

Connecting Parameter Magnitudes and Hessian Eigenspaces at Scale using Sketched Methods

Andres Fernandez
Tübingen AI Center
University of Tübingen

a.fernandez@uni-tuebingen.de

Frank Schneider
Tübingen AI Center
University of Tübingen

f.schneider@uni-tuebingen.de

Maren Mahsereci
Yahoo Research

maren.mahsereci@uni-tuebingen.de

Philipp Hennig
Tübingen AI Center
University of Tübingen

philipp.hennig@uni-tuebingen.de

Reviewed on OpenReview: <https://openreview.net/forum?id=yGGoOVpBVP>

Abstract

Recently, it has been observed that when training a deep neural net with SGD, the majority of the loss landscape’s curvature quickly concentrates in a tiny *top* eigenspace of the loss Hessian, which remains largely stable thereafter. Independently, it has been shown that successful magnitude pruning masks for deep neural nets emerge early in training and remain stable thereafter. In this work, we study these two phenomena jointly and show that they are connected: We develop a methodology to measure the similarity between arbitrary parameter masks and Hessian eigenspaces via Grassmannian metrics. We identify *overlap* as the most useful such metric due to its interpretability and stability. To compute *overlap*, we develop a matrix-free algorithm based on sketched SVDs that allows us to compute over 1000 Hessian eigenpairs for nets with over 10M parameters—an unprecedented scale by several orders of magnitude. Our experiments reveal an *overlap* between magnitude parameter masks and top Hessian eigenspaces consistently higher than chance-level, and that this effect gets accentuated for larger network sizes. This result indicates that *top Hessian eigenvectors tend to be concentrated around larger parameters*, or equivalently, that *larger parameters tend to align with directions of larger loss curvature*. Our work provides a methodology to approximate and analyze deep learning Hessians at scale, as well as a novel insight on the structure of their eigenspace.

1 Introduction

Deep learning (DL) benefits from overparametrization; but not all parameters are equally important. Often, a substantial portion of parameters can be *pruned*, i.e. removed, without compromising the model’s performance (see Blalock et al., 2020; Hoefler et al., 2021). One efficient and popular method to identify these subnetworks is via parameter magnitude (Han et al., 2015). Interestingly, these subnetworks materialize very early in training (Frankle & Carbin, 2019), and once they emerge, their topology stops changing significantly (Achille et al., 2019; You et al., 2020). In other words, competitive subnetworks *crystallize early* in training and remain *stable* thereafter (Section 2.2, Figures 15 and 16).

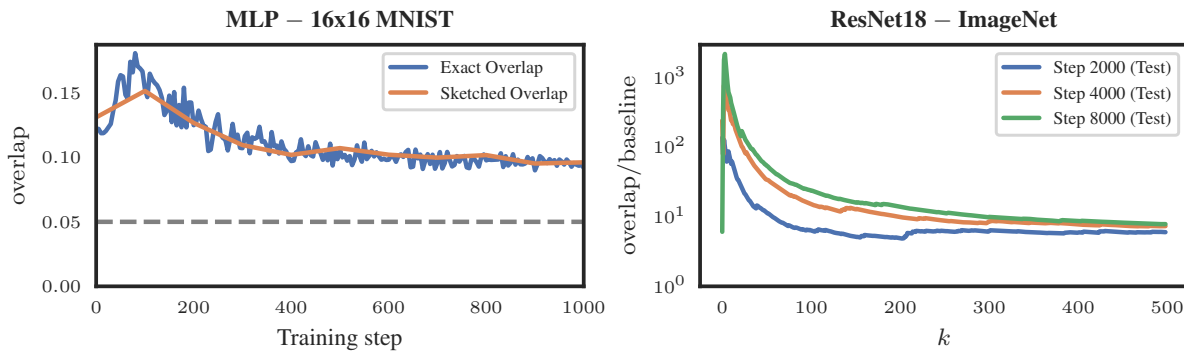


Figure 1: Overlap between top- k parameter magnitude masks and top- k Hessian eigenspaces is consistently and substantially above random chance. (Left) Measurements for a 7030-parameter network trained on 16×16 downsampled MNIST until convergence (see Section 6.1 and Figure 15): Exact *overlap* between top- k parameters and eigenvectors with $k = 350$ (—), approximate *overlap* via sketched eigendecomposition (—), and chance-level baseline (--) (see Section 4.2). Note how *overlap* is larger than chance, and sketched *overlap* is a good approximation. (Right) Lines show the ratio *sketched overlap vs. chance-level baseline* for a model with $>11\text{M}$ parameters trained on IMAGENET (see Section 6.2), at three different points during training, and as a function of k . Note how *overlap* is always higher than baseline, up to a factor of 1000.

Concurrent research focuses on the loss Hessian, characterizing the loss landscape’s curvature. Multiple studies found that empirically the Hessian spectrum separates into two parts: The *bulk* subspace, with many, near-zero eigenvalues, and the *top* subspace with a few eigenvalues of significantly larger magnitude (e.g. Dauphin et al., 2014; Sagun et al., 2018). Crucially, Gur-Ari et al. (2018) observed that, after initial training iterations, the gradient predominantly lies in this top subspace where it remains relatively stable throughout training. Analogous to the work on parameter masks, this indicates that the top Hessian eigenspace *crystallizes early* in training and tends to remain *stable* (Section 2.1, Figures 15 and 16).

In this work, we explore the connection between these largely independent lines of research—both reporting *early crystallization* of substructures. Our **contributions** include:

1. We establish a common theoretical ground for comparing binary parameter masks m_k (which select k parameters and discard the rest) and top- k eigenspaces of a Hessian. Both can be cast as rank- k orthogonal matrices belonging to the same Stiefel manifold (Section 3). This allows a direct comparison of their spans using *Grassmannian metrics*¹ (see Section 4.1). We review popular Grassmannian metrics and identify *overlap* (Eq. (2)) as the most meaningful metric (Section 4.2).
2. To efficiently compute *overlap*, we develop SEIGH (Section 5 and Alg. 2), a matrix-free eigendecomposition based on sketched SVDs (Tropp et al., 2019). Our open source implementation² allows to compute top- k Hessian eigendecompositions for $k > 10^3$ on neural networks with over 10M parameters, an unprecedented scale by orders of magnitude.
3. We provide empirical evidence that the similarities between the spaces induced by parameter magnitude masks and top- k Hessian eigenspaces are consistently and substantially higher than random chance (Figure 1). This suggests that, in DL, top Hessian eigenvectors tend to be concentrated around larger parameters throughout the training process (Section 6).

Seminal work on “Optimal Brain Damage” (LeCun et al., 1989) established a theoretical link between *individual parameter values* and loss curvature, but only *at convergence*. Recent work often focuses on loss curvature for *layer-wise* parameter groups (e.g. Sankar et al., 2021), or involves just a few ($k \approx 10$) Hessian

¹Grassmannians are manifolds of low-dimensional subspaces of a given vector space.

²https://github.com/andres-fr/hessian_overlap

eigenpairs without emphasis on parameters (e.g. Gur-Ari et al., 2018; Pappan, 2020; Dangel et al., 2022). In contrast, we connect Hessian *eigenspaces* to the *spans* of arbitrary parameter subsets, e.g. magnitude masks, throughout the *entire training*—including the relevant early stage—in an interpretable and scalable manner. We can thus jointly study the connection between these two phenomena and provide novel insights into a neural net’s Hessian structure. This connection implies being able to approximate expensive Hessian quantities via cheap parameter inspection, which bears potential for downstream tasks where the Hessian plays a prominent role, such as optimization (Martens, 2016), pruning (LeCun et al., 1989), or uncertainty estimation (Kristiadi et al., 2021).

2 Hessians, Parameters and Early Crystallization

We consider a supervised classification/regression setup with train set $\mathcal{D}_{\text{train}} := \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$ of labeled data $(\mathbf{x}_n, \mathbf{y}_n) \in \mathbb{X} \times \mathbb{Y}$, from an unknown data-generating distribution P . The neural net $f_{\boldsymbol{\theta}}(\mathbf{x}) : \mathbb{X} \rightarrow \mathbb{Y}$ maps inputs \mathbf{x} to predictions $\hat{\mathbf{y}}$ via parameters $\boldsymbol{\theta} \in \mathbb{R}^D$. A loss function $\ell : \mathbb{Y} \times \mathbb{Y} \rightarrow \mathbb{R}_{\geq 0}$ penalizes differences between prediction $\hat{\mathbf{y}}$ and true label \mathbf{y} . The goal is to minimize the inaccessible *risk* $L_P(\boldsymbol{\theta}) := \int \ell(f_{\boldsymbol{\theta}}(\mathbf{x}), \mathbf{y}) dP$ via the proxy *empirical risk* $L_{\mathcal{D}_{\text{train}}}(\boldsymbol{\theta}) := \frac{1}{N} \sum_{n=1}^N \ell(f_{\boldsymbol{\theta}}(\mathbf{x}_n), \mathbf{y}_n)$. For large N , we approximate $L_{\mathcal{D}_{\text{train}}}(\boldsymbol{\theta})$ using *mini-batches* $\mathcal{B} \stackrel{\text{iid}}{\sim} \mathcal{D}_{\text{train}}$ of $B \ll N$ samples. For f, ℓ twice differentiable, we use the *gradient* $\mathbf{g}(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\theta}} L(\boldsymbol{\theta}) \in \mathbb{R}^D$ and the *Hessian* $\mathbf{H}(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\theta}}^2 L(\boldsymbol{\theta}) \in \mathbb{R}^{D \times D}$. With \mathbf{g}, \mathbf{H} we refer to *any* loss gradient or Hessian. To emphasize the data domain, we use a subindex, e.g. $\mathbf{H}_{\mathcal{B}}$ for the Hessian of the mini-batch loss.

2.1 The Hessian in Deep Learning

The Hessian plays a prominent role in DL applications. A useful characterization of \mathbf{H} is its eigendecomposition $\mathbf{H} = \mathbf{U} \boldsymbol{\Lambda} \mathbf{U}^{\top} = \sum_{i=1}^D \lambda_i \mathbf{u}_i \mathbf{u}_i^{\top}$. Here, \mathbf{U} is orthogonal, with *eigenvectors* \mathbf{u}_i , and $\boldsymbol{\Lambda}$ is diagonal and real-valued with (ordered) *eigenvalues* $|\lambda_1| \geq \dots \geq |\lambda_D|$. We call $\mathbf{U}^{(k)} := \{\mathbf{u}_i\}_{i=1}^k$ the *top- k eigenbasis* of \mathbf{H} , and $\text{span}(\mathbf{U}^{(k)})$ the *top- k eigenspace*. The top- k eigendecomposition $\mathbf{H}^{(k)} = \sum_{i=1}^k \lambda_i \mathbf{u}_i \mathbf{u}_i^{\top}$ minimizes $\|\mathbf{H} - \mathbf{H}^{(k)}\|$ for all unitarily invariant norms (Golub & Van Loan, 2013, Th. 2.4.8).

Recent literature has extensively investigated the Hessian **spectrum** of neural nets, revealing that the eigenvalues are typically clustered into two parts: (1) The *bulk* of eigenvalues with near-zero magnitude and (2) a few *top* eigenvalues with significantly larger magnitude (e.g. Sagun et al., 2018; Pappan, 2019). Thus, \mathbf{H} can be well-approximated by its top eigenpairs. For the Hessian **eigenspace**, Li et al. (2018a) showed that projecting the whole space onto a few random, fixed dimensions still allows Stochastic Gradient Descent (SGD) to perform competitively—provided enough dimensions are given—leading to the idea of an intrinsic dimensionality of problems. Gur-Ari et al. (2018) observed that this restriction to a lower-dimensional, fixed subspace seems to happen spontaneously anyway: After a brief initial training period, the gradient predominately lies within a small subspace spanned by the few top Hessian eigenvectors and this space changes little over the remaining training process. However, these phenomena might rely on optimizer and model choices (Li et al., 2018a; Ghorbani et al., 2019).

One fundamental issue greatly limiting the use of \mathbf{H} for DL is its prohibitively large size, with D^2 entries. Consequently, most scalable methods are *matrix-free*, relying on Hessian-Vector Products (HVPs) to compute linear maps $\mathbf{w} = \mathbf{H}\mathbf{v}$ in just $\mathcal{O}(D)$ memory and time (Pearlmutter, 1994). Examples are the computation of spectral densities (Yao et al., 2020; Pappan, 2018) or top- k eigendecompositions (Gur-Ari et al., 2018; Dangel et al., 2022). To make Hessian properties more accessible, specialized DL libraries have been developed recently (e.g. Dangel et al., 2020; Yao et al., 2020; Elsayed & Mahmood, 2022; Dangel et al., 2025), but efficiently accessing large portions of the Hessian remains a major challenge (see Section 5).

2.2 Parameter Masks and Early Crystallization

Binary parameter masks $\mathbf{m}_k \in \mathbb{B}^D$, with $\mathbb{B} = \{0, 1\}$, consisting of k ones and the rest zeros, can be used to define subsets of parameters from a given model. A mask is *k -sparse* if it has k non-zero elements. We measure mask sparsity using the ratio $\rho := k/D$. For non-binary vectors $\mathbf{v} \in \mathbb{R}^D$, we instead measure whether a small subset of indices ι contains a large proportion of the total norm of the vector. When ι is known, this can be directly expressed as the ratio: $\kappa(\mathbf{v}) := \|\mathbf{v}_{\iota}\|_2^2 / \|\mathbf{v}\|_2^2$ (see Hurley & Rickard, 2009). A popular choice

for \mathbf{m}_k is to take the k -largest parameters by magnitude: it is computationally cheap, and has been shown to be effective for neural net pruning (Han et al., 2015), leading to smaller models that can often achieve competitive performances with $\rho \in [1\%, 10\%]$ (e.g. Gale et al., 2019; Blalock et al., 2020; Hoefer et al., 2021).

In this work, we do not perform any pruning, but rather focus on a key feature of parameter magnitude masks. Not only can they lead to competitive performance when used for pruning and be found very early in training (Frankle & Carbin, 2019; Frankle et al., 2020) but they also *stabilize soon afterwards*: You et al. (2020) compared Hamming distances between periodically extracted pruning masks and found they stop changing early in training, aligning with the loss of information plasticity reported in Achille et al. (2019). This *early crystallization* bears a striking parallel with observations made for the top Hessian subspace, and motivates our investigation: Are these two phenomena connected? How would one measure this connection? What would such a connection tell us about the parameters and the Hessian?

3 Parameter Masks as Orthogonal Projections

Quantifying the connection between parameter magnitude masks and top- k Hessian eigenspaces—both exhibiting early crystallization—requires a way to relate a mask to a subspace. In this section, we observe that both $\mathbf{U}^{(k)}$ (the top- k eigenbasis) and any k -sparse mask³ \mathbf{m}_k can be characterized as elements of the same compact Stiefel manifold $\mathbb{O}^{D \times k} = \{\mathbf{Q} : \mathbf{Q} \in \mathbb{R}^{D \times k}, \mathbf{Q}^\top \mathbf{Q} = \mathbf{I}_k\}$ (Absil et al., 2004), where $\mathbf{I}_k \in \mathbb{R}^{k \times k}$ denotes the identity. Specifically for masks, we further consider the subset $\mathbb{M}^{D \times k} \subset \mathbb{O}^{D \times k}$, with columns having exactly one 1, each row at most one 1, and the rest is zeros.

Reordering parameters: Recall from Section 2.1 that the Hessian eigenvalues are sorted by descending magnitude, exposing a single cutting point between *top* and *bulk* eigenspace at dimension k . To simplify notation, we impose a similar cutting point to the parameters, by defining a permutation matrix $\mathbf{P} \in \mathbb{M}^{D \times D}$ for any given mask \mathbf{m}_k , such that the mask entries are grouped in *selected* (i.e. k entries with $m_i = 1$) and *discarded* (i.e. $m_i = 0$), i.e. $\tilde{\mathbf{m}} := \mathbf{P}^\top \mathbf{m} = (1, \dots, 1, 0, \dots, 0)$. For a parameter magnitude mask, for example, we can define \mathbf{P} to order the parameters by nonincreasing magnitude such that $i \leq j \Rightarrow |\mathbf{P}\mathbf{m}|_i \geq |\mathbf{P}\mathbf{m}|_j$. We can now permute the parameters $\tilde{\boldsymbol{\theta}} = \mathbf{P}^\top \boldsymbol{\theta}$, as well as the Hessian rows and columns $\tilde{\mathbf{H}} := \mathbf{P}^\top \mathbf{H} \mathbf{P} = \tilde{\mathbf{U}} \boldsymbol{\Lambda} \tilde{\mathbf{U}}^\top$. This has no loss of generality, since $(\mathbf{m}, \boldsymbol{\theta}, \mathbf{H}) \cong (\tilde{\mathbf{m}}, \tilde{\boldsymbol{\theta}}, \tilde{\mathbf{H}})$ is an isomorphism, \mathbf{H} and $\tilde{\mathbf{H}}$ are *similar*, and the loss curvature remains unaltered ($\tilde{\boldsymbol{\theta}}^\top \tilde{\mathbf{H}} \tilde{\boldsymbol{\theta}} = \boldsymbol{\theta}^\top \mathbf{H} \boldsymbol{\theta}$). Then, any permuted k -sparse masking operation can be expressed as:

$$\mathbf{P}^\top (\mathbf{m}_k \circ \boldsymbol{\theta}) = \tilde{\mathbf{m}}_k \circ \tilde{\boldsymbol{\theta}} = \underbrace{\begin{pmatrix} \mathbf{I}_k & 0 \\ 0 & 0 \end{pmatrix}}_{:= \tilde{\Phi}} \tilde{\boldsymbol{\theta}} =: \mathbf{I}_{D,k} \mathbf{I}_{D,k}^\top \tilde{\boldsymbol{\theta}},$$

where $\mathbf{I}_{D,k} = \begin{pmatrix} \mathbf{I}_k \\ 0 \end{pmatrix} \in \mathbb{R}^{D \times k}$, which is clearly an element of the (binary) Stiefel manifold $\mathbb{M}^{D \times k}$. Also note that this holds for any \mathbf{P} , not only for those defined for parameter magnitude masks.

Partitioning $\tilde{\mathbf{H}}$: Consider now the following partition of the reordered Hessian with $\mathbf{V}, \mathbf{D} \in \mathbb{R}^{k \times k}$, $\tilde{\mathbf{W}}, \tilde{\mathbf{E}} \in \mathbb{R}^{(D-k) \times (D-k)}$, and $\tilde{\mathbf{V}}, \tilde{\mathbf{W}}^\top \in \mathbb{R}^{(D-k) \times (k)}$:

$$\tilde{\mathbf{H}} := \underbrace{\begin{pmatrix} \mathbf{V} & \mathbf{W} \\ \tilde{\mathbf{V}} & \tilde{\mathbf{W}} \end{pmatrix}}_{\tilde{\mathbf{U}} = \mathbf{P}^\top \mathbf{U}} \underbrace{\begin{pmatrix} \mathbf{D} & \\ & \tilde{\mathbf{E}} \end{pmatrix}}_{\boldsymbol{\Lambda}} \underbrace{\begin{pmatrix} \mathbf{V}^\top & \tilde{\mathbf{V}}^\top \\ \mathbf{W}^\top & \tilde{\mathbf{W}}^\top \end{pmatrix}}_{\tilde{\mathbf{U}}^\top = \mathbf{U}^\top \mathbf{P}}. \quad (1)$$

With this partition, $\tilde{\mathbf{U}}^{(k)} =: \begin{pmatrix} \mathbf{V} \\ \tilde{\mathbf{V}} \end{pmatrix}$ is the *top- k* Hessian eigenbasis, and $\begin{pmatrix} \mathbf{W} \\ \tilde{\mathbf{W}} \end{pmatrix}$ the *bulk* eigenbasis. Conversely, the rows of $(\mathbf{V} | \mathbf{W})$ correspond to the *selected* parameters, and $(\tilde{\mathbf{V}} | \tilde{\mathbf{W}})$ to the *discarded* ones. Since $\tilde{\mathbf{U}}$ is orthogonal, we have $\tilde{\mathbf{U}}^{(k)} \in \mathbb{O}^{D \times k}$, i.e. an element of the same Stiefel manifold.

4 Measuring Subspace Similarity via Grassmannian Metrics

We set out now to quantify the similarity between a k -sparse parameter mask and the top- k Hessian eigenspace. Specifically, we are only interested in the similarity of their spans, since the *spaces* are the ones

³This can be generalized to masks and eigenspaces with different k using Schubert varieties (Ye & Lim, 2016).

reported to undergo early crystallization. To connect parameter spaces to loss curvature, one may consider randomly perturbing the parameters, and empirically measuring the impact on the loss. This turns out to be problematic due to the non-PSD nature of \mathbf{H} (see Appendix A.1). Various workarounds to obtain nonnegative measurements can lead to promising results, but they also entail tradeoffs (see Appendix A.2). Instead, Grassmannian metrics provide a natural and theoretically grounded way of achieving our goal, given that k -sparse parameter masks and the top- k Hessian eigenspace can both be cast as elements of the same Stiefel manifold (Section 3). We review popular Grassmannian metrics in Section 4.1, and analyze them in more depth in Section 4.2, finding that the *overlap* metric, highlighted here, is both interpretable and stable:

$$\text{overlap}(\mathbf{Q}_1, \mathbf{Q}_2) = \frac{1}{k} \|\mathbf{Q}_1^\top \mathbf{Q}_2\|_F^2 = \frac{1}{k} \|\cos(\boldsymbol{\sigma})\|_F^2 \in [0, 1]. \quad (2)$$

Intuitively, we can see how a zero *overlap* between \mathbf{Q}_i and \mathbf{Q}_j indicates that the spans of both matrices are orthogonal to each other, and an *overlap* of 1 indicates that both spaces are identical. We can also interpret *overlap* as the rotation-invariant cosine similarity between \mathbf{Q}_1 and \mathbf{Q}_2 , as induced by the standard Euclidean metric on matrices.

4.1 Grassmann Manifolds and their Metrics

Grassmann manifolds are extensively studied (e.g. Witten, 1988; Absil et al., 2004; Bendokat et al., 2020) and have recently been applied to DL (e.g. Gur-Ari et al., 2018; Zhang et al., 2018; Dangel et al., 2022). A Grassmann manifold $\mathcal{G}_{k,D}$ is the set of all k -dimensional subspaces of a given D -dimensional Euclidean space. Two orthogonal matrices $\mathbf{Q}_1, \mathbf{Q}_2 \in \mathbb{O}^{D \times k}$ map to the same element $\mathfrak{g} \in \mathcal{G}_{k,D}$ if and only if their column span is identical. Thus, the subset of all matrices in $\mathbb{O}^{D \times k}$ that map to $\mathfrak{g}_i = \text{span}(\mathbf{Q}_i)$ forms an *equivalence class* $\mathcal{S}_i^\mathbb{O} := \{\mathbf{Q}_j: \mathbf{Q}_j \mathbf{Z}_j = \mathbf{Q}_i, \mathbf{Z}_j^\top \mathbf{Z}_j = \mathbf{I}_k\}$ (e.g. Edelman et al., 1998).

The *geodesics* (i.e. shortest paths) between two elements in $\mathcal{G}_{k,D}$ are available in closed-form and can be characterized in terms of the *principal angles* $\boldsymbol{\sigma} \in [0, \frac{\pi}{2}]^k$, i.e. the ‘‘amount of rotation’’ required to transition from one space to another. The principal angles between two matrices in $\mathbb{O}^{D \times k}$ can be obtained via an SVD of their product, satisfying the invariance to changes within $\mathcal{S}_i^\mathbb{O}$ (see Appendix A.3). Based on this, there are several *Grassmannian metrics* capturing different notions of distance between subspaces (e.g. Edelman et al., 1998; Qiu et al., 2005). In the following, we highlight a few. We abbreviate $\text{dist}_*(\mathfrak{g}_i, \mathfrak{g}_j) = f(\boldsymbol{\sigma}_{i \rightarrow j})$ as $\text{dist}_* = f(\boldsymbol{\sigma})$, where $*$ here parametrizes any unitarily invariant norm:

- a) **Geodesic distance:** This is the *arc length* of the geodesic between the respective spaces in \mathcal{G} , defined as $\text{dist}_g = \|\boldsymbol{\sigma}\|_2 \in [0, \frac{\pi}{2} \sqrt{k}]$.
- b) **Chordal norm:** $\text{dist}_{c,*}$, obtained by minimizing $\|\mathbf{Q}_i \mathbf{Z}_i - \mathbf{Q}_j \mathbf{Z}_j\|_*$ over orthogonal matrices $(\mathbf{Z}_1, \mathbf{Z}_2)$ (for that reason it is also called *Hausdorff distance*). The ℓ_2 and ℓ_F norms admit a closed-form solution in terms of principal angles: $\text{dist}_{c,2} = \|2 \sin(\frac{1}{2} \boldsymbol{\sigma})\|_\infty \in [0, \sqrt{2}]$ and $\text{dist}_{c,F} = \|2 \sin(\frac{1}{2} \boldsymbol{\sigma})\|_2 \in [0, \sqrt{2k}]$.
- c) **Projection norm:** Also called the *gap metric*, it uses the unique orthogonal projector representation of a given subspace, i.e. $\boldsymbol{\Psi}_i = \mathbf{Q}_i \mathbf{Q}_i^\top$, as follows: $\text{dist}_{p,*} = \|\boldsymbol{\Psi}_i - \boldsymbol{\Psi}_j\|_*$. Here, we also have closed-form expressions for the ℓ_2 and ℓ_F norms: $\text{dist}_{p,2} = \|\sin(\boldsymbol{\sigma})\|_\infty \in [0, 1]$ and $\text{dist}_{p,F} = \|\sin(\boldsymbol{\sigma})\|_2 \in [0, \sqrt{k}]$.
- d) **Fubini-Study:** This quantity is a measure of the *acute angle* between both spaces, generalized to higher dimensions: $\text{dist}_a = \arccos(|\det(\mathbf{Q}_i^\top \mathbf{Q}_j)|) = \arccos(\prod_i \cos(\sigma_i)) \in [0, \frac{\pi}{2}]$.
- e) **Overlap:** The *overlap* $= \frac{1}{k} \|\boldsymbol{\Psi}_i \mathbf{Q}_j\|_F^2 \in [0, 1]$ quantity was used in Gur-Ari et al. (2018) to measure subspace *similarity*. It is not a metric *per se*, since it is highest for equivalent subspaces and decreases with their distance, but it is a bijection of $\text{dist}_{p,F}$, as follows: $\frac{1}{k} \|\boldsymbol{\Psi}_i \mathbf{Q}_j\|_F^2 = \frac{1}{k} \|\mathbf{Q}_i^\top \mathbf{Q}_j\|_F^2 = \frac{1}{k} \|\cos(\boldsymbol{\sigma})\|_F^2 = 1 - \|\cos(\boldsymbol{\sigma})\|_F^2 = 1 - \frac{1}{k} \text{dist}_{p,F}^2$.

While the above metrics apply to any pair of matrices from $\mathbb{O}^{D \times k}$, there are also relevant metrics specific to $\mathbb{M}^{D \times k}$ (boolean subset of $\mathbb{O}^{D \times k}$ defined in Section 3), that can be characterized in a similar manner. Consider an arbitrary pair of k -sparse masks $(\mathbf{m}_i, \mathbf{m}_j)$, and their corresponding orthogonal projectors $\boldsymbol{\Phi} := \text{diag}(\mathbf{m})$. Then we have:




Figure 2: Behaviour of different Grassmannian metrics for random pairs of matrices and masks. (Left)

As a function of D , for different sparsity ratios $\rho := k/D$. (Right) As a function of ρ , for different ambient dimensions D . Lines show the median metric for 50 random pairs, the (almost imperceptible) shaded regions span the 5-95 percentiles (see Figure 11 for broader distributions). Each row shows a selected Grassmannian metric (see Section 4.1 for more metrics): $\text{dist}_{\overline{c,2}}$ is a representative example of a *collapsing* metric, being ≈ 0 almost everywhere. The *overlap* metric is non-collapsing, and its expectation equals ρ .

- i) **Intersection over Union (IoU):** Typically used as an evaluation metric, it is defined as the relative number of entries present in *both* masks, i.e. $\text{IoU} := \frac{\mathbf{m}_i \cap \mathbf{m}_j}{\mathbf{m}_i \cup \mathbf{m}_j} \in [0, 1]$. Then we have $\mathbf{m}_i \cap \mathbf{m}_j = \|\Phi_i \Phi_j\|_F^2 = k \cdot \text{overlap}$, and if both masks are k -sparse, we also have $\mathbf{m}_i \cup \mathbf{m}_j = 2k - (\mathbf{m}_i \cap \mathbf{m}_j) = k(2 - \text{overlap})$, yielding the bijection $\text{overlap} = \frac{2 \text{IoU}}{1 + \text{IoU}}$.
- ii) **Hamming distance:** This quantity, defined as the minimum number of bit-flips needed to pass from one mask to another, was used in You et al. (2020) to measure distances between pruning masks. It is in fact a Grassmannian metric: $\text{dist}_{\text{H}} := \|\mathbf{m}_i - \mathbf{m}_j\|_2^2 = \|\Phi_i - \Phi_j\|_F^2 = \text{dist}_{\rho, F}^2 \in [0, k]$, which means that the bijection $\text{overlap} = 1 - \frac{1}{k} \text{dist}_{\text{H}}$ also holds.

4.2 Comparing Grassmannian Metrics

Our aim is now to identify the most useful Grassmannian metric for comparing neural net parameter masks and Hessian eigenspaces. To empirically compare how different Grassmannian metrics change as a function of ambient dimension D (for fixed ρ) and sparsity ratio $\rho = k/D$ (for fixed D), we conduct a synthetic experiment: We randomly draw matrices from $\mathbb{O}^{D \times k}$ and masks from $\mathbb{M}^{D \times k}$ and compute various Grassmannian metrics between them, normalized to be in $[0, 1]$, with higher metric indicating higher similarity. Results are highlighted in Figure 2 (experimental procedure and further experiments in Appendix A.4). This empirical analysis is complemented with a theoretical result showing that the expected *overlap* equals exactly ρ , as seen in Figure 2 (proof in Appendix A.5). The main insights from this analysis (Appendix A.6 extends the discussion) are:

- I. For a fixed sparsity ratio ρ , the expectation of all metrics becomes *predictable* as D increases: Already in the $D \approx 10^3$ regime (much smaller than modern DL scenarios), expectations converge tightly to values that *seem to only depend on ρ* (Figure 2 left). We gathered empirical baseline values in Table 2.
- II. Not all metrics are equally *informative*: Some metrics quickly collapse to 0 as D increases and ρ decreases (see e.g. Figure 2 (top) or Figures 12 and 13). Looking at the metric definitions in Section 4.1, we can see the reason: Collapsing metrics (e.g. $\text{dist}_{\overline{c,2}}$, $\text{dist}_{\overline{p,2}}$ and $\text{dist}_{\overline{a}}$) are strongly influenced by individual angles, whereas non-collapsing metrics (e.g. $\text{dist}_{\overline{q}}$, $\text{dist}_{\overline{c,F}}$, $\text{dist}_{\overline{p,F}}$ and *overlap*) average over all angles.

With these insights, we identify *overlap* as the most suitable metric among the reviewed ones, since: (1) *It provides a stable and efficient baseline:* Among the non-collapsing metrics, it features the smoothest expectation, equaling exactly $\rho = k/D$. This provides us with a stable, simple and theoretically-grounded baseline, establishing a “**chance-level overlap**” to compare against⁴. This is particularly beneficial for deep

⁴In Appendix A.7 we also compare this analytical baseline to the empirical alternative of sampling random masks, showing it yields similar results.

learning, bypassing the need to compute empirical baselines for large D over many ρ . (2) *It is related to other metrics: overlap* can be mapped to other popular metrics such as the Hamming distance, IoU, or $\text{dist}_{p,F}$ via bijections (see Section 4.1). (3) *It has precedence in the literature:* The *overlap* metric has been used to measure eigenspace similarity in recent deep learning works (e.g. Gur-Ari et al., 2018; Dangel et al., 2022).

5 Sketched Hessian Eigendecompositions to Measure Overlap

To compute the *overlap* between parameter magnitude masks and top- k Hessian eigenspaces, we can now simply plug in the permutation from Section 3 into Eq. (2): \mathbf{Q}_1 becomes the permuted mask matrix $\mathbf{I}_{D,k}$, and \mathbf{Q}_2 the correspondingly permuted top eigenspace $\tilde{\mathbf{U}}^{(k)}$, yielding:

$$\text{overlap}(\mathbf{I}_{D,k}, \tilde{\mathbf{U}}^{(k)}) = \frac{1}{k} \|\mathbf{I}_{D,k}^\top \tilde{\mathbf{U}}^{(k)}\|_F^2 = \frac{1}{k} \|\mathbf{V}\|_F^2, \quad (3)$$

i.e. our problem reduces to computing the Frobenius norm of \mathbf{V} , a $k \times k$ submatrix of $\tilde{\mathbf{U}}^{(k)}$ (Eq. (1)).

A case for eigendecompositions: We aim to estimate $\|\mathbf{V}\|_F^2$, where \mathbf{V} is part of $\tilde{\mathbf{U}}^{(k)}$, a $D \times k$ orthogonal matrix spanning the top eigenspace. Since all standard algorithms to find such matrices (Gram–Schmidt, Householder transformation, Givens rotation) require to keep track of at least k full vectors (Golub & Van Loan, 2013), a *memory requirement of $\mathcal{O}(kD)$ seems inescapable in general*. Still, one may recognize that any orthogonal matrix $\tilde{\mathbf{Q}}$ with the same span as $\tilde{\mathbf{U}}^{(k)}$ yields the same Frobenius norm (due to unitary invariance). However, we need to ensure that the span of said $\tilde{\mathbf{Q}}$ does not overlap with the bulk eigenspace: It needs to be exclusively associated with the top- k eigenvalues. As a consequence, *knowledge of the eigenvalues also seems inescapable in general*. In conclusion, to measure the top- k *overlap*, we argue that a rank- k eigendecomposition is generally needed.

A case for sketched decompositions: Among the methods that satisfy the memory requirement of $\mathcal{O}(kD)$, **orthogonal iterations** (Golub & Van Loan, 2013), an extension of Rayleigh’s power method, is a popular and effective one, requiring $\mathcal{O}(k)$ measurements per iteration and τ iterations, to converge at a rate proportional to $|\lambda_{k+1}/\lambda_k|^\tau$ (Golub & Van Loan, 2013, Th. 7.3.1). This leads to large τ for smoothly decaying spectra, likely scenario for \mathbf{H} , thus $\mathcal{O}(\tau k)$ Hessian-Vector Products (HVPs) can become infeasible. Ritz accelerations provide better convergence but still suffer from this issue (Golub & Van Loan, 2013, 10.1). Alternatively, **Lanczos iterations** converge faster (Saad, 1980), (Golub & Van Loan, 2013, 10.1.6), and can be used to approximate eigenpairs via Ritz pairs (Golub & Van Loan, 2013, 8.1.4) as well as spectral densities of e.g. large-scale deep learning Hessians via Stochastic Lanczos Quadrature (Pappayan, 2018). The main issue here is that measurements must be done in sequential form, which can get very slow for \mathbf{H} , and they involve matrix powers, numerically unstable for rank-deficient matrices like \mathbf{H} . **Sketched methods** (e.g. Halko et al., 2011), based on random measurements, not only satisfy $\mathcal{O}(kD)$ memory, they also exhibit *good convergence for $\mathcal{O}(k)$ measurements* and are *parallelizable and numerically stable* (Halko et al., 2011, 1.4.2, 4.2, and 6.2). This, combined with the ability to perform matrix-free random measurements (Tropp et al., 2019, 3.2), makes it possible to compute Hessian eigendecompositions at unprecedented scales, since the main bottleneck is now the $\mathcal{O}(kD)$ memory requirement (Tropp et al., 2019, 7.1) (for example, storing 1000 eigenpairs for a ResNet-18 (He et al., 2016) with 13M float parameters requires roughly 50GB of memory, same as 100 eigenpairs for a $10\times$ larger model). Last but not least, affordable *a posteriori* methods allow to measure approximation error and rank of the recovered decomposition (Halko et al., 2011; Tropp et al., 2019).

The core idea behind the sketched SVD is that, given a linear operator $\mathbf{A} \in \mathbb{C}^{D_L \times D_R}$ of numerical rank k , any orthogonal matrices $\mathbf{P} \in \mathbb{C}^{D_L \times n_o}$, $\mathbf{Q} \in \mathbb{C}^{D_R \times n_o}$ that approximately capture the column and row space of \mathbf{A} , respectively, satisfy $\mathbf{A} \approx \mathbf{P}\mathbf{P}^* \mathbf{A}\mathbf{Q}\mathbf{Q}^*$ and can be efficiently obtained from $n_o = \mathcal{O}(k)$ random measurements followed by QR orthogonalization (Halko et al., 2011). Typically, $n_o \ll \min(D_L, D_R)$, which results in a *core matrix* $\mathbf{C} := \mathbf{P}^* \mathbf{A}\mathbf{Q} \in \mathbb{C}^{n_o \times n_o}$ that is much smaller⁵ than \mathbf{A} and can be decomposed via classical methods. In summary, a *potentially large and matrix-free linear operator \mathbf{A} is decomposed into two thin matrices (\mathbf{P}, \mathbf{Q}) and a small matrix \mathbf{C}* . A further development features an oversampled *inner matrix* $\mathbf{M}_I := \mathbf{\Upsilon}_I^* \mathbf{A}\mathbf{\Omega}_I \in \mathbb{C}^{n_i \times n_i}$ for random measurements $\mathbf{\Upsilon}_I, \mathbf{\Omega}_I$ of columns each (Boutsidis et al., 2016), where n_i is typically slightly larger

⁵We use lowercase n, k to indicate smaller dimension than D .

than n_o . The SSVD from (Tropp et al., 2019, Sec. 2), gathered in Alg. 1, follows this idea:

$$A \approx P(P^*AQ)Q^* \approx P \underbrace{(\Upsilon_I^*P)^\dagger \Upsilon_I^* A \Omega_I [(\Omega_I^*Q)^\dagger]^*}_{C=U\Sigma V^* \text{ (svd)}} Q^* = (PU)\Sigma(V^*Q^*), \quad (4)$$

This yields an SVD, since PU and QV are orthogonal. It requires n_i+2n_o measurements—see lines 1-3 in Alg. 1: outer measurements are used to produce P and Q , and inner measurements to produce C —, followed by thin matrix operations only. The memory cost is dominated by storing P and Q , i.e. $\mathcal{O}(n_o(D_L+D_R))$. Arithmetic is dominated by the QR orthogonalizations needed to obtain P and Q , as well as the two least-squares problems needed to solve the pseudoinverses (see (Tropp et al., 2019)). Crucially, this approximation only requires a single pass over A , yields tight bounds (Tropp et al., 2019, Th. 5.1) and leads to superior performance (Tropp et al., 2019, Sec. 7), due to its numerical stability, oversampled M_I , as well as uncorrelated measurements between P , Q and M_I (Halko et al., 2011, 5.5) (Tropp et al., 2019, 2.8.1).

Algorithm 1: SSVD (from Tropp et al. (2019))

```

Input:  $n_i \in \mathbb{N}$  // No. of inner measurements
Input:  $n_o \leq n_i$  // No. of outer measurements
Input:  $A \in \mathbb{C}^{D_L \times D_R}$  // Linear operator
Input:  $\Upsilon_I \in \mathbb{C}^{D_L \times n_i}$  // Left inner random matrix
Input:  $\Upsilon_O \in \mathbb{C}^{D_L \times n_o}$  // Left outer random matrix
Input:  $\Omega_I \in \mathbb{C}^{D_R \times n_i}$  // Right inner random matrix
Input:  $\Omega_O \in \mathbb{C}^{D_R \times n_o}$  // Right outer random matrix
//  $P \in \mathbb{C}^{D_L \times n_o}$ ,  $Q \in \mathbb{C}^{D_R \times n_o}$ ,  $U \in \mathbb{C}^{n_o \times n_o}$ 
Output:  $(P, U, \Sigma, V^*, Q^*)$  with  $PU\Sigma V^* Q^* \approx A$ 
// Perform outer measurements
1  $M_{LO}^* \leftarrow \Upsilon_O^* A$  //  $\mathbb{C}^{n_o \times D_L}$ 
2  $M_{RO} \leftarrow A \Omega_O$  //  $\mathbb{C}^{D_R \times n_o}$ 
// Perform inner measurements
3  $M_I \leftarrow \Upsilon_I^* A \Omega_I$  //  $\mathbb{C}^{n_i \times n_i}$ 
// Orthogonalize outer measurements
4  $(Q, \_) \leftarrow \text{qr}(M_{LO})$  //  $P$  has orthonormal columns
5  $(P, \_) \leftarrow \text{qr}(M_{RO})$  //  $Q$  has orthonormal columns
// Solve core matrix via least squares and SVD
6  $C \leftarrow (\Upsilon_I^* P)^\dagger M_I$ 
7  $C \leftarrow [(\Omega_I^* Q)^\dagger C]^*$  //  $\mathbb{C}^{n_o \times n_o}$ 
8  $(U, \Sigma, V^*) \leftarrow \text{svd}(C)$  //  $\Sigma$  are the singular values
9 return  $(P, U, \Sigma, V^*, Q^*)$ 

```

Algorithm 2: SEIGH

```

Input:  $n_i \in \mathbb{N}$  // No. of inner measurements
Input:  $n_o \leq n_i$  // No. of outer measurements
Input:  $A \in \mathbb{C}^{D \times D}$  // Hermitian linear operator
Input:  $\Upsilon \in \mathbb{C}^{D \times n_i}$  // Left inner random matrix
Input:  $\Omega_I \in \mathbb{C}^{D \times (n_i - n_o)}$  // Right inner random matrix
Input:  $\Omega_O \in \mathbb{C}^{D \times n_o}$  // Right outer random matrix
//  $Q \in \mathbb{C}^{D \times n_o}$ ,  $U \in \mathbb{C}^{n_o \times n_o}$ 
Output:  $(Q, U, \Lambda)$  with  $QU\Lambda U^* Q^* \approx A$ 
// Perform measurements.  $[X, Y]$  means matrix concatenation.
// Note the recycled  $\Omega_O$  measurements.
1  $M_O \leftarrow A \Omega_O$  //  $\mathbb{C}^{D \times n_o}$ 
2  $M_I \leftarrow \Upsilon^* [A \Omega_I, M_O]$  //  $M_I = \Upsilon^* A [\Omega_I, \Omega_O] \in \mathbb{C}^{n_i \times n_i}$ 
// Orthogonalize outer measurements
3  $(Q, \_) \leftarrow \text{qr}(M_O)$  //  $Q$  has orthonormal columns
// Solve core matrix via least squares and EIGH
4  $(\bar{U}, \bar{\Lambda}, \bar{V}^*) \leftarrow \text{svd}(M_I)$ 
5  $C_L \leftarrow (\Upsilon^* Q)^\dagger \bar{U}$ 
6  $C_R \leftarrow (\Omega^* Q)^\dagger \bar{V}$ 
7  $C \leftarrow C_L \bar{\Lambda} C_R^*$  //  $C \in \mathbb{C}^{n_o \times n_o}$  is Hermitian
8  $(U, \Lambda) \leftarrow \text{eigh}(C)$  //  $\Lambda$  are the eigenvalues
9 return  $(Q, U, \Lambda)$ 

```

Leveraging Symmetry: While previous works studied Hermitian extensions (Halko et al., 2011; Clarkson & Woodruff, 2017; Tropp et al., 2017), none of them proposes a single-pass sketched eigendecomposition via oversampled and uncorrelated inner matrix. In this work, we propose SEIGH (Alg. 2), a variant of SSVD that retains core oversampling while leveraging conjugate symmetry to drastically reduce the required measurements and memory. Without loss of generality (Halko et al., 2011), we assume $Q=P$, so only Q needs to be computed via $A\Omega$ outer measurements. Then, C is also Hermitian and can be eigendecomposed:

$$A \approx Q(Q^*AQ)Q^* \approx Q \underbrace{(\Upsilon_I^*Q)^\dagger \Upsilon_I^* A \Omega_I [(\Omega_I^*Q)^\dagger]^*}_{M_I = \bar{U}\bar{\Lambda}\bar{V}^*} Q^* = Q \underbrace{(\Upsilon_I^*Q)^\dagger \bar{U} \bar{\Lambda} \bar{V}^* [(\Omega_I^*Q)^\dagger]^*}_{C=C^*=U\Lambda U^* \text{ (eigh)}} Q^*$$

Setting $\Upsilon_I = \Omega_I$ would further simplify the structure, saving one pseudoinverse, but there is one caveat: We also wish to recycle part of the $A\Omega_I$ inner measurements to obtain Q , since this is a major bottleneck for the Hessian. But recycling measurements and setting $\Upsilon_I = \Omega_I$ would lead to correlated measurements between Q and M_I which, as mentioned, has negative impact in the quality of this procedure (Tropp et al., 2019, 2.8.1). We opt to recycle the measurements, while keeping the uncorrelated $\Upsilon_I = \Omega_I$ and the two pseudoinverses.

This way, only n_i measurements are needed, memory is roughly halved, and arithmetic is reduced by one QR decomposition and increased by one inner SVD. With SEIGH, we are now able to approximate $\|\mathbf{V}\|_F^2$ for large Hessians, leading to the *sketched overlap* approximation.

6 Magnitude Masks and Hessian Eigenspaces Overlap Substantially

We now investigate the similarity between the space spanned by the parameter magnitude masks and the top Hessian eigenspaces over the course of neural network training. We first study a small-scale toy problem (Section 6.1), where all involved quantities can still be computed exactly. This allows us to verify the existence of early crystallization, as well as the quality of our sketched *overlap* approximation. We observe an *overlap* consistently above random chance. We then move onto larger problems (up to $D > 11\text{M}$ and $k = 1500$), where we also confirm that the similarity between parameter magnitude masks and top Hessian eigenspaces is orders of magnitude above random chance (Section 6.2), and that this effect is accentuated with network size. We report *overlap* results up to step 8000: since larger problems take longer to train, this covers different stages of training, but all of them reach well past initialization (see Table 1 for detailed accuracy at terminal steps). This provides us with a sufficient and convenient range to compare the observed high *overlap* across all problems.

6.1 Exact Computations on 16×16 MNIST

Table 1: Overview of experimental settings, detailing number of model parameters (D), learning rate (η), batch size (B), steps per epoch (T), test accuracy (acc) at step t , number of train/test samples used to compute H_{train}/H_{test} (N_{train}/N_{test} respectively), and number of SEIGH outer measurements (n_o , see Alg. 2)

Problem	Model	D	η	B	T	acc	N_{train}/N_{test}	n_o
16×16 MNIST	tanh-MLP Martens & Grosse (2015)	7030	0.3	500	100	95.78% ($t = 1000$)	500/500	355
CIFAR-10	3c3D-CNN Schneider et al. (2019)	895,210	0.0226	128	312	74.52% ($t = 8000$)	500/500	1000
CIFAR-100	ALL-CNN-C Springenberg et al. (2015)	1,387,108	0.1658	256	156	40.50% ($t = 8000$)	n.a./1000	1000
ImageNet	RESNET-18 He et al. (2016)	11,689,512	0.1	150	8207	17.33% ($t = 8000$)	n.a./5000	1500

For our toy problem, we seek a setup that is overparametrized enough to exhibit mask crystallization, but small enough so that full Hessian eigendecompositions can be computed. This is achieved by the setup from Martens & Grosse (2015), which is able to achieve zero training loss on downsampled MNIST digit classification with only 7030 parameters. We train using SGD as detailed in Table 1, reaching a test accuracy of 95.78% after 10 epochs (Figure 15). To compute the training and test Hessians, we use a fixed set of 500 random samples (50 per class) from each respective data split. At every fifth training step t_i , we compute the sparsity ratio κ (defined in Section 2.2) for the parameters θ and the Hessian spectrum Λ , observing that *both subspaces experience early collapse* (Figure 15). Then, for pairs of steps (t_i, t_j) , we compute the IoU between successive parameter magnitude masks, as well as the overlap between successive Hessian eigenspaces. We observe that *both subspaces also experience early crystallization* (Figure 16). We reiterate that we merely *inspect* the parameter magnitude masks, but we never *apply* them, i.e., we never prune the network. Once we established our desired scenario, we compute several Grassmannian metrics between the spans of parameter magnitude masks and top Hessian eigenspaces (Figure 17). All non-collapsing metrics report that *both subspaces show a substantial and consistent similarity*. In particular, our sketched *overlap* approximation closely tracks the exact *overlap* (Figure 1, left), which is largest shortly after initialization and then decays, but still stabilizes well above random chance.

6.2 Sketched *overlap* on larger problems (CIFAR-10, CIFAR-100 and ImageNet)

We aim now to verify our discovery of substantial overlap in larger setups. Using the DEEPOBS framework (Schneider et al., 2019) and SGD, we train larger models on CIFAR-10/100 Krizhevsky & Hinton (2009) and on IMAGENET Deng et al. (2009) (see Table 1 for details). At steps $\{2000, 4000, 8000\}$, we gather the




Figure 3: Overlap between parameter magnitude masks and top Hessian eigenspaces as a function of k at three different points during the training process. Shown is the sketched *overlap* between the mask of the largest k parameters by magnitude and the top- k Hessian eigenspace. Consistently across the training process and for all problems and values of k , we observe an *overlap* substantially larger than the random chance baseline (--) introduced in Section 4.2.

parameters and compute the top Hessian eigenspaces using SEIGH, which allows us to compute the sketched *overlap* for every $k \in \{1, \dots, n_o\}$. To compute the Hessians, we choose N_{train}/N_{test} to have several (balanced) samples per class, and n_o to be substantially larger than the number of classes, which has been shown to be linked to the top eigenspace dimensionality (e.g. Gur-Ari et al., 2018; Pappayan, 2020). We elaborate on these choices and report runtimes in Appendix B.2. Note that training CIFAR-100 and IMAGENET relies on noisy data augmentations, which lead to nondeterministic behaviour in \mathbf{H}_{train} . For this reason, we compute their *overlap* for \mathbf{H}_{test} only. This choice is supported by the CIFAR-10 results, showing minimal difference in *overlap* for either \mathbf{H}_{train} or \mathbf{H}_{test} .

We observe that the measured *overlap* is *significant for all problems, steps and dimensions* (Figure 3 and Figure 4, left), confirming the findings from our 16×16 MNIST toy problem. Furthermore, *this effect is accentuated by model size, surpassing a factor of 10^3 times the baseline for the RESNET-18 on IMAGENET* (Figure 4, right). The observed sketched *overlap* also *raises in the very early steps*, perhaps linked to the crystallization covered in Section 2. We also note that, while the precision of sketched decompositions is reported to decay sharply for eigenpairs close to n_o (Tropp et al., 2019, 7.9), this does not affect our observation, since the obtained *overlap* is *particularly high for the lower regimes of k* , far from n_o and where the method is most reliable.

7 Conclusion

We started with the observation that, at the early stages of neural network training, both parameter magnitude masks and loss Hessian eigenspaces collapse and crystallize. To connect these two phenomena throughout the training process, we proposed a principled methodology that compares these two fundamental deep learning quantities using Grassmannian metrics, and identified *overlap* as a particularly advantageous one. We further developed SEIGH, a matrix-free sketched eigendecomposition that works for Hessians of over 10^7 parameters, allowing us to approximate *overlap*. Our experiments reveal an *overlap* between parameter magnitude masks and top Hessian eigenspaces well above chance level for all observed problems, training steps and dimensionalities, indicating that, in DL, *top Hessian eigenvectors tend to be concentrated around larger parameters*, or equivalently, that *large parameters tend to align with directions of large loss curvature*.

While the obtained *overlap* may be considered small in absolute terms, the fact that it is orders of magnitude above chance could play a crucial role at large scales, analogously to how mini-batch gradients produce

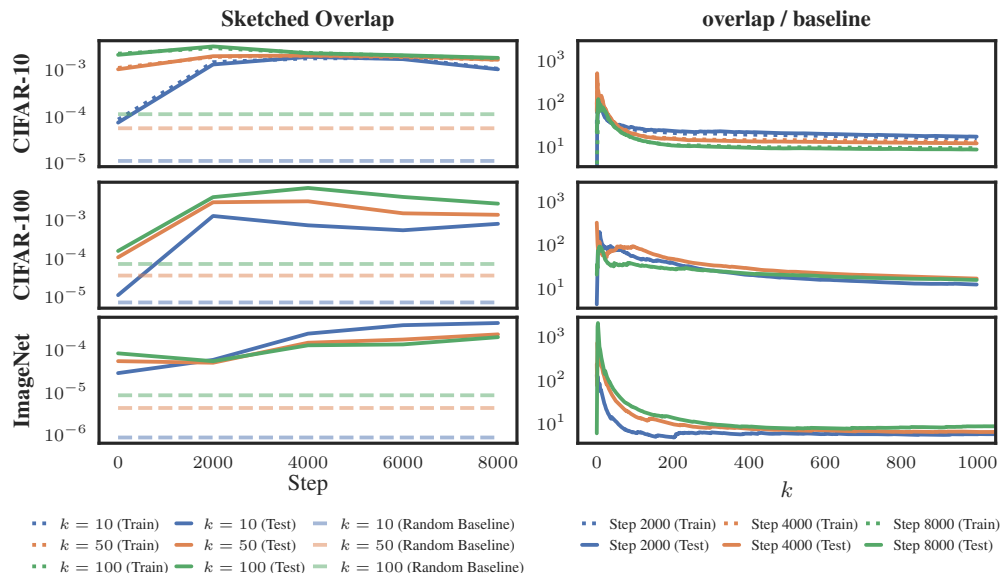


Figure 4: (Left) **Overlap between parameter magnitude masks and top Hessian eigenspaces as a function of training step.** Shown is the sketched *overlap* between masks of k -largest parameters by magnitude and top- k Hessian eigenspaces, for different values of k at different training steps. The observed *overlap* raises early on, and is consistently and substantially above random chance baseline (---, introduced in Section 4.2). (Right) **Factor by which the measured sketched *overlap* is greater than the random baseline.** For most of the measured training process and choices of k , the observed *overlap* is at least $10\times$ larger than random chance baseline. This multiple factor over the random baseline is largest for small k and increases with network size, surpassing 10^3 .

very noisy fluctuations in the short term, but lead to qualitative performance and generalization over larger training spans. The relevance of this connection could also bear potential for downstream tasks where the Hessian plays a prominent role (e.g. by approximating expensive Hessian quantities via cheap parameter inspection).

One main limitation of this method is that it still does not scale up to contemporary network sizes (see Appendix B.2). In this sense, we emphasize that our methodology is primarily meant as an *analysis* technique to understand and leverage the connection between parameters and loss curvature, and not necessarily part of the downstream applications themselves. In Appendix A.2, we have also explored alternative approaches that show potential as more scalable alternatives to computing *overlap*.

Acknowledgments

The authors thank Joel A. Tropp for background on the distance between random subspaces. AF was supported by the DFG SPP 2298 (Project HE 7114/5-1). FS was supported by funds from the Cyber Valley Research Fund of the State of Baden-Württemberg. Further, the authors gratefully acknowledge financial support by the European Research Council through ERC CoG Action 101123955 ANUBIS; the DFG Cluster of Excellence "Machine Learning - New Perspectives for Science", EXC 2064/1, project number 390727645; the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ: 01IS18039A); and the Carl Zeiss Foundation, (project "Certification and Foundations of Safe Machine Learning Systems in Healthcare"), as well as funds from the Ministry of Science, Research and Arts of the State of Baden-Württemberg.

References

- P. A. Absil, Robert Mahony, and Rodolphe Sepulchre. Riemannian Geometry of Grassmann Manifolds with a View on Algorithmic Computation. *Acta Applicandae Mathematicae*, 2004.
- Alessandro Achille, Matteo Rovere, and Stefano Soatto. Critical Learning Periods in Deep Networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- Thomas Bendokat, Ralf Zimmermann, and Pierre-Antoine Absil. A Grassmann Manifold Handbook: Basic Geometry and Computational Aspects, 2020.
- Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Gutttag. What is the State of Neural Network Pruning? In *Proceedings of Machine Learning and Systems (MLSys)*, 2020.
- Christos Boutsidis, David P. Woodruff, and Peilin Zhong. Optimal Principal Component Analysis in Distributed and Streaming Models. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, 2016.
- Kenneth L. Clarkson and David P. Woodruff. Low-rank psd approximation in input-sparsity time. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2017.
- Felix Dangel, Frederik Kunstner, and Philipp Hennig. BackPACK: Packing more into Backprop. In *International Conference on Learning Representations (ICLR)*, 2020.
- Felix Dangel, Lukas Tatzel, and Philipp Hennig. ViViT: Curvature Access Through The Generalized Gauss-Newton’s Low-Rank Structure. *Transactions on Machine Learning Research (TMLR)*, 2022.
- Felix Dangel, Runa Eschenhagen, Weronika Ormaniec, Andres Fernandez, Lukas Tatzel, and Agustinus Kristiadi. Position: Curvature matrices should be democratized via linear operators, 2025.
- Yann N. Dauphin, Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, Surya Ganguli, and Yoshua Bengio. Identifying and Attacking the Saddle Point Problem in High-Dimensional Non-Convex Optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- Persi Diaconis and Mehrdad Shahshahani. The Subgroup Algorithm for Generating Uniform Random Variables. *Probability in the Engineering and Informational Sciences*, 1987.
- Alan Edelman, Tomás A. Arias, and Steven T. Smith. The Geometry of Algorithms with Orthogonality Constraints. *SIAM Journal on Matrix Analysis and Applications*, 1998.
- Mohamed Elsayed and A. Rupam Mahmood. HesScale: Scalable Computation of Hessian Diagonals, 2022.
- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Machine Learning (ICML)*, 2021.
- Jonathan Frankle and Michael Carbin. The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin. Linear Mode Connectivity and the Lottery Ticket Hypothesis. In *International Conference on Machine Learning (ICML)*, 2020.
- Trevor Gale, Erich Elsen, and Sara Hooker. The State of Sparsity in Deep Neural Networks, 2019.
- Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An Investigation into Neural Net Optimization via Hessian Eigenvalue Density. In *International Conference on Machine Learning (ICML)*, 2019.
- Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, 2013.
- Guy Gur-Ari, Daniel A. Roberts, and Ethan Dyer. Gradient Descent Happens in a Tiny Subspace, 2018.

- N. Halko, P. G. Martinsson, and J. A. Tropp. Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions. *SIAM Review*, 2011.
- Song Han, Jeff Pool, John Tran, and William Dally. Learning both Weights and Connections for Efficient Neural Network. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Torstén Høefer, Dan Alistarh, Tal Ben-Nun, Nikoli Dryden, and Alexandra Peste. Sparsity in Deep Learning: Pruning and growth for efficient inference and training in neural networks. *Journal of Machine Learning Research*, 2021.
- Niall Hurley and Scott Rickard. Comparing Measures of Sparsity. *IEEE Transactions on Information Theory*, 2009.
- Agustinus Kristiadi, Matthias Hein, and Philipp Hennig. Learnable uncertainty under Laplace approximations. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2021.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report, 2009.
- Yann LeCun, John Denker, and Sara Solla. Optimal Brain Damage. In *Advances in Neural Information Processing Systems (NeurIPS)*, 1989.
- Chunyuan Li, Heerad Farkhoor, Rosanne Liu, and Jason Yosinski. Measuring the Intrinsic Dimension of Objective Landscapes. In *International Conference on Learning Representations (ICLR)*, 2018a.
- Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018b.
- James Martens. Deep learning via Hessian-free optimization. In *International Conference on Machine Learning (ICML)*, 2010.
- James Martens. *Second-order Optimization for Neural Networks*. PhD dissertation, University of Toronto, 2016.
- James Martens and Roger Grosse. Optimizing Neural Networks with Kronecker-Factored Approximate Curvature. In *International Conference on Machine Learning (ICML)*, 2015.
- Yuri Neretin. On Jordan Angles and the Triangle Inequality in Grassmann Manifolds. *Geometriae Dedicata*, 2001.
- Vardan Papyan. The Full Spectrum of Deepnet Hessians at Scale: Dynamics with SGD Training and Sample Size. *arXiv: Learning*, 2018.
- Vardan Papyan. Measurements of Three-Level Hierarchical Structure in the Outliers in the Spectrum of Deepnet Hessians. In *International Conference on Machine Learning (ICML)*, 2019.
- Vardan Papyan. Traces of Class/Cross-Class Structure Pervade Deep Learning Spectra. *Journal of Machine Learning Research*, 2020.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library, 2019.
- Barak A. Pearlmutter. Fast Exact Multiplication by the Hessian. *Neural Computation*, 1994.
- Li Qiu, Y. X. Zhang, and Chi-Kwong Li. Unitarily Invariant Metrics on the Grassmann Space. *SIAM Journal on Matrix Analysis and Applications*, 2005.

- Yousef Saad. On the Rates of Convergence of the Lanczos and the Block-Lanczos Methods. *SIAM Journal on Numerical Analysis*, 1980.
- Levent Sagun, Utku Evci, V. Ugur Guney, Yann Dauphin, and Leon Bottou. Empirical Analysis of the Hessian of Over-Parametrized Neural Networks, 2018. ICLR Workshop.
- Adepu Ravi Sankar, Yash Khasbage, Rahul Vigneswaran, and Vineeth N Balasubramanian. A Deeper Look at the Hessian Eigenspectrum of Deep Neural Networks and its Applications to Regularization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- Frank Schneider, Lukas Balles, and Philipp Hennig. DeepOBS: A Deep Learning Optimizer Benchmark Suite. In *International Conference on Learning Representations (ICLR)*, 2019.
- Nicol N Schraudolph. Fast curvature matrix-vector products for second-order gradient descent. *Neural Computation*, 2002.
- Jost T. Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for simplicity: The all convolutional net. In *ICLR Workshop*, 2015.
- Joel A. Tropp, Alp Yurtsever, Madeleine Udell, and Volkan Cevher. Fixed-rank approximation of a positive-semidefinite matrix from streaming data. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- Joel A. Tropp, Alp Yurtsever, Madeleine Udell, and Volkan Cevher. Streaming low-rank matrix approximation with an application to scientific simulation. *SIAM Journal on Scientific Computing*, 2019.
- Edward Witten. Quantum field theory, Grassmannians, and algebraic curves. *Communications in Mathematical Physics*, 1988.
- Zhewei Yao, Amir Gholami, Kurt Keutzer, and Michael W. Mahoney. PyHessian: Neural Networks Through the Lens of the Hessian. In *IEEE International Conference on Big Data (Big Data)*, 2020.
- Ke Ye and Lek-Heng Lim. Schubert Varieties and Distances between Subspaces of Different Dimensions. *SIAM Journal on Matrix Analysis and Applications*, 2016.
- Haoran You, Chaojian Li, Pengfei Xu, Yonggan Fu, Yue Wang, Xiaohan Chen, Richard G. Baraniuk, Zhangyang Wang, and Yingyan Lin. Drawing Early-Bird Tickets: Toward More Efficient Training of Deep Networks. In *International Conference on Learning Representations (ICLR)*, 2020.
- Jiayao Zhang, Guangxu Zhu, Robert W. Heath, and Kaibin Huang. Grassmannian Learning: Embedding Geometry Awareness in Shallow and Deep Learning, 2018.

A Supplementary Material for Grassmannian Metrics

A.1 Perturbing the Parameters to Measure their Importance

In order to prove that a certain subset of parameters is “sensitive” without using the theoretical framework we propose, one may consider a perturbation study where large parameters are perturbed, and the resulting loss variation is measured. More variation would then relate to more sensitivity. Here we show that, in general, such an experiment does not yield an informative quantity towards this goal (unlike our proposed Grassmannians), due to the non-PSD nature of \mathbf{H} , (i.e. negative and positive directions of curvature may cancel out for arbitrary sets of parameters) and the lack of details about the loss landscape regularity.

At any parameter vector $\boldsymbol{\theta} \in \mathbb{R}^D$, and under additive perturbations $\boldsymbol{\delta} \in \mathbb{R}^D$, the loss can be expressed via its Taylor expansion (for gradient \mathbf{g} , Hessian \mathbf{H} and Lagrange remainder R_3):

$$L(\boldsymbol{\theta} + \boldsymbol{\delta}) = L(\boldsymbol{\theta}) + \mathbf{g}^\top \boldsymbol{\delta} + \frac{1}{2} \boldsymbol{\delta}^\top \mathbf{H} \boldsymbol{\delta} + R_3(\boldsymbol{\delta}) \quad (5)$$

Then, given a parameter mask $\mathbf{m} \in \mathbb{B}^D$, the proposed perturbation study would consist in drawing random perturbations $\boldsymbol{\delta}$ from any i.i.d. distribution \mathcal{N} with zero mean $\boldsymbol{\mu}$ and unit variance, and then mask said perturbations via $\boldsymbol{\delta}_m = (\mathbf{m} \circ \boldsymbol{\delta})$ and add them to $\boldsymbol{\theta}$ in order to estimate the expectation:

$$\mathbb{E}_{\mathcal{N}} [L(\boldsymbol{\theta} + \boldsymbol{\delta}_m)] = L(\boldsymbol{\theta}) + \underbrace{\mathbf{g}^\top \mathbb{E}_{\mathcal{N}} [\boldsymbol{\delta}_m] + \frac{1}{2} \langle \mathbf{H}, \mathbb{E}_{\mathcal{N}} [\boldsymbol{\delta}_m \boldsymbol{\delta}_m^\top] \rangle + \mathbb{E}_{\mathcal{N}} [R_3(\boldsymbol{\delta}_m)]}_{\varepsilon}$$

This way, we are perturbing only the selected parameters in \mathbf{m} , and measuring their impact on the (known) loss $L(\boldsymbol{\theta})$: If $|\varepsilon|$ is larger, this can be interpreted as the subset ξ being more ‘‘sensitive’’.

Since we assumed $\boldsymbol{\mu}_m := \mathbb{E}_{\mathcal{N}} [\boldsymbol{\delta}_m] = 0$ (otherwise reparametrize $(\boldsymbol{\theta} + \boldsymbol{\delta}_m) =: (\boldsymbol{\theta} + \boldsymbol{\mu}_m) + (\boldsymbol{\delta}_m - \boldsymbol{\mu}_m)$ such that the reparametrized mean of the perturbations will be 0), and we also assumed a (masked) unit covariance $\text{Cov} [\boldsymbol{\delta}_m] = \mathbf{I}_m$, we have:

$$\text{Cov} [\boldsymbol{\delta}_m] := \mathbb{E}_{\mathcal{N}} [(\boldsymbol{\delta}_m - \boldsymbol{\mu}_m)(\boldsymbol{\delta}_m - \boldsymbol{\mu}_m)^\top] = \mathbb{E}_{\mathcal{N}} [\boldsymbol{\delta}_m \boldsymbol{\delta}_m^\top] = \mathbf{I}_m$$

Then, the expectation of the perturbed loss simplifies:

$$\begin{aligned} \mathbb{E}_{\mathcal{N}} [L(\boldsymbol{\theta} + \boldsymbol{\delta}_m)] &= L(\boldsymbol{\theta}) + \underbrace{\mathbf{g}^\top \mathbb{E}_{\mathcal{N}} [\boldsymbol{\delta}_m]}_0 + \frac{1}{2} \langle \mathbf{H}, \underbrace{\mathbb{E}_{\mathcal{N}} [\boldsymbol{\delta}_m \boldsymbol{\delta}_m^\top]}_{\mathbf{I}_m} \rangle + \mathbb{E}_{\mathcal{N}} [R_3(\boldsymbol{\delta}_m)] \\ &= L(\boldsymbol{\theta}) + \underbrace{\frac{1}{2} \text{Tr}_m(\mathbf{H}) + \mathbb{E}_{\mathcal{N}} [R_3(\boldsymbol{\delta}_m)]}_{\varepsilon} \end{aligned}$$

We see now how, in general, this technique cannot be relied upon: since \mathbf{H} is non-PSD in general, any of its subtraces could cancel R_3 , leading to $\varepsilon = 0$. Therefore, this procedure is not guaranteed to be informative in measuring the impact of parameter subsets on the loss curvature. In contrast, our proposed method is guaranteed to yield larger similarities whenever the selected parameters in \mathbf{m} tend to align with directions of larger loss curvature, as shown in the paper.

Alternatively, one could consider to sample perturbations for parameters *one by one*, i.e. by setting \mathbf{m} to equal 1 at exactly m_i , and 0 otherwise. Then, the expectation of the perturbed loss becomes:

$$\mathbb{E}_{\mathcal{N}} [L(\boldsymbol{\theta} + \boldsymbol{\delta}_m)] = L(\boldsymbol{\theta}) + \underbrace{\frac{1}{2} \mathbf{H}_{ii}}_{\varepsilon_i} + \mathbb{E}_{\mathcal{N}} [R_3(\boldsymbol{\delta}_m)] \quad (6)$$

The idea here is that the aggregated perturbation $\varepsilon' = \sum_i |\varepsilon_i|$ is now more likely to be informative, since positive and negative diagonal entries don’t cancel anymore. This still entails a problematic trade-off that favours our Grassmannian characterization, for the following reasons: (1) In DL, the loss landscape may be far from convex or regular (e.g. Li et al., 2018b, Fig. 1), and to interpret ε' meaningfully we would still require some characterization of R_3 similar in spirit to the theoretically founded Grassmannian that we use. (2) This method requires to estimate one expectation per scalar parameter. For a set of k parameters and n samples per expectation, this means performing $\mathcal{O}(n \cdot k)$ forward propagations over a given dataset. In contrast, our method is theoretically grounded and our proposed sketched method requires just $\mathcal{O}(k)$ HVPs (see Section 5 and Halko et al. (2011, 4.2) for more discussion).

A.2 Beyond Perturbing the Parameters to Measure their Importance

Despite the various sources of error and computational issues related to parameter perturbations (see Appendix A.1), one may still feel compelled to explore related strategies, seeking a more efficient and conceptually simple alternative to our proposed Grassmannian methodology. In this section, we introduce three such alternatives which, despite of some issues, constitute a promising line of investigation. We discuss here some of their main aspects in terms of model error, computation and empirical behaviour.

Nonnegative Gauss-Newton Subtraces: A main issue with \mathbf{H} is its non-PSD nature. Instead, the idea here is to target the Gauss-Newton matrix \mathbf{G} , which is a PSD approximation to \mathbf{H} (Schraudolph, 2002).

Since PSD matrices have nonnegative diagonal entries, \mathbf{G}_{ii} could serve as a proxy measurement, denoting parameter sensitivity for θ_i . In that case, a larger trace for a given subset of indices could mean that the corresponding parameter space is associated with larger loss curvature.

Instead of casting this as a parameter perturbation study, which would be exposed to the R_3 approximation error as well as the impossibility of disentangling \mathbf{G} from \mathbf{H} (see [Appendix A.1](#)), we propose here to *measure the diagonal entries of \mathbf{G} one by one using Gauss-Newton Vector Products (GNVPs)*. Since we are just interested in the entries corresponding to the k largest parameters by magnitude, this just requires $\mathcal{O}(D)$ memory and $\mathcal{O}(k)$ GNVPs, thus being fully parallelizable and substantially faster than computing *overlap*. Furthermore, GNVPs are roughly twice as fast as HVPs, and require half the memory ([Martens, 2010, 4.2](#)).

Once we have obtained the top- k diagonal entries (i.e. the diagonal entries of \mathbf{G} corresponding to the k -largest parameters), we need a way to compare them with *overlap*. For this, we propose to compute the following *normalized trace summation*:

$$\xi_k := \frac{1}{\text{Tr}(\mathbf{G})} \sum_{i \in \{i_1 \dots i_k\}} \mathbf{G}_{ii} \in [0, 1] \quad (7)$$

Where the i_1, \dots, i_k indices run over the parameters by descending magnitude. We can see that ξ increases monotonically with k , reaching $\xi_D = 1$. Furthermore, we define the baseline corresponding to a uniformly distributed trace as $\bar{\xi}_k := k/D$. This baseline behaves analogously to our “*chance-level*” *overlap* baseline from [Section 4.2](#): if $\xi_k > \bar{\xi}_k$, the top- k parameters are associated with an above-chance trace (which, as already mentioned, we treat as a proxy for curvature in the case of \mathbf{G}). See [Figures 6 and 7](#) for an illustration.

Nonnegative Squared-Hessian Subtraces: The rationale and computations involved here are identical as for the previously described \mathbf{G} , but using \mathbf{H}^2 instead. This is viable, since \mathbf{H}^2 is also PSD and $[\mathbf{H}^2]_{ii}$ can be efficiently and exactly computed with a single HVP as follows:

$$[\mathbf{H}^2]_{ii} = \mathbf{e}_i^\top \mathbf{H}^2 \mathbf{e}_i = \|\mathbf{H} \mathbf{e}_i\|_2^2$$

This technique is also substantially faster than *overlap*, requiring only $\mathcal{O}(D)$ memory and $\mathcal{O}(k)$ parallelizable HVPs. Although we do not assume $\mathbf{H}^2 \approx \mathbf{H}$, this experiment aims to explore whether subtraces of \mathbf{H}^2 can also be treated as proxy for curvature, in the same fashion as previously discussed for \mathbf{G} .

Adversarial Perturbations: Recall from [Appendix A.1](#) and [Eq. \(6\)](#) that a random perturbation study would require $\mathcal{O}(n \cdot k)$ forward propoagations, for potentially large n . Instead of random sampling, the idea here is to *use a single “optimal” perturbation, such that curvature is maximized in that direction*. This is precisely the idea behind Sharpness-Aware Minimization (SAM) ([Foret et al., 2021](#)), where the perturbation with constrained radius $\lambda \in \mathbb{R}_+$ such that the loss $L(\boldsymbol{\theta} + \lambda \epsilon^*)$ is maximized, has the approximate closed-form solution $\epsilon^* = \mathbf{g}/\|\mathbf{g}\|_2$. In our setting from [Eq. \(6\)](#), we want to evaluate the following *masked loss delta*:

$$\Delta_{\lambda,i} = |L(\boldsymbol{\theta}) - L(\boldsymbol{\theta} + \lambda(\mathbf{e}_i \circ \epsilon^*))|$$

where the i index covers each of the top- k parameters by magnitude. This doesn’t rely on PSD-ness anymore, since Δ comprises absolute values for individual parameters. It is also very fast to compute, since we just need $\mathcal{O}(D)$ memory, one forward- and backpropagation to obtain $L(\boldsymbol{\theta})$ and ϵ , plus one additional forward propagation per i . Once we have the *deltas*, we can define a normalized summation, analogous to [Eq. \(7\)](#), that can be used to compare against $\bar{\xi}_k$ (see [Figures 6 and 7](#)):

$$\zeta_{\lambda,k} := \frac{1}{\sum_{i=1}^D \Delta_{\lambda,i}} \sum_{i \in \{i_1 \dots i_k\}} \Delta_{\lambda,i} \in [0, 1] \quad (8)$$

Experimental results: So far, we introduced three measurement techniques that are more efficient than *overlap*. We now ask: *Can they be used as a reliable proxy for parameter sensitivity? Do they yield measurements above baseline? How do they relate to overlap?* To answer this, we gathered the above metrics for steps 0, 100 and 1000 of the MNIST setup described in Section 6.1. The steps were chosen to cover initialization, mid-training and convergence, as illustrated in Figure 5. Results can be found in Figures 6 and 7, together with the corresponding *overlap* and uniform baseline $\bar{\xi}$ for comparison. We chose 3 different values of λ to cover different perturbation scales.




Figure 5: Training evolution of the model used to report Figures 6, 7 and 14. It corresponds to the 7030-parameter model trained on 16×16 downsampled MNIST, as described in Section 6.1.

Comparing both figures, we first observe that training and test data look quite similar throughout the whole training, which is consistent with the small gap between training and test accuracies from Figure 5. Crucially, we observe that, like *overlap*, all 3 proxy measurements tend to be well above chance-level baseline in most of their domain, supporting their potential use as scalable replacements for *overlap*, and suggesting a possible connection to the Hessian eigenspace. This is promising, but it comes with the caveat of *more unstable behaviour*, both across k —all proxies tend to be below baseline for small k —, and across training step—they can go well above *overlap* early on, but then they steeply collapse towards baseline. The instability of these *overlap* proxies, together with their approximate nature, makes them a challenging but potentially rewarding target for future research.

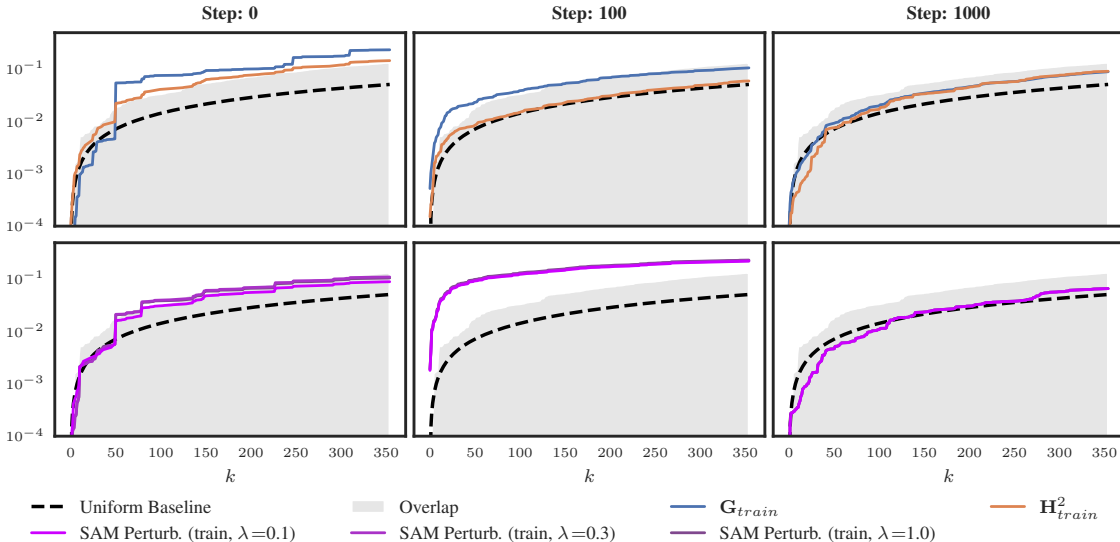


Figure 6: Proxy *overlap* measurements, training split: (top row) The \mathbf{G} normalized trace summation feature from Eq. (7), as well as the corresponding \mathbf{H}^2 normalized trace summation. (bottom row) The SAM perturbation feature from Eq. (8), obtained for three different values of λ , to cover different perturbation scales. (both rows) All features are compared to the uniform baseline k/D and the *overlap* associated to the k -largest parameters by magnitude. See Figure 5 for respective accuracies at given steps.

A.3 Main Properties of Grassmannian Metrics

A desirable property of Grassmann manifolds is the availability of closed-form expressions for their *geodesics* (i.e. the shortest paths between any two elements $\mathbf{g}_i, \mathbf{g}_j \in \mathcal{G}$) and metric functions. We can thus measure distances between subspaces in an interpretable and computationally amenable manner: geodesics from \mathbf{g}_i to

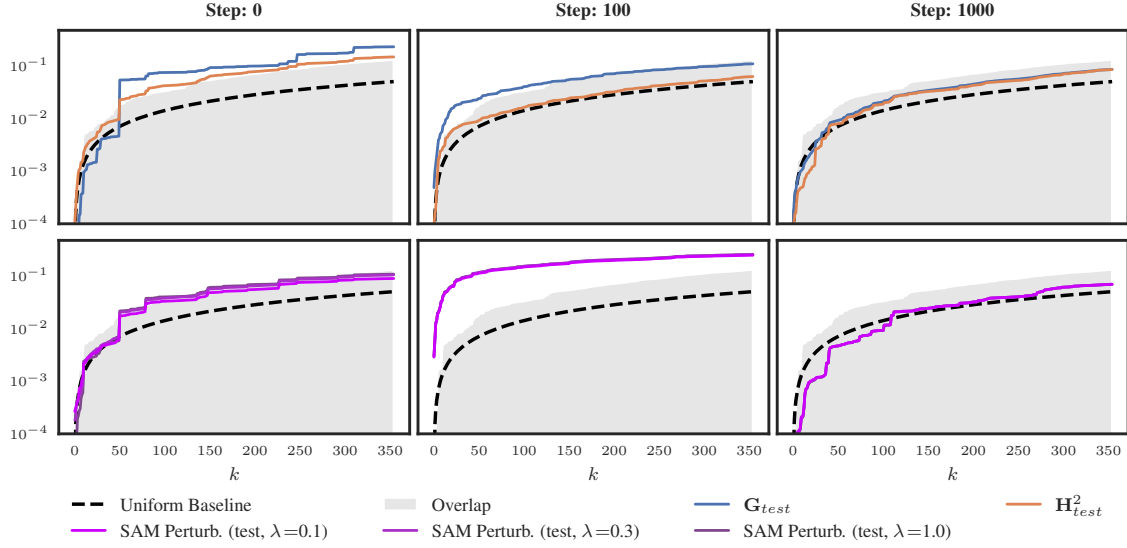


Figure 7: Proxy *overlap* measurements, test split: (*top row*) The \mathbf{G} normalized trace summation feature from Eq. (7), as well as the corresponding \mathbf{H}^2 normalized trace summation. (*bottom row*) The SAM perturbation feature from Eq. (8), obtained for three different values of λ , to cover different perturbation scales. (*both rows*) All features are compared to the uniform baseline k/D and the *overlap* associated to the k -largest parameters by magnitude. See Figure 5 for respective accuracies at given steps.

\mathbf{g}_j follow circular trajectories, and therefore their *distance* can be interpreted as the “amount of rotation” needed to go from one space to another. Such rotations can be succinctly expressed in terms of *principal angles* $\sigma_{i \rightarrow j} \in [0, \frac{\pi}{2}]^k$, and they can be efficiently obtained from $\mathbf{Q} \in \mathbb{O}^{D \times k}$ matrices via their Singular Value Decomposition (SVD) (e.g. Edelman et al., 1998, s. 4.3):

$$\mathbf{Q}_i^\top \mathbf{Q}_j =: \mathbf{L}_{i \rightarrow j} \text{diag}(\cos(\sigma_{i \rightarrow j})) \mathbf{R}_{i \rightarrow j}^\top, \quad \mathbf{L}, \mathbf{R} \text{ orthogonal.} \quad (9)$$

Since $\mathbf{Q}_i, \mathbf{Q}_j$ have orthonormal columns, $\cos(\sigma_{i \rightarrow j}) \in [0, 1]^k$ (e.g. Neretin, 2001), and more similar spans will yield larger singular values, which translate to smaller rotations. Importantly, singular values are invariant under similarity:

$$\text{diag}(\cos(\sigma_{i \rightarrow j})) = \mathbf{L}_{i \rightarrow j}^\top \mathbf{Q}_i^\top \mathbf{Q}_j \mathbf{R}_{i \rightarrow j} \quad (10)$$

$$= \mathbf{L}'_{i \rightarrow j} (\mathbf{Z}_i^\top \mathbf{Q}_i^\top) (\mathbf{Q}_j \mathbf{Z}_j) \mathbf{R}'_{i \rightarrow j}. \quad (11)$$

In other words, they are invariant under the action of any orthogonal \mathbf{Z} , which means that σ does not change if we replace an input matrix with any other matrix from the same equivalence class (see definition for $\mathcal{S}_i^\mathbb{O}$ in Section 4.1). The family of functions that satisfy this invariance, plus the axioms of metric spaces, form the family of *Grassmannian metrics* (e.g. Qiu et al., 2005), each capturing a different notion of distance between subspaces (largest principal angle, sum of principal angles...). See Section 4.1 for more details on specific metrics.

A.4 Synthetic Experiment on Grassmannian Metrics

Here we provide details about the synthetic experiment discussed in Section 4.2. We compute the reviewed Grassmannian metrics **a**) to **e**) from Section 4.1 between randomly drawn matrices from $\mathbb{O}^{D \times k}$ and masks from $\mathbb{M}^{D \times k}$. We use the subgroup algorithm (Diaconis & Shahshahani, 1987) to sample matrices uniformly from $\mathbb{O}^{D \times k}$ and column permutations of $\mathbf{I}_{D,k}$ to sample uniformly from $\mathbb{M}^{D \times k}$. To determine if the metrics behave differently for masks than for general orthogonal matrices, we inspect three different *modalities*: pairs of matrices (\mathbb{O} -to- \mathbb{O} , Figures 8 and 9), masks (\mathbb{M} -to- \mathbb{M} , Figures 10 and 11), and matrix-mask pairs (\mathbb{O} -to- \mathbb{M} , Figures 12 and 13). We normalize all metrics, denoted by $\text{dist}_{\mathbb{F}}$, to be in $[0, 1]$, with 1 indicating highest similarity (i.e. smallest distance). The overall procedure is gathered in Alg. 3.

Algorithm 3: Synthetic experiment on Grassmannian metrics (see Section 4.2 for details).

```

Input:  $\{D_1, D_2, \dots\}$  // Matrix height ( $D_i \in \mathbb{N}$ )
Input:  $\{r_1, r_2, \dots\}$  // Width-to-height ratio ( $r_i \in [0, 1]$ )
Input:  $\{(\mathbb{O}, \mathbb{O}), (\mathbb{O}, \mathbb{M}), (\mathbb{M}, \mathbb{M})\}$  // Modality
Input:  $\{\text{dist}_{\bar{g}}, \text{dist}_{\bar{c}, \bar{c}}, \text{dist}_{\bar{c}, \bar{F}}, \text{dist}_{\bar{p}, \bar{c}}, \text{dist}_{\bar{p}, \bar{F}}, \text{dist}_{\bar{a}}, \text{overlap}\}$  // Metric (normalized)
Input:  $T$  // Number of random samples
1  $\mathcal{R} \leftarrow \emptyset$  // Result (a dictionary)
2 for  $d \in \{D_1, D_2, \dots\}$  do
3   for  $r \in \{r_1, r_2, \dots\}$  do
4      $k \leftarrow \max(1, \text{round}(r \cdot d))$ 
5     for  $\text{dist} \in \{\text{dist}_{\bar{g}}, \text{dist}_{\bar{c}, \bar{c}}, \text{dist}_{\bar{c}, \bar{F}}, \text{dist}_{\bar{p}, \bar{c}}, \text{dist}_{\bar{p}, \bar{F}}, \text{dist}_{\bar{a}}, \text{overlap}\}$  do
6       for  $(\mathcal{M}_1, \mathcal{M}_2) \in \{(\mathbb{O}, \mathbb{O}), (\mathbb{O}, \mathbb{M}), (\mathbb{M}, \mathbb{M})\}$  do
7          $\mathcal{H} \leftarrow \emptyset$  // Collection of samples
8         for  $\{1, \dots, T\}$  do
9            $(\mathbf{Q}_1, \mathbf{Q}_2) \stackrel{\text{unif.}}{\sim} (\mathcal{M}_1^{d \times k}, \mathcal{M}_2^{d \times k})$   $\mathcal{H} \leftarrow \mathcal{H} \cup \text{dist}(\text{span}(\mathbf{Q}_1), \text{span}(\mathbf{Q}_2))$ 
10        end
11         $\mathcal{R}_{[d, r, \text{dist}, \mathcal{M}_1, \mathcal{M}_2]} \leftarrow \mathcal{H}$  // Gather samples into result
12      end
13    end
14  end
15 end
16 return  $\mathcal{R}$ 

```

We want to inspect how the distribution of the metrics changes as a function of sparsity k and sparsity ratio $\rho = \frac{k}{D}$. First, given fixed values of ρ , we study the distribution of all Grassmannian metrics as a function of D (Figures 8, 10 and 12). Additionally, given fixed values of D we study how the Grassmannian metrics change as a function of ρ (Figures 9, 11 and 13). We used the following values:

- Number of random (matrix or mask) samples: $T = 50$
- For Figures 8, 10 and 12 we investigate four different (fixed) ratios $\rho \in \{0.4, 0.2, 0.05, 0.01\}$ at several increasing dimensions $d \in \{16, 32, 64, 128, 256, 512, 1024, 2048\}$.
- For Figures 9, 11 and 13 we investigated four (fixed) dimensions $d \in \{128, 256, 512, 1024\}$ at several increasing ratios $\rho \in \{0.005, 0.01, 0.05, 0.1, 0.33, 0.66, 0.9, 0.95, 0.99, 1\}$.

Note how, in the second case, we concentrate the sparsity ratios around the extremes. This is to better capture the behaviour of collapsing metrics, as discussed in Section 4.2.

A.5 Expectation of *overlap* for Uniformly Random Matrices

In Section 4.2 we empirically observed that the expectation under uniformly distributed random matrices becomes *predictable* for all reviewed metrics, converging to a baseline value. Here we show that, for *overlap*, such expectation equals exactly $\rho = \frac{k}{D}$. This lemma depends on a standard calculation that was communicated to us by Joel A. Tropp.

Lemma A.1

Let $\mathbf{Q}_1, \mathbf{Q}_2$ be random matrices drawn uniformly from the Stiefel manifold $\mathbb{O}^{D \times k} := \{\mathbf{Q} : \mathbf{Q} \in \mathbb{R}^{D \times k}, \mathbf{Q}^\top \mathbf{Q} = \mathbf{I}_k\}$. Then,

$$\mathbb{E}[\text{overlap}(\text{span}(\mathbf{Q}_1), \text{span}(\mathbf{Q}_2))] = \frac{k}{D} \quad (12)$$

Proof. We start by rewriting the definition of *overlap*, presented in Section 4.1, in terms of the trace of orthogonal projectors $\Psi = \mathbf{Q}\mathbf{Q}^\top \in \mathbb{R}^{D \times D}$. We make use of the idempotence of Ψ and the unitary invariance

of the Frobenius norm:

$$\text{overlap}(\text{span}(\mathbf{Q}_1), \text{span}(\mathbf{Q}_2)) = \frac{1}{k} \|\mathbf{Q}_1^\top \mathbf{Q}_2\|_F^2 = \frac{1}{k} \|\Psi_1 \Psi_2\|_F^2 = \frac{1}{k} \text{Tr}(\Psi_1 \Psi_2^2 \Psi_1) \quad (13)$$

$$= \frac{1}{k} \text{Tr}(\Psi_1 \Psi_2 \Psi_1) \quad (14)$$

We further observe that, if \mathbf{Q} is drawn uniformly from $\mathbb{O}^{D \times k}$, the marginal distribution of every column \mathbf{q} is the uniform distribution over the Euclidean unit sphere, hence it is isotropic:

$$\mathbb{E}[\mathbf{q}\mathbf{q}^\top] = \frac{1}{D} \mathbf{I}_D \quad (15)$$

Where \mathbf{I}_D is the rank- D identity matrix. Then, leveraging linearity of expectation, the orthogonal projector $\Psi = \mathbf{Q}\mathbf{Q}^\top$ can be expressed as follows:

$$\mathbb{E}[\Psi] = \sum_{i=1}^k \mathbb{E}[\mathbf{q}_i \mathbf{q}_i^\top] = \frac{k}{D} \mathbf{I}_D \quad (16)$$

Now, given two independent realizations $\mathbf{Q}_1, \mathbf{Q}_2$, we form the associated orthogonal projectors Ψ_1, Ψ_2 . Write $\mathbb{E}_1, \mathbb{E}_2$ for the expectations of the respective distributions of \mathbf{Q}_1 and \mathbf{Q}_2 . Then, leveraging independence of \mathbf{Q}_1 and \mathbf{Q}_2 , idempotence of Ψ , linearity of expectation, and $\text{Tr}(\Psi) = k$, we have:

$$\mathbb{E}[\text{Tr}(\Psi_1 \Psi_2 \Psi_1)] = \mathbb{E}_1[\text{Tr}(\Psi_1 \mathbb{E}_2[\Psi_2] \Psi_1)] = \frac{k}{D} \mathbb{E}_1[\text{Tr}(\Psi_1 \mathbf{I}_D \Psi_1)] = \frac{k}{D} \text{Tr}(\Psi_1) = \frac{k^2}{D} \quad (17)$$

Replacing in the definition of *overlap* concludes the proof. \square

Table 2: Empirical baselines for Grassmannian metrics. Shown are the measured expectations (averaged over 50 samples at $D=2048$) of different Grassmannian metrics between two uniformly random matrices in \mathbb{O} for different values of $\rho = \frac{k}{D}$. Note how *overlap* converges to exactly ρ (see Appendix A.5).

Metric	ρ				
	0.005	0.01	0.05	0.2	0.4
$\text{dist}_{\bar{g}}$	0.03711	0.05191	0.11909	0.24534	0.36536
$\text{dist}_{\bar{c}, \bar{z}}$	0.00197	0.00161	0.00048	0.00046	0.00026
$\text{dist}_{\bar{c}, \bar{F}}$	0.02983	0.04202	0.09984	0.21754	0.33783
$\text{dist}_{\bar{p}, \bar{z}}$	10^{-5}	10^{-5}	0	0	0
$\text{dist}_{\bar{p}, \bar{F}}$	0.00248	0.00479	0.02513	0.1057	0.22527
$\text{dist}_{\bar{a}}$	0	0	0	0	0
<i>overlap</i>	0.00495	0.00955	0.04962	0.20022	0.39980

A.6 Extended Discussion on Grassmannian Metrics

We extend here the main insights presented in Section 4.2:

Mask-vs-mask metrics have larger variance, and in some cases lower expectations: All distributions for the mask-vs-mask modality have higher variance compared to the other modalities (see Appendix A.4). Furthermore, the distributions for the $\text{dist}_{\bar{g}}$ and $\text{dist}_{\bar{c}, \bar{F}}$ metrics seem to follow lower trajectories for the mask-vs-mask modality, compared to the other modalities. This is not the case for $\text{dist}_{\bar{p}, \bar{F}}$ and *overlap*, whose expectation does not seem to be affected by the modality.

Extremal values of ρ lead to saturation, except for *overlap*: Most metrics exhibit a nonlinear behaviour near the extremes (Figure 9). This is particularly so for collapsing metrics, but saturation can also be observed in the non-collapsing ones. The only exception is *overlap*, whose expectation is linear as shown in Appendix A.5.

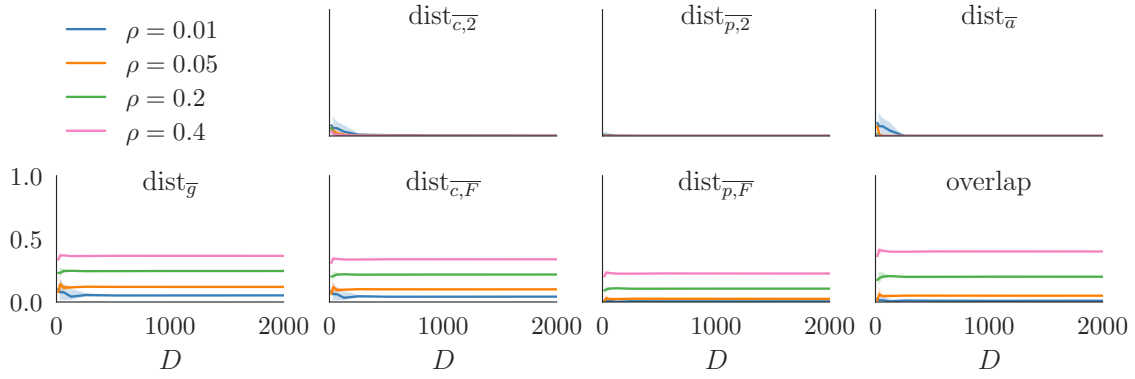


Figure 8: Grassmannian metrics for random pairs of matrices $(Q_1, Q_2) \stackrel{\text{unif.}}{\sim} (\mathbb{O}^{D \times k}, \mathbb{O}^{D \times k})$ as a function of D . Each subplot shows a different metric, and each color corresponds to a different ratio $r = \frac{k}{D}$. For each (D, k) , we sample 50 random pairs and report the median of the resulting distribution (line plot) as well as the 5-95 percentiles (shaded regions).




Figure 9: Grassmannian metrics for random pairs of matrices $(Q_1, Q_2) \stackrel{\text{unif.}}{\sim} (\mathbb{O}^{D \times k}, \mathbb{O}^{D \times k})$ as a function of $\rho = \frac{k}{D}$. Each subplot shows a different metric, and each color corresponds to a different dimension D . For each (D, k) , we sample 50 random pairs and report the median of the resulting distribution (line plot) as well as the 5-95 percentiles (shaded regions).




Figure 10: Grassmannian metrics for random pairs of masks $(Q_1, Q_2) \stackrel{\text{unif.}}{\sim} (\mathbb{M}^{D \times k}, \mathbb{M}^{D \times k})$ as a function of D and for fixed ρ . Each subplot shows a different Grassmannian metric with the different lines indicating four different ratios ρ . For each value of D , we report the median metric over 50 random pairs with the shaded regions showing the 5-95 percentiles.




Figure 11: Grassmannian metrics for random pairs of masks $(Q_1, Q_2) \stackrel{\text{unif.}}{\sim} (\mathbb{M}^{D \times k}, \mathbb{M}^{D \times k})$ as a function of ρ and for fixed D . Each subplot shows a different Grassmannian metric with the different lines indicating four different dimensions D . For each value of ρ , we report the median metric over 50 random pairs with the shaded regions showing the 5-95 percentiles.




Figure 12: Grassmannian metrics for random pairs of matrices and masks $(Q_1, Q_2) \stackrel{\text{unif.}}{\sim} (\mathbb{O}^{D \times k}, \mathbb{M}^{D \times k})$ as a function of D and for fixed ρ . Each subplot shows a different Grassmannian metric with the different lines indicating four different ratios ρ . For each value of D , we report the median metric over 50 random pairs with the shaded regions showing the 5-95 percentiles.




Figure 13: Grassmannian metrics for random pairs of matrices and masks $(Q_1, Q_2) \stackrel{\text{unif.}}{\sim} (\mathbb{O}^{D \times k}, \mathbb{M}^{D \times k})$ as a function of ρ and for fixed D . Each subplot shows a different Grassmannian metric with the different lines indicating four different dimensions D . For each value of ρ , we report the median metric over 50 random pairs with the shaded regions showing the 5-95 percentiles.

A.7 Analytical *overlap* Baseline vs. Sampling Random DL Masks

In Section 4.2 we justify the use of $\rho := k/D$ as an analytical, “chance-level” baseline for *overlap*, that we use in Section 6 to observe that the overlap between spaces spanned by magnitude pruning masks and top Hessian eigenspaces in DL is substantially large (see e.g. Figures 3 and 4).

In this context, one natural question one may ask is: *Do random DL masks also lead to higher-than-baseline overlaps?* To answer this question, we trained the 16×16 MNIST setup following Appendix A.2, and also for steps $t \in \{0, 100, 1000\}$, we computed the *overlap* between the top- k Hessian eigenspace and 10 randomly chosen masks of k nonzeros (the steps were also chosen to cover initialization, mid-training and convergence, as illustrated in Figure 5). Results are shown in Figure 14: We observe that, indeed, *random DL masks tend to behave like our analytical baseline*, which supports the use of the analytical baseline in our experiments.

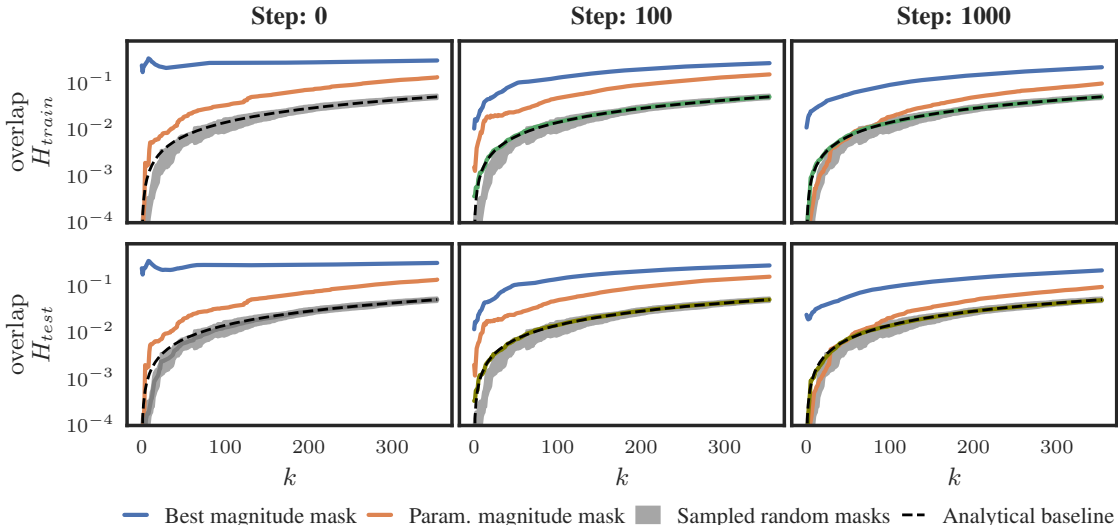


Figure 14: *overlap* measurements for different masking criteria against H_{train} (top row) and H_{test} (bottom row). *Best magnitude mask* represents the *overlap* obtained if we pick the k parameters with largest *overlap* with the top- k eigenspace. *Param. magnitude mask* corresponds to the k -largest parameters by magnitude. The *random mask* grey area covers \pm one standard deviation for the *overlap* between H and 10 randomly sampled masks. See Figure 5 for respective accuracies at given steps.

Another natural follow-up question is: *How special are the parameter magnitude masks? Would some other non-random mask also achieve a high overlap?* To tackle this question, we also plot in Figure 14 the *overlap* for the top- k parameter magnitude masks, and the maximal *overlap* achievable by the best possible k -mask, serving as upper bound. We observe that, *while the magnitude pruning mask is substantially larger than baseline, it is still far from optimal*.

Still, there is a limited number of masks that can yield high *overlap*: Due to the orthogonality of the Hessian eigenbasis, each column of $\tilde{U}^{(k)}$ must have an ℓ_2 norm of 1. This means that if a given set of parameters has an *overlap* of α , the *overlap* for any non-intersecting set of parameters is at most $1 - \alpha$. Intuitively, the *overlap* is a limited resource, and therefore, masks with an *overlap* significantly above chance-level are bound to be scarce.

B Supplementary Material for the Experiments in Section 6

B.1 Results for 16×16 MNIST

This section includes supplementary material as part of the exhaustive experiments done for the MNIST toy problem. Figure 15 and Figure 16 verify the existence of early collapse and stabilization in both Hessian top space and parameter masks. They also show that SGD was able to successfully train the model under this

setting. Figure 17 corroborates that the behaviour of the Grassmannian metrics analyzed in Section 4 is also present for DL problems.




Figure 15: Both the parameters of the a neural network as well as their Hessian spectrum collapse early during training. (*top*) The top two subplots show the mini-batch loss L_B as well as the train/validation/test accuracy of the 7030-parameter model trained on 16×16 downsampled MNIST (see Section 6.1). (*middle*) Looking at the top 20%, 5%, 1%, 0.5% of parameters by magnitude, we can see that very early during training, most of the energy is concentrated on a small subset of the parameters (see Section 2.2 for a definition of κ). For example, shortly after initialization, the top 0.5% largest parameters by magnitude have roughly $1/4$ of the total ℓ_2 norm of all parameters. (*bottom*) We can observe a similar behaviour for the Hessian spectrum on both the training set (fourth subplot) and the test set (fifth subplot). Only a few steps after training, most of the energy is concentrated in only 0.5% of the eigenvalues.

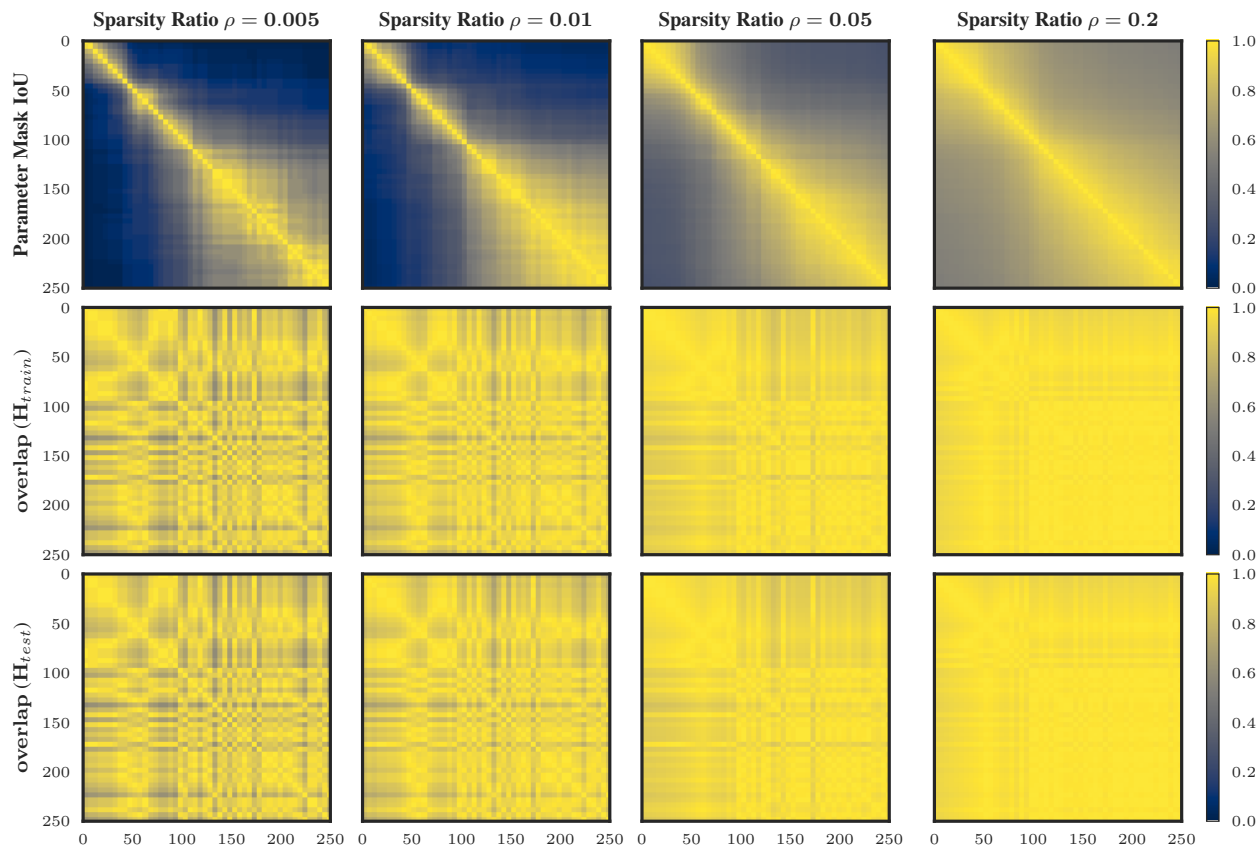


Figure 16: Both parameter pruning masks (top row) and Hessians (bottom two rows) remain relatively stable after an initial training phase. Following You et al. (2020), the depicted matrices represent all pairwise similarities (i.e. higher is better) from the beginning of training (top left corners) until step 250 (bottom right corners), for the 7030-parameter model trained on 16×16 downsampled MNIST (see Section 6.1). For this reason, all matrices are symmetric and have unit diagonals. (*top*) Even when selecting only 0.5% of the parameters (left column), masks collected at different training steps show a remarkable similarity after an initial phase of training. (*middle and bottom*) The *overlap* metric for top Hessian eigenspaces on the train (middle) and test set (bottom) extracted at different training steps and for different subspace sizes (columns). Starting from initialization, the Hessian eigenspaces do not change significantly over the course of training. No substantial differences in behaviour between \mathbf{H}_{test} and \mathbf{H}_{train} can be observed.




Figure 17: Different Grassmannian metrics between pruning masks and top Hessian subspaces. Each line shows a particular sparsity level ρ , i.e., the ratio of unpruned parameters or the size of the top Hessian subspace relative to the full Hessian, as a function of training progress for the 7030-parameter model trained on 16×16 downsampled MNIST (see Section 6.1). All non-collapsing metrics reveal a consistent and substantial similarity between spaces spanned by pruning masks and top Hessian subspaces well above random chance (random baselines gathered from our synthetic experiments are shown in dashed lines for $\rho = 0.2$), while collapsing metrics are effectively zero due to the large D . For better visibility, we limit the plot to the first 1000 steps but all metrics remain stable afterwards.

B.2 Discussion on Computational Resources

We now discuss the resources involved in the computation of the Hessian eigendecompositions, as well as the experimental design choices surrounding them. The goal is to reflect the scale of computations involved in our experiments, and to convey an idea of the overall resources needed to reproduce them. For more rigorous asymptotics involving memory, arithmetic and measurement costs, readers are invited to review the references exhaustively provided in Section 5, particularly (Tropp et al., 2019, Sec. 4).




Figure 18: Runtimes of main SEIGH operations to compute a single Hessian eigendecomposition, assuming a single computer with 400GB RAM equipped with an NVIDIA A100 (40GB) graphics card. Note that the three involved eigendecompositions feature different networks, datasets, and rank, so a blind one-to-one comparison of runtimes is not possible: refer to Appendix B.2 for an interpretation of this figure.

Figure 18 gathers the average times it took to compute a single sketched Eigendecomposition, for each one of the Hessians covered in Section 6.2. Specifically:

- Resources needed to tune or train the model are here ignored. We used PyTorch Paszke et al. (2019).
- The *layout* step allocates the $\mathcal{O}(Dk)$ storage memory in the form of virtual HDF5 datasets⁶. This allows to efficiently write numerical data in parallel, which is instrumental for the scalability of the procedure.
- The HVPs, computed with the help of CurvLinOps⁷, correspond to the measurements in lines 1-2 of Alg. 2. As we can see, they take the bulk of the runtime: each HVP requires two forward and backward passes over the Hessian dataset (consisting of 500, 1000 and 5000 samples for the three respective problems, see N_{train}/N_{test} in Table 1). While this can be done in batched fashion, the HVP runtime is linearly affected by dataset size. Another linear factor is the number of measurements (see n_o in Table 1). Crucially, those can be fully parallelized, so the runtimes reported under the *HVPs* column can be divided by the number of available machines. This property of sketched methods, combined with the HDF5 technology, is the main enabler for scalability here.
- The *merge* step stems from a technicality: most operative systems don't allow a process to open thousands of files at the same time. Hence, we need to merge the virtual HDF5 dataset into a monolithic file. This can be done with only $\mathcal{O}(D)$ memory overhead.
- Although (Tropp et al., 2019, Sec. 4) points out that the *QR decomposition* (line 3 of Alg. 2) dominates the (asymptotic) arithmetic cost for the whole procedure, interestingly this a relatively fast step. We also note that the main reason why we needed so much RAM memory was the need to load $\mathcal{O}(Dk)$ entries in-memory to perform this step. Out-of-core QR orthogonalizations would greatly help reducing RAM requirements.
- The remaining operations (lines 4-8 in Alg. 2) are here gathered under the *lstsq* label, which takes a considerable portion of the overall runtime. To improve this, lines 5 and 6 could be parallelized, roughly halving the runtime, but still faster solutions would greatly help towards more scalable Hessian eigendecompositions.

⁶<https://github.com/HDFGroup/hdf5>

⁷<https://github.com/f-dangel/curvlinops>

Figure 18 also reports the *TOTAL* runtime of the steps purely involved with the eigendecomposition. Further *overheads*, related e.g. to the use of a computational cluster, are reported separately. In practice, computing e.g. a single RESNET-18 Hessian eigendecomposition with up to fifteen A100s took a bit over a day. While this may seem like a heavyweight computation, it is still a substantial improvement, enabling previously intractable computations. In this section we have also outlined a few ways in which this could be further improved.

Due to the scale of the experiments, we had to be conservative with the choice of hyperparameters. We chose dataset sizes N_{train}/N_{test} such that all classes are balanced, and each class has at least 5 examples. While the chosen number of measurements n_o is much smaller than the ambient dimension D , the tremendously rank-deficient structure of \mathbf{H} (see Section 2.1) allows for such a disparity. Still, it has been noted that the numerical rank of \mathbf{H} may be linked to the number of classes (eg Gur-Ari et al., 2018; Papyan, 2020; Dangel et al., 2022). For this reason we chose n_o values well above this limit.