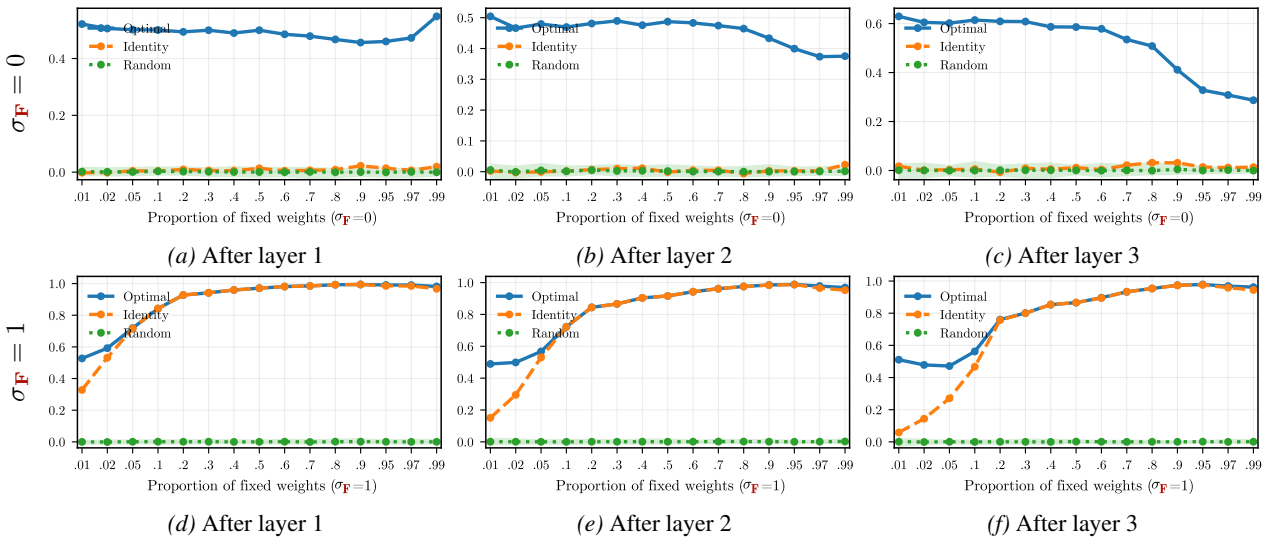


Figure 11. Activation matching objectives per layer (W -asymmetric MLPs on MNIST), sweeping over sparsity parameter (proportion of fixed weights) for both $\sigma_F = 0$ (standard sparse training) and $\sigma_F = 1$.

H.2. Empirical Validation of Center Separation Rate

We also examine how the minimum distance between projected neuron centers changes with layer width m . For randomly sampled \mathbf{F} we observe the expected scaling behavior from [Thm. C.12](#). To this end, we sweep the hidden dimension m and intrinsic dimension $k \in \{2, 8, 32\}$ on the Gaussian mixture dataset. The constant c_k in [Fig. 12](#) is chosen per k such that the rate is anchored at the respective γ_{out} for $m = 16$. [Thm. C.12](#) predicts a rate of $\Theta(\sigma_{\mathbf{F}}\sqrt{k}m^{-2/k})$, where $\sigma_{\mathbf{F}}\sqrt{k}$ and further constants hidden by Θ are taken care of by the anchoring. This helps explain that neuron swaps are only plausible when projected centers are close, which is mainly the case for small intrinsic dimension or extremely wide layers. For realistic values of k , layers would need to be astronomically wide for such close centers to become likely.




Figure 12. Minimum pairwise projected center distance γ_{out} (solid, mean \pm std) vs. predicted asymptotic rate $c_k m^{-2/k}$ (dotted).

H.3. Neuron Swap Costs via Ridge Regression and Additional Pairwise Neuron Swap Costs

In [§ 5](#), we estimated the realization cost, and derived from that the neuron swap cost, via the Mahalanobis distance [\(5\)](#) in a PCA space with 90% explained variance. [Appendix § E.1](#) suggests a different way to estimate realization costs, via a ridge regression objective [Def. E.3](#). We compare the two methods in the following figures in more detail, separated by architecture, layer, and estimation method: We provide distribution box plots, comparing neuron swap costs for the four tested MLP-based architectures (MNIST) / zero fixed weights vs. center-dominated regime \mathbf{W} -asym. MLPs (Gaussian mixture data), and detailed pairwise neuron swap cost plots showing large/small/negative costs.

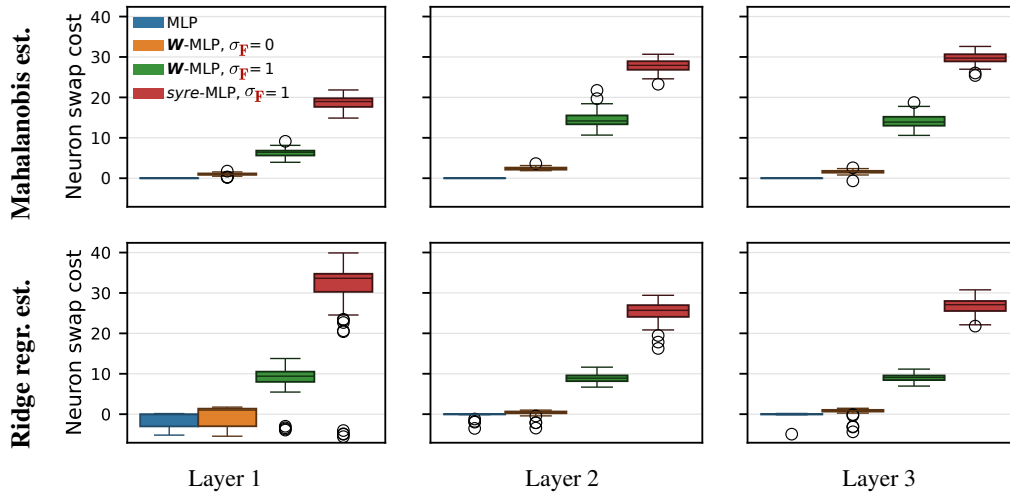


Figure 13. Distributions of neuron swap costs by architecture on MNIST. Plotted are signed sqrt. transformed $\Delta_{(ij)}^{\text{out}}$ for disjoint consecutive pairs (i, j) of neurons, estimated via Mahalanobis distance [\(5\)](#) and ridge regression ([Def. E.3](#), $\beta = 0.01$).




Figure 14. Neuron swap costs for \mathbf{W} -MLPs with $\sigma_{\mathbf{F}} \in \{0, 1\}$ on Gaussian mixture data. Plotted are signed sqrt. transformed $\Delta_{(ij)}^{\text{out}}$ for disj. consecutive pairs (i, j) of neurons, estimated via Mahalanobis dist. [\(5\)](#) and ridge regression ([Def. E.3](#), $\beta = 0.01$).




Figure 15. Neuron swap costs $\Delta_{(ij)}^{\text{out}}$ (signed square-root transformed) estimated via **Mahalanobis distance** on **Gaussian mixture data**, for varying *intrinsic dimension* k and fixed *weight scale* $\sigma_{\mathbf{F}}$ (same as Fig. 5).




Figure 16. Neuron swap costs $\Delta_{(ij)}^{\text{out}}$ (signed square-root transformed) estimated via **ridge regression**, $\beta = 0.01$ (see Def. E.3) on **Gaussian mixture data**, for varying *intrinsic dimension* k and fixed *weight scale* $\sigma_{\mathbf{F}}$.

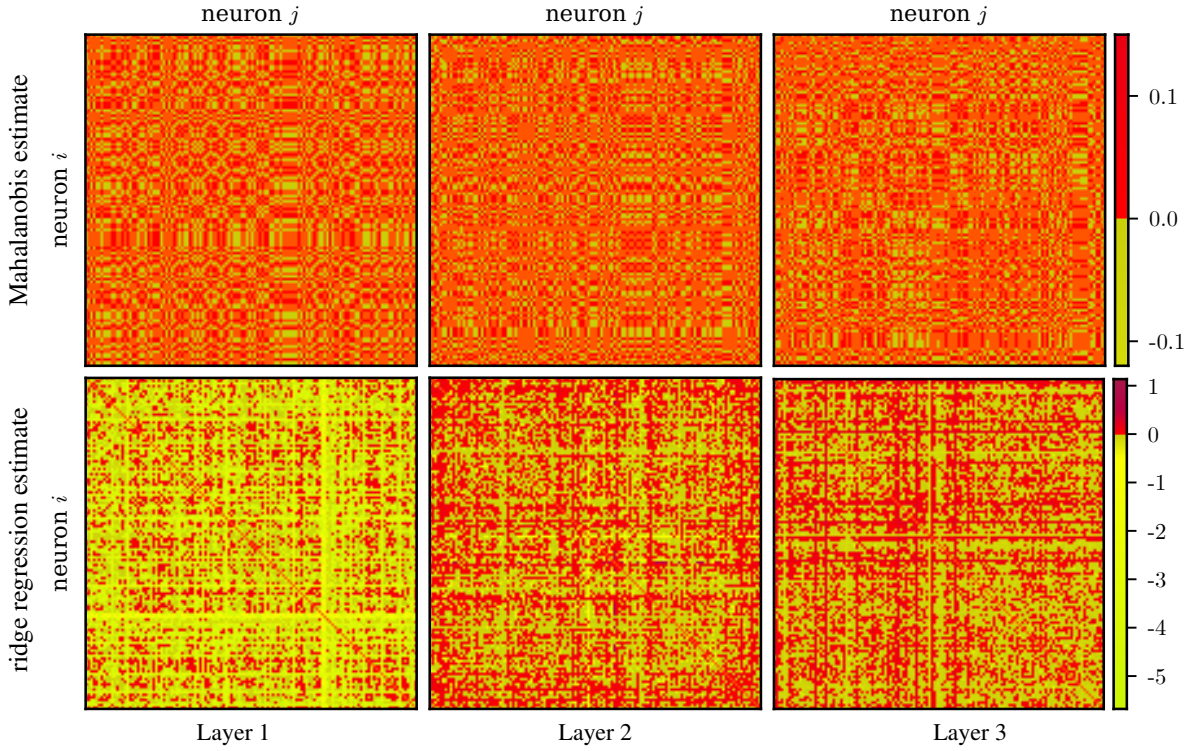


Figure 17. Neuron swap costs $\Delta_{(ij)}^{\text{out}}$ (signed square-root transformed) per layer for **standard MLP** trained on **MNIST**, estimated via *Mahalanobis distance* (5) vs. *ridge regression* (Def. E.3, $\beta = 0.01$). Layers 1 – 3 have hidden dimension 512, shown are costs for first 128×128 neuron pairs.




Figure 18. Neuron swap costs $\Delta_{(ij)}^{\text{out}}$ (signed square-root transformed) per layer for **W-asymmetric MLP** ($\sigma_{\mathbf{F}} = 0$) trained on **MNIST**, estimated via *Mahalanobis distance* (5) and *ridge regression* (Def. E.3, $\beta = 0.01$). Layers 1 – 3 have hidden dimension 512, shown are costs for first 128×128 neuron pairs.

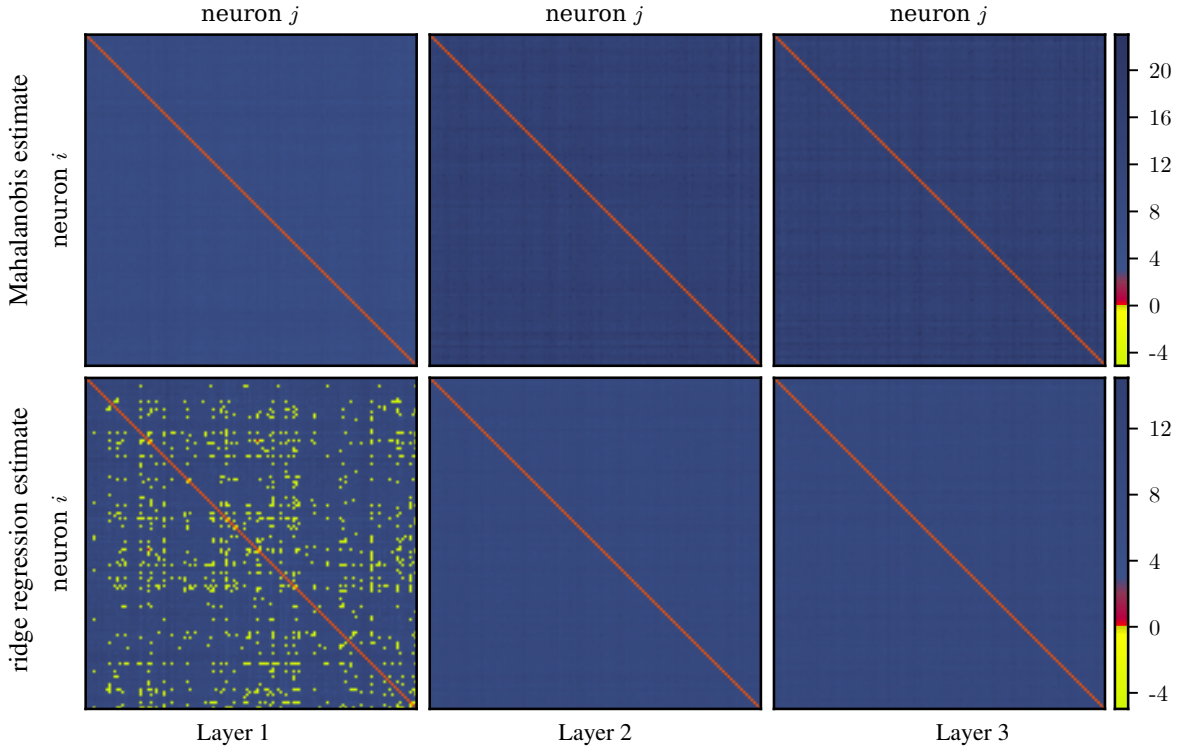


Figure 19. Neuron swap costs $\Delta_{(ij)}^{\text{out}}$ (signed square-root transformed) per layer for **W-asymmetric MLP** ($\sigma_{\mathbf{F}} = 1$) trained on **MNIST**, estimated via *Mahalanobis distance* (5) and *ridge regression* (Def. E.3, $\beta = 0.01$). Layers 1 – 3 have hidden dimension 512, shown are costs for first 128×128 neuron pairs.

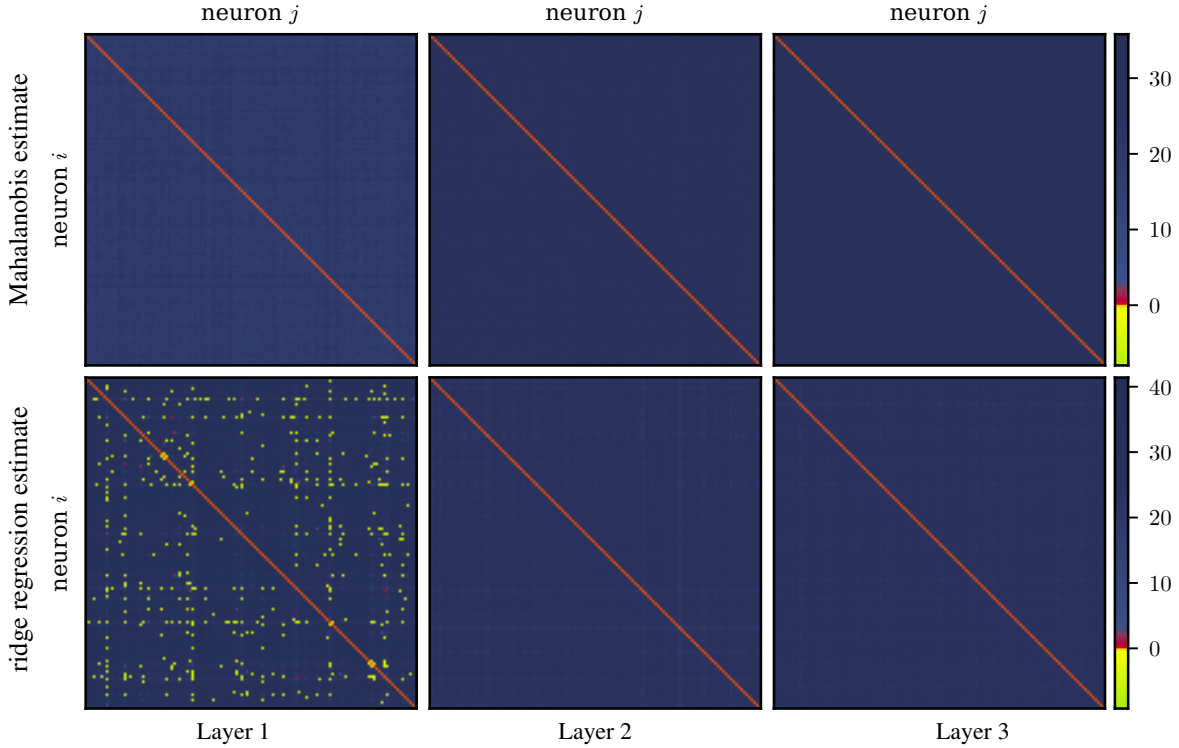


Figure 20. Neuron swap costs $\Delta_{(ij)}^{\text{out}}$ (signed square-root transformed) per layer for **syre-MLP** ($\sigma_{\mathbf{F}} = 1$) trained on **MNIST**, estimated via *Mahalanobis distance* (5) and *ridge regression* (Def. E.3, $\beta = 0.01$). Layers 1 – 3 have hidden dimension 512, shown are costs for first 128×128 neuron pairs.

H.4. Transformers

§ 5 studies the effect of neural parameter symmetries on optimization and linear mode connectivity for MLPs and ResNets. We additionally examine Vision Transformers (ViTs) (Dosovitskiy et al., 2021) on CIFAR-10 to test whether the same symmetry breaking mechanism extends to attention-based architectures. A transformer is composed of linear maps, including the query, key, value, and output projections in multi-head self-attention, the two linear layers in each feed-forward block, the patch embedding layer, and the final classification head. Transformers are an especially interesting extension because attention admits richer $GL(\mathbb{R}^n)$ symmetries beyond the permutation setting in § 5. We note that our goal is not to introduce a highly tuned state-of-the-art ViT, but to use a trainable transformer baseline whose training accuracy reached high performance. The goal is to investigate whether the same parameter-level intervention can reduce the interpolation barrier.

The ViT setting is more challenging for several reasons. First, transformer symmetries are not limited to hidden-unit or channel permutations. Multi-head attention introduces head-permutation symmetries, while the feed-forward blocks contain MLP-like hidden-unit symmetries. Second, LayerNorm, residual connections, positional embeddings, and the class token may interact with these symmetries in ways not captured by a simple linear-map masking scheme. Third, our implementation applies the \mathbf{W} -asymmetric parameterization directly to the transformer’s linear maps. This is a natural extension of the original construction, but it may not be the optimal transformer-specific way to break symmetries. For example, breaking individual entries of \mathbf{W}_Q , \mathbf{W}_K , \mathbf{W}_V , \mathbf{W}_O may not fully eliminate head-level equivalences unless the asymmetry is coordinated with the internal head structure. Similarly, asymmetrizing all linear maps may introduce additional optimization difficulty because fixed entries constrain the parameterization throughout training.

We use a CIFAR-10 ViT with patch size 4, embedding dimension 384, depth 8, 6 attention heads, and MLP ratio 4. Following the experimental setup in § 5, we apply a random 80%/10%/10% train/val/test split. For each linear map, we replace the ordinary trainable weight matrix \mathbf{W} with a \mathbf{W} -asymmetric parameterization in the form of (1). The fixed entries are sampled once at initialization, and the number of fixed entries is 8. This construction is a natural extension of the \mathbf{W} -asymmetric parameterization used for MLPs and ResNets, but it is not necessarily the optimal way to break transformer symmetries. For example, a more tailored approach might use head-aware masks or a design that explicitly respects the coupled structure of \mathbf{W}_Q , \mathbf{W}_K , \mathbf{W}_V , \mathbf{W}_O . We leave such transformer-specific variants for future work. We train the model using AdamW (Loshchilov & Hutter, 2019) with cosine learning-rate decay for 200 epochs using a batch size of 128. In our preliminary tuning, we found that larger fixed-weight scales can make transformer optimization more difficult. In particular, settings such as $\sigma_{\mathbf{F}} = 2$ were harder to train. When their scale is too large, they can dominate the trainable entries, alter the scale of attention logits and MLP activations, and reduce the flexibility of the model during optimization. Therefore, unlike in an idealized symmetry breaking setting, increasing $\sigma_{\mathbf{F}}$ should not be interpreted as a monotonic improvement in practice.

Fig. 21 reports the training accuracy along the linear interpolation path between independently trained transformer endpoints. We see that the standard transformer exhibits a pronounced interpolation barrier and the \mathbf{W} -Transformer with $\sigma_{\mathbf{F}} = 0$ closely matches this behavior, suggesting that masking alone, when the fixed entries are zero, is not sufficient to meaningfully improve unaligned LMC. In contrast, the \mathbf{W} -Transformer with $\sigma_{\mathbf{F}} = 1$ yields a noticeably flatter interpolation curve and substantially higher mid-point accuracy. However, the interpolation curve still remains well below the endpoint accuracy. The improvement at $\sigma_{\mathbf{F}} = 1$ suggests that nonzero \mathbf{W} -asymmetric fixed entries can reduce the unaligned interpolation barrier in transformers, qualitatively matching the mechanism observed for MLPs and ResNets in § 5. At the same time, the remaining barrier indicates that the current entrywise \mathbf{W} -asymmetric construction does not fully resolve the symmetries or optimization effects present in transformer architectures. This may be due to residual attention-head symmetries, richer $GL(\mathbb{R}^n)$ invariances, interactions with LayerNorm and residual connections, or the fact that our masks are inherited from the MLP and ResNet setting rather than designed specifically for attention. Thus, this provides evidence that the mechanism identified by our theory carries over qualitatively to transformers, while also motivating future work on transformer-specific symmetry breaking schemes.




Figure 21. LMC of Transformer and \mathbf{W} -Transformer on CIFAR-10, measured by training accuracy along the interpolation path.

H.5. Subspace Coherence of MNIST Inputs and Hidden Representations

We empirically measure the subspace coherence $\nu(\mathcal{U})$ of model inputs, hidden representations and the final output for standard and asymmetric MLPs on MNIST data variants over training. Fig. 22 shows the results: in particular, variants

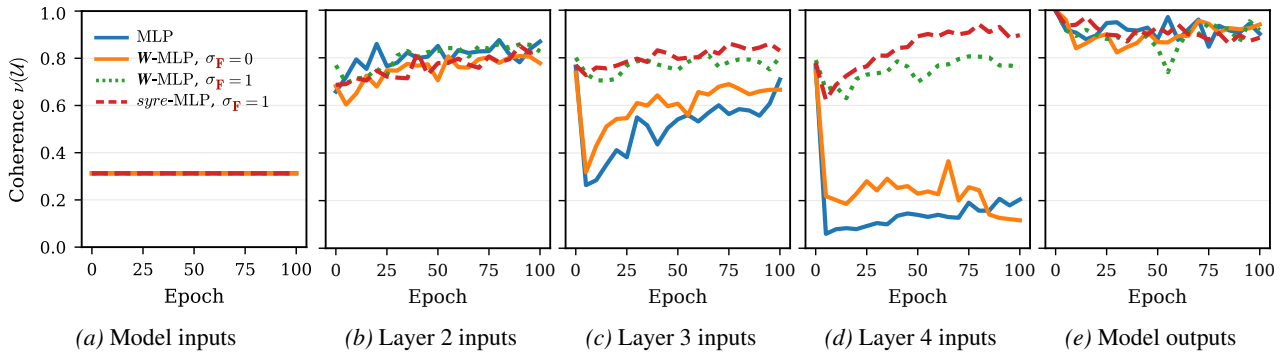


Figure 22. Subspace coherences $\nu(\mathcal{U})$ computed for model inputs, outputs, and hidden representations, on MNIST data, for standard and asymmetric MLP variants over training.

without effective symmetry breaking (MLP and \mathcal{W} -asymmetric MLP with zero fixed weights) exhibit low subspace coherence, in particular in later layers, while \mathcal{W} -asymmetric MLPs and *syre*-MLPs with large fixed weights both tend to exhibit higher subspace coherence values.