

---

# Relative Flatness and Generalization

---

**Henning Petzka\***

Lund University, Sweden  
henning.petzka@math.lth.se

**Michael Kamp\***

CISPA Helmholtz Center for Information Security,  
Germany and Monash University, Australia  
michael.kamp@monash.edu

**Linara Adilova**

Ruhr University Bochum, Germany  
and Fraunhofer IAIS

**Cristian Sminchisescu**

Lund University, Sweden  
and Google Research, Switzerland

**Mario Boley**

Monash University, Australia

## Abstract

Flatness of the loss curve is conjectured to be connected to the generalization ability of machine learning models, in particular neural networks. While it has been empirically observed that flatness measures consistently correlate strongly with generalization, it is still an open theoretical problem why and under which circumstances flatness is connected to generalization, in particular in light of reparameterizations that change certain flatness measures but leave generalization unchanged. We investigate the connection between flatness and generalization by relating it to the interpolation from representative data, deriving notions of representativeness, and feature robustness. The notions allow us to rigorously connect flatness and generalization and to identify conditions under which the connection holds. Moreover, they give rise to a novel, but natural relative flatness measure that correlates strongly with generalization, simplifies to ridge regression for ordinary least squares, and solves the reparameterization issue.

## 1 Introduction

Flatness of the loss curve has been identified as a potential predictor for the generalization abilities of machine learning models [6, 10, 11]. In particular for neural networks, it has been repeatedly observed that generalization performance correlates with measures of flatness, i.e., measures that quantify the change in loss under perturbations of the model parameters [4, 8, 16, 21, 34, 39, 41, 44]. In fact, Jiang et al. [14] perform a large-scale empirical study and find that flatness-based measures have a higher correlation with generalization than alternatives like weight norms, margin-, and optimization-based measures. It is an open problem why and under which circumstances this correlation holds, in particular in the light of negative results on reparameterizations of ReLU neural networks [5]: these reparameterizations change traditional measures of flatness, yet leave the model function and its generalization unchanged, making these measures unreliable. We present a novel and rigorous approach to understanding the connection between flatness and generalization by relating it to the interpolation from representative samples. Using this theory we, for the first time, identify conditions under which flatness explains generalization. At the same time, we derive a measure of *relative flatness* that simplifies to ridge/Tikhonov regularization for ordinary least squares [36], and resolves

---

\*equal contribution

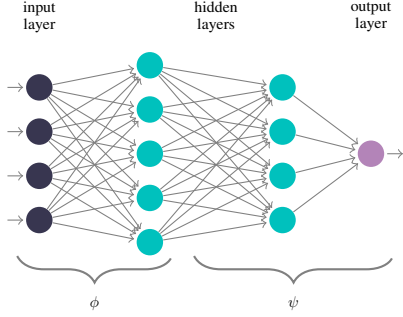


Figure 1: Decomposition of  $f = \psi \circ \phi$  into a feature extractor  $\phi$  and a model  $\psi$  for neural networks.

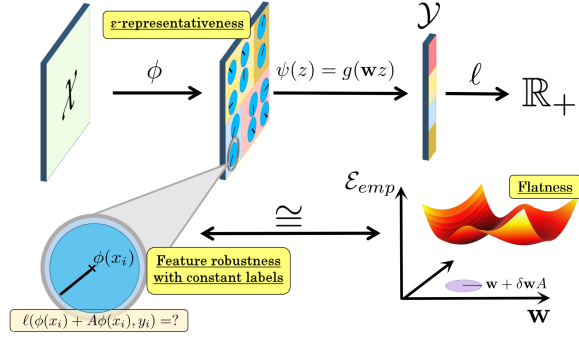


Figure 2: Overview: We theoretically connect a notion of representative data with a notion of feature robustness and a novel measure of flatness of the loss surface.

the reparametrization issue for ReLU networks [5] by appropriately taking the norm of parameters into account as suggested by Neyshabur et al. [25].

Formally, we connect flatness of the loss surface to the *generalization gap*  $\mathcal{E}_{gen}(f, S) = \mathcal{E}(f) - \mathcal{E}_{emp}(f, S)$  of a *model*  $f : \mathcal{X} \rightarrow \mathcal{Y}$  from a *model class*  $\mathcal{H}$  with respect to a twice differentiable *loss function*  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$  and a finite sample set  $S \subseteq \mathcal{X} \times \mathcal{Y}$ , where

$$\mathcal{E}(f) = \mathbb{E}_{(x,y) \sim \mathcal{D}} [\ell(f(x), y)] \quad \text{and} \quad \mathcal{E}_{emp}(f, S) = \frac{1}{|S|} \sum_{(x,y) \in S} \ell(f(x), y) .$$

That is,  $\mathcal{E}_{gen}(f, S)$  is the difference between the *risk*  $\mathcal{E}(f)$  and the *empirical risk*  $\mathcal{E}_{emp}(f, S)$  of  $f$  on a finite sample set  $S$  drawn iid. according to a *data distribution*  $\mathcal{D}$  on  $\mathcal{X} \times \mathcal{Y}$ . To connect flatness to generalization, we start by decomposing the generalisation gap into two terms, a *representativeness* term that quantifies how well a distribution  $\mathcal{D}$  can be approximated using distributions with local support around sample points and a *feature robustness* term describing how small changes of feature values affect the model’s loss. Here, feature value refers to the implicitly represented features by the model, i.e., we consider models that can be expressed as  $f(x) = \psi(\mathbf{w}, \phi(x)) = g(\mathbf{w}\phi(x))$  with a feature extractor  $\phi$  and a model  $\psi$  (which includes linear and kernel models, as well as most neural networks, see Fig. 1). With this decomposition, we measure the generalization ability of a particular model by how well its interpolation between samples in feature space fits the underlying data distribution. We then connect feature robustness (a property of the feature space) to flatness (a property of the parameter space) using the following key identity: Multiplicative perturbations in feature space by arbitrary matrices  $A \in \mathbb{R}^{m \times m}$  correspond to perturbations in parameter space, i.e.,

$$\psi(\mathbf{w}, \phi(x) + A\phi(x)) = g(\mathbf{w}(\phi(x) + A\phi(x))) = g((\mathbf{w} + \mathbf{w}A)\phi(x)) = \psi(\mathbf{w} + \mathbf{w}A, \phi(x)) . \quad (1)$$

Using this key equation, we show that feature robustness is approximated by a novel, but natural, loss Hessian-based *relative flatness* measure under the assumption that the distribution can be approximated by locally constant labels. Under this assumption and if the data is representative, then flatness is the main predictor of generalization (see Fig. 2 for an illustration).

This offers an explanation for the correlation of flatness with generalization on many real-world data distributions for image classification [14, 21, 26], where the assumption of locally constant labels is reasonable (the definition of adversarial examples [35] even hinges on this assumption). This dependence on locally constant labels has not been uncovered by previous theoretical analysis [26, 37]. Moreover, we show that the resulting relative flatness measure is invariant to linear reparameterization and has a stronger correlation with generalization than other flatness measures [14, 21, 26]. Other measures have been proposed that similarly achieve invariance under reparameterizations [21, 37], but the Fisher-Rao norm [21] is lacking a strong theoretical connection to generalization, our measure sustains a more natural form than normalized sharpness [37] and for neural networks, it considers only a single layer, given by the decomposition of  $f$  (including the possibility of choosing the input layer when  $\phi = id_{\mathcal{X}}$ ). An extended comparison to related work is provided in Appdx. A.

The limitations of our analysis are as follows. We assume a noise-free setting where for each  $x \in \mathcal{X}$  there is a unique  $y = y(x) \in \mathcal{Y}$  such that  $P_{x,y \sim \mathcal{D}}(y|x) = 1$ , and this assumption is also extended to

the feature space of the given model, i.e., we assume that  $\phi(x) = \phi(x')$  implies  $y(x) = y(x')$  for all  $x, x' \in \mathcal{X}$  and write  $y(x) = y(\phi(x))$ . Moreover, we assume that the marginal distribution  $\mathcal{D}_{\mathcal{X}}$  is described by a density function  $p_{\mathcal{D}}(x)$ , that  $f(x) = \psi(\mathbf{w}, \phi(x)) = g(\mathbf{w}\phi(x))$  is a local minimizer of the empirical risk on  $S$ , and that  $g, \psi, \phi$  are twice differential. Quantifying the representativeness of a dataset precisely is challenging since the data distribution is unknown. Using results from density estimation, we derive a worst-case bound on representativeness for all data distributions that fulfill mild regularity assumptions in feature space  $\phi(\mathcal{X})$ , i.e., a smooth density function  $p_{\phi(\mathcal{D})}$  such that  $\int_{z \in \phi(\mathcal{X})} |\nabla^2 (p_{\phi(\mathcal{D})}(z) \|z\|^2)| dz$  and  $\int_{z \in \phi(\mathcal{X})} p_{\phi(\mathcal{D})}(z) / \|z\|^m dz$  are well-defined and finite. This yields a generalization bound incorporating flatness. In contrast to the common bounds of statistical learning theory, the bound depends on the feature dimension. The dimension-dependence is a result of the interpolation approach (applying density estimation uniformly over all distributions that satisfy the mild regularity assumptions). The bound is consistent with the no-free-lunch theorem and the convergence rate derived by Belkin et al. [2] for a model based on interpolations. In practical settings, representativeness can be expected to be much smaller than the worst-case bound, which we demonstrate by a synthetic example in Sec. 6. Generally, it is a bound that remains meaningful in the interpolation regime [1, 3, 24], where traditional measures of generalization based on the empirical risk and model class complexity are uninformative [22, 42].

**Contribution.** In summary, this paper rigorously connects flatness of the loss surface to generalization and shows that this connection requires feature representations such that labels are (approximately) locally constant, which is also validated in a synthetic experiment (Sec. 6). The empirical evaluation shows that this flatness and an approximation to representativeness can tightly bound the generalization gap. Our contributions are: (i) the rigorous connection of flatness and generalization; (ii) novel notions of representativeness and feature robustness that capture the extent to which a model’s interpolation between samples fits the data distribution; and (iii) a novel flatness measure that is layer- and neuron-wise reparameterization invariant, reduces to ridge regression for ordinary least squares, and outperforms state-of-the-art flatness measures on CIFAR10.

## 2 Representativeness

In this section, we formalize when a sample set  $S$  is representative for a data distribution  $\mathcal{D}$ .

**Partitioning the input space.** We choose a partition  $\{V_i \mid i = 1, \dots, |S|\}$  of  $\mathcal{X}$  such that each element of this partition  $V_i$  contains exactly one of the samples  $x_i$  from  $S$ . The distribution can then be described by a set of densities  $p_i(x) = \frac{1}{\alpha_i} \cdot p_{\mathcal{D}}(x) \cdot \mathbf{1}_{V_i}(x)$  with support contained in  $V_i$  (where  $\mathbf{1}_{V_i}(x) = 1$  if  $x \in V_i$  and 0 otherwise) and with normalizing factor  $\alpha_i = \int_{V_i} p_{\mathcal{D}}(x) dx$ . Then the risk decomposes as  $\mathcal{E}(f) = \sum_{i=1}^{|S|} \alpha_i \cdot \mathbb{E}_{x \sim p_i} [\ell(f(x), y(x))]$ . Since  $x_i \in V_i$  for each  $i$ , we can change variables and consider density functions  $\lambda_i^*(\xi) = p_i(x_i + \xi)$  with support in a neighborhood around the origin of  $\mathcal{X}$ . The risk then decomposes as

$$\mathcal{E}(f) = \sum_{i=1}^{|S|} \alpha_i \cdot \mathbb{E}_{\xi \sim \lambda_i^*} [\ell(f(x_i + \xi), y(x_i + \xi))] . \quad (2)$$

Starting from this identity, we formalize an approximation to the risk: In a practical setting, the distribution  $p_{\mathcal{D}}$  is unknown and hence, in the decomposition (2), we have unknown densities  $\lambda_i^*$  and unknown normalization factors  $\alpha_i$ . We assume that each neighborhood contributes equally to the loss, i.e., we approximate each  $\alpha_i$  with  $\frac{1}{|S|}$ . Then, given a sample set  $S$  and an  $|S|$ -tuple  $\Lambda = (\lambda_i)_{1 \leq i \leq |S|}$  of “local” probability density functions on  $\mathcal{X}$  with support  $\text{supp}(\lambda_i)$  in a neighborhood around the origin  $0_{\mathcal{X}}$ , we call the pair  $(S, \Lambda)$   $\epsilon$ -representative for  $\mathcal{D}$  with respect to a model  $f$  and loss  $\ell$  if  $|\mathcal{E}_{Rep}(f, S, \Lambda)| \leq \epsilon$ , where

$$\mathcal{E}_{Rep}(f, S, \Lambda) = \mathcal{E}(f) - \sum_{i=1}^{|S|} \frac{1}{|S|} \cdot \mathbb{E}_{\xi \sim \lambda_i} [\ell(f(x_i + \xi), y(x_i + \xi))] . \quad (3)$$

If the partitions  $V_i$  and the distributions  $\lambda_i$  are all chosen optimal so that the approximation  $\alpha_i = \frac{1}{|S|}$  is exact and  $\lambda_i = \lambda_i^*$ , then  $\mathcal{E}_{Rep}(f, S, \Lambda) = 0$  by (2). If the support of each  $\lambda_i$  is decreased to the origin so that  $\lambda_i = \delta_0$  is a Dirac delta function, then  $\mathcal{E}_{Rep}(f, S, \Lambda) = \mathcal{E}_{gen}(f, S)$  equals the

generalization gap. For density functions with an intermediate support, the generalization gap can be decomposed into representativeness and the expected deviation of the loss around the sample points:

$$\mathcal{E}_{gen}(f, S) = \mathcal{E}_{Rep}(f, S, \Lambda) + \sum_{i=1}^{|S|} \frac{1}{|S|} \cdot \mathbb{E}_{\xi \sim \lambda_i} [\ell(f(x_i + \xi), y(x_i + \xi)) - \ell(f(x_i), y_i)]$$

The main idea of our approach to understand generalization is to use this equality and to control both representativeness and expected loss deviations for a suitable  $|S|$ -tuple of distributions  $\Lambda$ .

**From input to feature space.** An interesting aspect of  $\epsilon$ -representativeness is that it can be considered in a feature space instead of the input space. For a model  $f = (\psi \circ \phi) : \mathcal{X} \rightarrow \mathcal{Y}$ , we can apply our notion to the feature space  $\phi(\mathcal{X})$  (see Fig. 1 for an illustration). This leads to the notion of  $\epsilon$ -representativeness in feature space defined for an  $|S|$ -tuple  $\Lambda^\phi = (\lambda_i^\phi)_{1 \leq i \leq |S|}$  of densities on  $\phi(\mathcal{X})$  by replacing  $x_i$  with  $\phi(x_i)$  in (3), which we denote by  $\mathcal{E}_{Rep}^\phi(f, S, \Lambda^\phi)$ . By measuring representativeness in a feature space, this becomes a notion of both data and feature representation. In particular, it assumes that a target output function  $y(\phi(x))$  also exists for the feature space. We can then decompose the generalization gap  $\mathcal{E}_{gen}(f)$  of  $f = (\psi \circ \phi)$  into

$$\mathcal{E}_{Rep}^\phi(f, S, \Lambda^\phi) + \left( \frac{1}{|S|} \sum_{i=1}^{|S|} \mathbb{E}_{\xi \sim \lambda_i^\phi} [\ell(\psi(\phi(x_i) + \xi), y(\phi(x_i) + \xi)) - \ell(f(x_i), y_i)] \right)$$

The second term is determined by how the loss changes under small perturbations in the feature space for the samples in  $S$ . As before, for  $\lambda_i = \delta_0$  the term in the bracket vanishes and  $\mathcal{E}_{Rep}^\phi(f, S, \Lambda^\phi) = \mathcal{E}_{gen}$ . But the decomposition becomes more interesting for distributions with support of nonzero measure around the origin. If the true distribution can be interpolated efficiently in feature space from the samples in  $S$  with suitable  $\lambda_i^\phi$  so that  $\mathcal{E}_{Rep}^\phi(f, S, \Lambda^\phi) \approx 0$ , then the term in the bracket approximately equals the generalization gap and the generalization gap can be estimated from local properties in feature space around sample points.

### 3 Feature Robustness

Having decomposed the generalisation gap into a representativeness and a second term of loss deviation, we now develop a novel notion of feature robustness that is able to bound the second term for specific families of distributions  $\Lambda$  using key equation (1). Our definition of *feature robustness* for a model  $f = (\psi \circ \phi) : \mathcal{X} \rightarrow \mathcal{Y}$  depends on a small number  $\delta > 0$ , a sample set  $S$  and a *feature selection* defined by a matrix  $A \in \mathbb{R}^{m \times m}$  of *operator norm*  $\|A\| \leq 1$ . With feature perturbations  $\phi_A(x) = (I + A)\phi(x)$  and

$$\mathcal{E}_{\mathcal{F}}^\phi(f, S, A) := \frac{1}{|S|} \sum_{i=1}^{|S|} [\ell(\psi(\phi_A(x_i)), y[\phi_A(x_i)]) - \ell(f(x_i), y_i)], \quad (4)$$

the definition of feature robustness is given as follows.

**Definition 1.** Let  $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$  denote a loss function,  $\epsilon$  and  $\delta$  two positive (small) real numbers,  $S \subseteq \mathcal{X} \times \mathcal{Y}$  a finite sample set, and  $A \in \mathbb{R}^{m \times m}$  a matrix. A model  $f(x) = (\psi \circ \phi)(x)$  with  $\phi(\mathcal{X}) \subseteq \mathbb{R}^m$  is called  $((\delta, \mathbf{S}, \mathbf{A}), \epsilon)$ -feature robust, if  $\left| \mathcal{E}_{\mathcal{F}}^\phi(f, S, \alpha A) \right| \leq \epsilon$  for all  $0 \leq \alpha \leq \delta$ . More generally, for a probability distribution  $\mathcal{A}$  on perturbation matrices in  $\mathbb{R}^m$ , we define

$$\mathcal{E}_{\mathcal{F}}^\phi(f, S, \mathcal{A}) = \mathbb{E}_{A \sim \mathcal{A}} \left[ \mathcal{E}_{\mathcal{F}}^\phi(f, S, A) \right],$$

and call the model  $((\delta, \mathbf{S}, \mathbf{A}), \epsilon)$ -feature robust on average over  $\mathcal{A}$ , if  $\left| \mathcal{E}_{\mathcal{F}}^\phi(f, S, \alpha A) \right| \leq \epsilon$  for  $0 \leq \alpha \leq \delta$ .

Given a feature extractor  $\phi$ , feature robustness measures the performance of  $\psi$  when feature values are perturbed (with constant feature extractor  $\phi$ ). This local robustness at sample points differs from the robustness of Xu and Mannor [40] that requires a data-independent partitioning of the input space. The matrix  $A$  in feature robustness determines which feature values shall be perturbed. For

each sample, the perturbation is linear in the expression of the feature. Thereby, we only perturb features that are relevant for the output for a given sample and leave feature values unchanged that are not expressed. For  $\phi$  mapping into an intermediate layer of a neural network, traditionally, the activation values of a neuron are considered as feature values, which corresponds to a choice of  $A$  as a projection matrix. However, it was shown by Szegedy et al. [35] that, for any other direction  $v \in \mathbb{R}^m, \|v\| = 1$ , the values  $\langle \phi(x), v \rangle$  obtained from the projection  $\phi(x)$  onto  $v$ , can be likewise semantically interpreted as a feature. This motivates the consideration of general *feature matrices*  $A$ .

**Distributions on feature matrices induce distributions on the feature space** Feature robustness is defined in terms of feature matrices (suitable for an application of (1) to connect perturbations of features with perturbations of weights), while the approach exploiting representative data from Section 2 considers distributions on feature vectors, cf. (3). To connect feature robustness to the notion of  $\epsilon$ -representativeness, we specify for any distribution  $\mathcal{A}$  on matrices  $A \in \mathbb{R}^{m \times m}$  an  $|S|$ -tuple  $\Lambda_{\mathcal{A}} = (\lambda_i)$  of probability density functions  $\lambda_i$  on the feature space  $\mathbb{R}^m$  with support containing the origin. Multiplication of a feature matrix with a feature vector  $\phi(x_i)$  defines a feature selection  $A\phi(x_i)$ , and for each  $z \in \mathbb{R}^m$  there is some feature matrix  $A$  with  $\phi(x_i) + z = \phi(x_i) + A\phi(x_i)$  (unless  $\phi(x_i) = 0$ ). Our choice for distributions  $\lambda_i$  on  $\mathbb{R}^m$  are therefore distributions that are induced via multiplication of feature vectors  $\phi(x_i) \in \mathbb{R}^m$  with matrices  $A \in \mathbb{R}^{m \times m}$  sampled from a distribution on feature matrices  $\mathcal{A}$ . Formally, we assume that a Borel measure  $\mu_{\mathcal{A}}$  is defined by a probability distribution  $\mathcal{A}$  on matrices  $\mathbb{R}^{m \times m}$ . We then define Borel measures  $\mu_i$  on  $\mathbb{R}^m$  by  $\mu_i(C) = \mu_{\mathcal{A}}(\{A \mid A\phi(x_i) \in C\})$  for Borel sets  $C \subseteq \mathbb{R}^m$ . Then  $\lambda_i$  is the probability density function defined by the Borel measure  $\mu_i$ . As a result, we have for each  $i$  that

$$\mathbb{E}_{A \sim \mathcal{A}} \left[ \ell(\psi(\phi_A(x_i)), y(\phi_A(x_i))) \right] = \mathbb{E}_{z \sim \lambda_i} \left[ \ell(\psi(\phi(x_i) + z), y(\phi(x_i) + z)) \right]$$

**Feature robustness and generalization.** With this construction and a distribution  $\Lambda_{\mathcal{A}}$  on the feature space induced by a distribution  $\mathcal{A}$  on feature matrices, we have that

$$\mathcal{E}(f) = \mathcal{E}_{emp}(f, S) + \mathcal{E}_{Rep}^{\phi}(f, S, \Lambda_{\mathcal{A}}) + \mathcal{E}_{\mathcal{F}}^{\phi}(f, S, \mathcal{A}) \quad (5)$$

Here,  $\mathcal{A}$  can be any distribution on feature matrices, which can be chosen suitably to control how well the corresponding mixture of local distributions approximates the true distribution. The third term then measures how robust the model is in expectation over feature changes for  $A \sim \mathcal{A}$ . In particular, if  $\mathcal{E}_{Rep}^{\phi}(f, S, \Lambda_{\mathcal{A}}) \approx 0$ , then  $\mathcal{E}_{gen}(f, S) \approx \mathcal{E}_{\mathcal{F}}^{\phi}(f, S, \mathcal{A})$  and the generalization gap is determined by feature robustness. We end this section by illustrating how distributions on feature matrices induce natural distributions on the feature space. The example will serve in Sec. 5 to deduce a bound on  $\mathcal{E}_{Rep}^{\phi}(f, S, \Lambda_{\mathcal{A}})$  from kernel density estimation.

**Example: Truncated isotropic normal distributions** are induced by a suitable distribution on feature matrices. We consider probability distributions  $\mathcal{K}_{\delta \|\phi(x_i)\|}$  on feature vectors  $z \in \mathbb{R}^m$  in the feature space defined by densities  $k_{\delta \|\phi(x_i)\|}(0, z)$  with smooth rotation-invariant kernels, bounded support and bandwidth  $h$ :

$$k_h(z_i, z) = \frac{1}{h^m} \cdot k\left(\frac{\|z_i - z\|}{h}\right) \cdot \mathbb{1}_{\|z_i - z\| < h} \quad (6)$$

with  $\mathbb{1}_{\|z_i - z\| < h} = 1$  when  $\|z - z_i\| < h$  and 0 otherwise, and such that  $\int_{z \in \mathbb{R}^m} k_h(z_0, z) dz = 1$  for all  $z_0$ . An example for such a kernel is a truncated isotropic normal distribution with variance  $h^2 \sigma^2 I$ ,  $k_h(z_i, z) = \mathcal{N}(z_i, h^2 \sigma^2)(z)$ . The following result states that the densities in (6) can indeed be induced by distributions on feature matrices, which will enable us to connect feature robustness with  $\epsilon$ -representativeness.

**Proposition 2.** *Let  $S^{\phi} = \{\phi(x_i) \mid x_i \in S\}$  be a set of feature vectors in  $\mathbb{R}^m$ . With  $k_h$  defined as in (6), let  $\lambda_i(z) = k_{\delta \|\phi(x_i)\|}(0, z)$  define an  $|S|$ -tuple  $\Lambda_{\delta}$  of densities. Then there exists a distribution  $\mathcal{A}_{\delta}$  on matrices in  $\mathbb{R}^{m \times m}$  of norm less than  $\delta$  such that for each  $i = 1, \dots, |S|$ ,*

$$\mathbb{E}_{A \sim \mathcal{A}_{\delta}} \left[ \ell(\psi(\phi_A(x_i)), y(\phi_A(x_i))) \right] = \mathbb{E}_{\xi \sim \lambda_i} \left[ \ell(\psi(\phi(x_i) + \xi), y(\phi(x_i) + \xi)) \right]$$

The technical proof is deferred to the appendix, but we describe the distribution  $\mathcal{A}_{\delta}$  on matrices for later use: The desired distribution is defined on the set of matrices of the form  $rO$  for a real number  $r$

and an orthogonal matrix  $O$  (i.e.  $OO^T = O^T O = I$ ) as a product measure combining the (unique) Haar measure on the set of orthogonal matrices  $\mathcal{O}(m)$  with a suitable distribution on  $\mathbb{R}$ . The Haar measure on  $\mathcal{O}(m)$  induces the uniform measure on a sphere of radius  $r$  via multiplication with a vector of length  $r$  [17], and we choose a measure on  $\mathbb{R}$  to match the radial change of the kernel  $k_h$ .

## 4 Relative Flatness of the Loss Surface

Flatness is a property of the parameter space quantifying the change in loss under small parameter perturbations, classically measured by the trace of the loss Hessian  $Tr(H)$ , where  $H$  is the matrix containing the partial second derivatives of the empirical risk with respect to all parameters of the model. In order to connect feature robustness (a property of the feature space) to flatness, we present how key equation (1) translates to the empirical risk: For a model  $f(x, \mathbf{w}) = \psi(\mathbf{w}, \phi(x)) = g(\mathbf{w}\phi(x))$  with parameters  $\mathbf{w} \in \mathbb{R}^{d \times m}$  and  $g: \mathbb{R}^d \rightarrow \mathcal{Y}$  a function on a matrix product of parameters  $\mathbf{w}$  and a feature representation  $\phi: \mathcal{X} \rightarrow \mathbb{R}^m$  and any feature matrix  $A \in \mathbb{R}^{m \times m}$  we have that

$$\begin{aligned} \mathcal{E}_{emp}(\mathbf{w} + \mathbf{w}A, \phi(S)) &= \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\psi(\mathbf{w} + \mathbf{w}A, \phi(x_i)), y_i) \\ &= \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\psi(\mathbf{w}, \phi(x_i) + A\phi(x_i)), y_i) = \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\psi(\mathbf{w}, \phi_A(x_i)), y_i) \end{aligned} \quad (7)$$

Subtracting  $\mathcal{E}_{emp}(\mathbf{w}, \phi(S)) = \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\psi(\mathbf{w}, \phi(x_i)), y_i)$ , we can recognize feature robustness (4) on the right side of this equality when labels are constant under perturbations of the features, i.e.  $y(\phi_A(x_i)) = y_i$ . In other words, flatness  $\mathcal{E}_{emp}(\mathbf{w} + \mathbf{v}, \phi(S)) - \mathcal{E}_{emp}(\mathbf{w}, \phi(S))$  describes the performance of a model function on perturbed feature vectors while holding labels constant. We proceed to introduce a novel, but natural, loss Hessian-based flatness measure that approximates feature robustness, given that the underlying data distribution  $\mathcal{D}$  satisfies the assumption of locally constant labels.

With  $\mathbf{w}_s = (w_{s,t})_t \in \mathbb{R}^{1 \times m}$  denoting the  $s$ -th row of the parameter matrix  $\mathbf{w}$ , we let  $H_{s,s'}(\mathbf{w}, \phi(S)) \in \mathbb{R}^{m \times m}$  denote the Hessian matrix containing all partial second derivatives of the empirical risk  $\mathcal{E}_{emp}(\mathbf{w}, \phi(S))$  with respect to weights in rows  $\mathbf{w}_s$  and  $\mathbf{w}_{s'}$ , i.e.

$$H_{s,s'}(\mathbf{w}, \phi(S)) = \left[ \frac{\partial^2 \mathcal{E}_{emp}(\mathbf{w}, \phi(S))}{\partial w_{s,t} \partial w_{s',t'}} \right]_{1 \leq t, t' \leq m}. \quad (8)$$

**Definition 3.** For a model  $f(x, \mathbf{w}) = g(\mathbf{w}\phi(x))$ ,  $\mathbf{w} \in \mathbb{R}^{d \times m}$ , with a twice differentiable function  $g$ , a twice differentiable loss function  $\ell$  and a sample set  $S$ , relative flatness is defined by

$$\kappa_{Tr}^\phi(\mathbf{w}) := \sum_{s,s'=1}^d \langle \mathbf{w}_s, \mathbf{w}_{s'} \rangle \cdot Tr(H_{s,s'}(\mathbf{w}, \phi(S))), \quad (9)$$

where  $Tr$  denote the trace and  $\langle \mathbf{w}_s, \mathbf{w}_{s'} \rangle = \mathbf{w}_s \mathbf{w}_{s'}^T$  the scalar product of two row vectors.

**Properties of relative flatness** (i) Relative flatness simplifies to ridge regression for linear models  $f(x, \mathbf{w}) = \mathbf{w}x \in \mathbb{R}$  ( $\mathcal{X} = \mathbb{R}^d$ ,  $g = id$  and  $\phi = id$ ) and squared loss: To see this, note that for any loss function  $\ell$ , the second derivatives with respect to the parameters  $\mathbf{w} \in \mathbb{R}^d$  computes to  $\frac{\partial^2 \ell}{\partial w_i \partial w_j} = \frac{\partial^2 \ell}{\partial (f(x, \mathbf{w}))^2} x_i x_j$ . For  $\ell(\hat{y}, y) = \|\hat{y} - y\|^2$  the squared loss function,  $\partial^2 \ell / \partial \hat{y}^2 = 2$  and the Hessian is independent of the parameters  $\mathbf{w}$ . In this case,  $\kappa_{Tr}^{id} = c \cdot \|\mathbf{w}\|^2$  with a constant  $c = \sum_{x \in S} 2Tr(xx^T)$ , which is the well-known Tikhonov (ridge) regression penalty.

(ii) Invariance under reparameterization: We consider neural network functions

$$f(x) = \mathbf{w}^L \sigma(\dots \sigma(\mathbf{w}^2 \sigma(\mathbf{w}^1 x + b^1) + b^2) \dots) + b^L \quad (10)$$

of a neural network of  $L$  layers with nonlinear activation function  $\sigma$ . By letting  $\phi^l(x)$  denote the composition of the first  $l-1$  layers, we obtain a decomposition  $f(x, \mathbf{w}^l) = g^l(\mathbf{w}^l \phi^l(x))$  of the network. Using (9) we obtain a relative flatness measure  $\kappa_{Tr}^l(\mathbf{w})$  for the chosen layer.

For a well-defined Hessian of the loss function, we require the network function to be twice differentiable. With the usual adjustments (equations only hold almost everywhere in parameter space), we can also consider neural networks with ReLU activation functions. In this case, Dinh et al. [5] noted that the network function—and with it the generalization performance—remains unchanged under linear reparameterization, i.e., multiplying layer  $l$  with  $\alpha > 0$  and dividing layer  $k \neq l$  by  $\alpha$ , but common measures of the loss Hessian change. Our measure fixes this issue in relating flatness to generalization since the change of the loss Hessian is compensated by multiplication with the scalar products of weight matrices and is therefore invariant under layer-wise reparameterizations [cf. 26]. It is also invariant to neuron-wise reparameterizations, i.e., multiplying all incoming weights into a neuron by a positive number  $\alpha$  and dividing all outgoing weights by  $\alpha$  [23], except for neuron-wise reparameterizations of the feature layer  $\phi^l$ . Using a simple preprocessing step (a neuron-wise reparameterization with the variance over the sample), our proposed measure becomes independent of all neuron-wise reparameterizations.

**Theorem 4.** *Let  $\sigma_i$  denote the variance of the  $i$ -th coordinate of  $\phi^l(x)$  over samples  $x \in S$  and  $V = \text{diag}(\sigma_1, \dots, \sigma_{n_{l-1}})$ . If the relative flatness measure  $\kappa_{Tr}^l$  is applied to the representation*

$$f(x) = \mathbf{w}^L \sigma(\dots \sigma(\mathbf{w}^l V \sigma(V^{-1} \mathbf{w}^{l-1} \sigma(\dots \sigma(\mathbf{w}^1 x + b^1)) \dots) + V^{-1} b^{l-1}) + b^l) \dots) + b^L$$

*then  $\kappa_{Tr}^l$  is invariant under all neuron-wise (and layer-wise) reparameterizations*

We now connect flatness with feature robustness: Relative flatness approximates feature robustness for a model at a local minimum of the empirical risk, when labels are approximately constant in neighborhoods of the training samples  $(\phi(x), y) \in \phi(S)$  in feature space.

**Theorem 5.** *Consider a model  $f(x, \mathbf{w}) = g(\mathbf{w}\phi(x))$  as above, a loss function  $\ell$  and a sample set  $S$ , and let  $O_m \subset \mathbb{R}^{m \times m}$  denote the set of orthogonal matrices. Let  $\delta$  be a positive (small) real number and  $\mathbf{w} = \omega \in \mathbb{R}^{d \times m}$  denote parameters at a local minimum of the empirical risk on a sample set  $S$ . If the labels satisfy that  $y(\phi_{\delta A}(x_i)) = y(\phi(x_i)) = y_i$  for all  $(x_i, y_i) \in S$  and all  $\|A\| \leq 1$ , then  $f(x, \omega)$  is  $((\delta, S, O_m), \epsilon)$ -feature robust on average over  $O_m$  for  $\epsilon = \frac{\delta^2}{2m} \kappa_{Tr}^\phi(\omega) + \mathcal{O}(\delta^3)$ .*

Applying the theorem to Eq. 5 implies that if the data is representative, i.e.,  $\mathcal{E}_{Rep}^\phi(f, S, \Lambda_{\mathcal{A}_\delta}) \approx 0$  for the distribution  $\mathcal{A}_\delta$  of Prop. 2, then  $\mathcal{E}_{gen}(f(\cdot, \omega), S) \lesssim \frac{\delta^2}{2m} \kappa_{Tr}^\phi(\omega) + \mathcal{O}(\delta^3)$ . The assumption on locally constant labels in Thm. 5 can be relaxed to approximately locally constant labels without unraveling the theoretical connection between flatness and feature robustness. Appendix B investigates consequences from even dropping the assumption of approximately locally constant labels.

## 5 Flatness and Generalization

Combining the results from sections 2–4, we connect flatness to the generalization gap when the distribution can be represented by smooth probability densities on a feature space with approximately locally constant labels. By approximately locally constant labels we mean that, for small  $\delta$ , the loss in  $\delta \|\phi(x_i)\|$ -neighborhoods around the feature vector of a training sample  $x_i$  is approximated (on average over all training samples) by the loss for constant label  $y(x_i)$  on these neighborhoods. This and the following theorem connecting flatness and generalization are made precise in Appendix D.4.

**Theorem 6 (informal).** *Consider a model  $f(x, \mathbf{w}) = g(\mathbf{w}\phi(x))$  as above, a loss function  $\ell$  and a sample set  $S$ , let  $m$  denote the dimension of the feature space defined by  $\phi$  and let  $\delta$  be a positive (small) real number. Let  $\omega$  denote a local minimizer of the empirical risk on a sample set  $S$ . If the distribution  $\mathcal{D}$  has a smooth density  $p_{\mathcal{D}}^\phi$  on the feature space  $\mathbb{R}^m$  with approximately locally constant labels around the points  $x \in S$ , then it holds with probability  $1 - \Delta$  over sample sets  $S$  that*

$$\mathcal{E}_{gen}(f(\cdot, \omega), S) \lesssim |S|^{-\frac{2}{4+m}} \left( \frac{\kappa_{Tr}^\phi(\omega)}{2m} + C_1(p_{\mathcal{D}}^\phi, L) + \frac{C_2(p_{\mathcal{D}}^\phi, L)}{\sqrt{\Delta}} \right)$$

*up to higher orders in  $|S|^{-1}$  for constants  $C_1, C_2$  that depend only on the distribution in feature space  $p_{\mathcal{D}}^\phi$  induced by  $\phi$ , the chosen  $|S|$ -tuple  $\Lambda_\delta$  and the maximal loss  $L$ .*

To prove Theorem 6 we bound both  $\epsilon$ -representativeness and feature robustness in Eq. 5. For that, the main idea is that the family of distributions considered in Proposition 2 has three key properties: (i) it

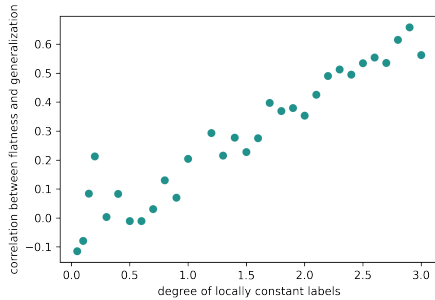


Figure 3: The correlation between flatness and generalization increases with the degree of locally constant labels.

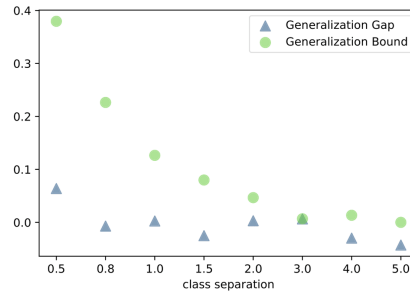


Figure 4: Approximation of representativeness via KDE together with relative flatness leads to a tight generalization bound.

provides an explicit link between the distributions on feature matrices  $\mathcal{A}_\delta$  used in feature robustness and the family of distributions  $\Lambda_\delta$  of  $\epsilon$ -representativeness (Proposition 2) (ii) it allows us to bound feature robustness using Thm. 5; and (iii) it is simple enough that it allows us to use standard results of kernel density estimation (KDE) to bound representativeness.

Our bound suffers from the curse of dimensionality, but for the chosen feature space instead of the (usually much larger) input space. The dependence on the dimension is a result of using KDE uniformly over all distributions satisfying mild regularity assumptions. In practice, for a given distribution and sample set  $S$ , representativeness can be much smaller, which we showcase in a toy example in Sec. 6. In the so-called interpolation regime, where datasets with arbitrarily randomized labels can be fit by the model class, the obtained convergence rate is consistent with the no free lunch theorem and the convergence rate derived by Belkin et al. [2] for an interpolation technique using nearest neighbors.

A combination of our approach with prior assumptions on the hypotheses or the algorithm in accordance to statistical learning theory could potentially achieve faster convergence rate. Our herein presented theory is instead based solely on interpolation and aims to understand the role of flatness (a local property) in generalization: If the data is representative in feature layers and if the distribution can be approximated by locally constant labels in these layers, then flatness of the empirical risk surface approximates the generalization gap. Conversely, Equation 7 shows that flatness measures the performance under perturbed features only when labels are kept constant. As a result, we offer an explanation for the often observed correlation between flatness and generalization: Real-world data distributions for classification are benign in the sense that small perturbations in feature layers do not change the target class, i.e., they can be approximated by locally constant labels. (Note that the definition of adversarial examples hinges on this assumption of locally constant labels.) In that case, feature robustness is approximated by flatness of the loss surface. If the given data and its feature representation are further  $\epsilon$ -representative for small  $\epsilon \approx 0$ , then flatness becomes the main contributor to the generalization gap leading to their noisy, but steady, correlation.

## 6 Empirical Validation

We empirically validate the assumptions and consequences of the theoretical results derived above <sup>2</sup>. For that, we first show on a synthetic example that the empirical correlation between flatness and generalization decreases if labels are not locally constant, up to a point when they are not correlated anymore. We then show that the novel relative flatness measure correlates strongly with generalization, also in the presence of reparameterizations. Finally, we show in a synthetic experiment that while representativeness cannot be computed without knowing the true data distribution, it can in practice be approximated. This approximation—although technically not a bound anymore—tightly bounds the generalization gap. Synthetic data distributions for binary classification are generated by sampling 4 Gaussian distributions in feature space (two for each class) with a given distance between their means (class separation). We then sample a dataset in feature space  $S^\phi$ , train a linear classifier

<sup>2</sup> Code is available at <https://github.com/kampmichael/relativeFlatnessGeneralization>.



$\psi$  on the sample, randomly draw the weights of a 4-layer MLP  $\phi$ , and generate the input data as  $S = (\phi^{-1}(S_x^\phi), S_y^\phi)$ . This yields a dataset  $S$  and a model  $f = \phi \circ \psi$  such that  $\phi(S)$  has a given class separation. Details on the experiments are provided in Appdx. C.

**Locally constant labels:** To validate the necessity of locally constant labels, we measure the correlation between the proposed relative flatness measure and the generalization gap for varying degrees of locally constant labels, as measured by the class separation on the synthetic datasets. For each chosen class separation, we sample 100 random datasets of size 500 on which we measure relative flatness and the generalization gap. Fig. 3 shows the average correlation for different degrees of locally constant labels, showing that the higher the degree, the more correlated flatness is with generalization. If labels are not locally constant, flatness does not correlate with generalization.

### Approximating representativeness:

While representativeness cannot be calculated without knowing the data distribution, it can be approximated from the training sample  $S$  by the error of a density estimation on that sample. For that, we use multiple random splits of  $S$  into a training set  $S_{\text{train}}$  and a test set  $S_{\text{test}}$ , train a kernel density estimation on  $S_{\text{train}}$  and measure its error on  $S_{\text{test}}$ . Again, details can be found in Appx. C. The lower the class separation of the synthetic datasets, the harder the learning problem and the less representative a random sample will be. For each sample and its distribution, we compute the generalization gap and the approximation to the generalization bound. The results in Fig. 4 show that the approximated generalization bound tightly bounds the generalization error (note that this approximation is technically not a bound anymore). Moreover, as expected, the bound decreases the easier the learning problems become.

**Relative flatness correlates with generalization:** We validate the correlation of relative flatness to the generalization gap in practice by measuring it for 110 different local minima—achieved via different learning setups, such as initialization, learning rate, batch size, and optimization algorithm—of LeNet5 [19] on CIFAR10 [18]. We compare this correlation to the classical Hessian-based flatness measures using the trace of the loss-Hessian, the Fisher-Rao norm [21], the PACBayes flatness measure that performed best in the extensive study of Jiang et al. [14] and the  $L_2$ -norm of the weights. The results in Fig. 5 show that indeed relative flatness has higher correlation than all the competing measures. Of these measures, only the Fisher-Rao norm is reparameterization invariant but shows the weakest correlation in the experiment. In Appdx C we show how reparameterizations of the network significantly reduce the correlation for non-reparameterization invariant measures.

## 7 Discussion and Conclusion

Contributing to the trustworthiness of machine learning, this paper provides a rigorous connection between flatness and generalization. As to be expected for a local property, our association between flatness and generalization requires the samples and its representation in feature layers to be representative for the target distribution. But our derivation uncovers a second, usually overlooked condition. Flatness of the loss surface measures the performance of a model close to training points when labels are kept locally constant. If a data distribution violates this, then flatness cannot be a good indicator for generalization.

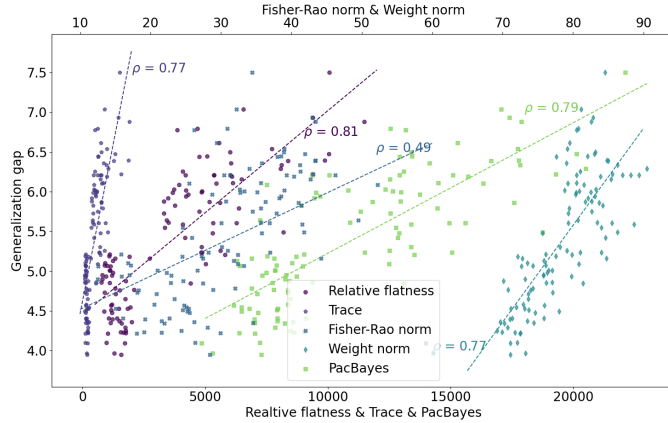


Figure 5: The generalization gap for various local minima correlates stronger with relative flatness than standard flatness, Fisher-Rao norm, PacBayes based measure and weights norm (points corresp. to local minima).

Whenever we consider feature representations other than the input features, the derivation of our results makes one strong assumption: the existence of a target output function  $y(\phi(x))$  on the feature space  $\phi(\mathcal{X})$ . By moving assumptions on the distribution from the input space to the feature space, we achieve a bound based on interpolation that depends on the dimension of the feature layer instead of the input space. Hence, we assume that the feature representation is reasonable and does not lose information that is necessary for predicting the output. To achieve faster convergence rates independent of any involved dimensions, future work could aim to combine our approach of interpolation with a prior-based approach of statistical learning theory.

Our measure of relative flatness may still be improved in future work. Better estimates for the generalization gap are possible by improving the representativeness of local distributions in two ways: The support shape of the local distributions can be improved and their volume-parameter  $\delta$  can be optimally chosen. Both improvements will affect the derivation of the measure of relative flatness as an estimation of feature robustness for the corresponding distributions on feature matrices. Whereas different support shapes change the trace to a weighted average of the Hessian eigenvalues, the volume parameter can provide a correcting scaling factor. Both approaches seem promising to us, as our relative measure from Definition 3 already outperforms the competing measures of flatness in our empirical validation.

## Acknowledgements

Cristian Sminchescu was supported by the European Research Council Consolidator grant SEED, CNCS-UEFISCDI (PN-III-P4-ID-PCE-2016-0535, PN-III-P4-ID-PCCF-2016-0180), the EU Horizon 2020 grant DE-ENIGMA (688835), and SSF.

Mario Boley was supported by the Australian Research Council (under DP210100045).

We would like to thank Julia Rosenzweig, Dorina Weichert, Jilles Vreeken, Thomas Gärtner, Asja Fischer, Tatjana Turova and Alexandru Aleman for the great discussions.

## References

- [1] Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- [2] Mikhail Belkin, Daniel J Hsu, and Partha Mitra. Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. In *Advances in Neural Information Processing Systems*, pages 2300–2311, 2018.
- [3] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [4] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer T. Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. In *Proceedings of the International Conference of Learning Representations*, 2017.
- [5] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1019–1028. JMLR. org, 2017.
- [6] Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *AAAI*, 2017.
- [7] Gintare Karolina Dziugaite and Daniel M Roy. Data-dependent pac-bayes priors via differential privacy. In *Advances in Neural Information Processing Systems*, pages 8430–8441, 2018.
- [8] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *Proceedings of the International Conference on Learning Representations*, 2021.

- [9] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 249–256. PMLR, 2010.
- [10] Sepp Hochreiter and Jürgen Schmidhuber. Simplifying neural nets by discovering flat minima. In *Advances in Neural Information Processing Systems*, pages 529–536, 1995.
- [11] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural Computation*, 9(1):1–42, 1997.
- [12] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *34th Conference on Uncertainty in Artificial Intelligence*, 2018.
- [13] Stanisław Jastrzębski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in sgd. *arXiv preprint arXiv:1711.04623*, 2017.
- [14] Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. In *Proceedings of the International Conference on Learning Representations*, 2020.
- [15] MC Jones, IJ McKay, and T-C Hu. Variable location and scale kernel density estimation. *Annals of the Institute of Statistical Mathematics*, 46(3):521–535, 1994.
- [16] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *Proceedings of the International Conference on Learning Representations*, 2017.
- [17] Steven G. Krantz and Harold R. Parks. *Geometric integration theory*. Springer Science and Business Media, 2008.
- [18] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [19] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. Technical report, AT&T Labs, 2010.
- [20] Yann LeCun, Bernhard E Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne E Hubbard, and Lawrence D Jackel. Handwritten digit recognition with a back-propagation network. In *Advances in Neural Information Processing Systems*, pages 396–404, 1990.
- [21] Tengyuan Liang, Tomaso Poggio, Alexander Rakhlin, and James Stokes. Fisher-rao metric, geometry, and complexity of neural networks. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- [22] Vaishnavh Nagarajan and J Zico Kolter. Uniform convergence may be unable to explain generalization in deep learning. In *Advances in Neural Information Processing Systems*, pages 11611–11622, 2019.
- [23] Behnam Neyshabur, Ruslan Salakhutdinov, and Nathan Srebro. Path-sgd: Path-normalized optimization in deep neural networks. In *Advances in Neural Information Processing Systems*, volume 28, 2015.
- [24] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *Workshop contribution at the International Conference on Learning Representations*, 2015.
- [25] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pages 1376–1401, 2015.
- [26] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, pages 5947–5956, 2017.

- [27] Roman Novak, Yasaman Bahri, Daniel A Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. Sensitivity and generalization in neural networks: an empirical study. In *Proceedings of the International Conference on Learning Representations*, 2018.
- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary 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. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035, 2019.
- [29] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [30] Henning Petzka, Linara Adilova, Michael Kamp, and Cristian Sminchisescu. A reparameterization-invariant flatness measure for deep neural networks. In *Workshop on Science meets Engineering of Deep Learning at NeurIPS*, 2019.
- [31] Akshay Rangamani, Nam H. Nguyen, Abhishek Kumar, Dzung T. Phan, Sang H. Chin, and Trac D. Tran. A scale invariant flatness measure for deep network minima. *arXiv preprint arXiv:1902.02434*, 2019.
- [32] Levent Sagun, Utku Evci, V Ugur Guney, Yann Dauphin, and Leon Bottou. Empirical analysis of the hessian of over-parametrized neural networks. *Workshop contribution at the International Conference on Learning Representations*, 2019.
- [33] Bernard W. Silverman. *Density estimation for statistics and data analysis*. Monographs on Statistics and Applied Probability. Chapman and Hall, 1986.
- [34] Xu Sun, Zhiyuan Zhang, Xuancheng Ren, Ruixuan Luo, and Liangyou Li. Exploring the vulnerability of deep neural networks: A study of parameter corruption. *arXiv preprint arXiv:2006.05620*, 2020.
- [35] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [36] Andrey Nikolayevich Tikhonov, A. Goncharsky, V. V. Stepanov, and Anatolij Grigorevic Yagola. *Numerical Methods for the Solution of Ill-Posed Problems*. Mathematics and Its Applications. Springer Netherlands, 1995.
- [37] Yusuke Tsuzuku, Issei Sato, and Masashi Sugiyama. Normalized flat minima: Exploring scale invariant definition of flat minima for neural networks using PAC-Bayesian analysis. In *Proceedings of the 37th International Conference on Machine Learning*, pages 9636–9647, 2020.
- [38] Huan Wang, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. Identifying generalization properties in neural networks. *arXiv preprint arXiv:1809.07402*, 2018.
- [39] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. In *Advances in Neural Information Processing Systems*, volume 33, pages 2958–2969, 2020.
- [40] Huan Xu and Shie Mannor. Robustness and generalization. *Machine learning*, 86(3):391–423, 2012.
- [41] Zhewei Yao, Amir Gholami, Qi Lei, Kurt Keutzer, and Michael W Mahoney. Hessian-based analysis of large batch training and robustness to adversaries. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

- [42] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *Proceedings of the International Conference on Learning Representations*, 2017.
- [43] Chiyuan Zhang, Qianli Liao, Alexander Rakhlin, Brando Miranda, Noah Golowich, and Tomaso Poggio. Theory of deep learning iib: Optimization properties of sgd. *arXiv preprint arXiv:1801.02254*, 2018.
- [44] Yaowei Zheng, Richong Zhang, and Yongyi Mao. Regularizing neural networks via adversarial model perturbation. *arXiv preprint arXiv:2010.04925*, 2020.

**Organization of the Appendix** The appendix is organized as follows:

**A – Related work** contains an extended discussion on related work.

**B – The effect of local label changes** discusses consequences for the association of flatness and generalization for general output function  $y(x)$  without the assumption of locally constant labels.

**C – Details on the Empirical Validation** contains a detailed description of the experiments.

**D – Proofs** contains the full proofs to all statements. In detail:

D.1: Proposition 2,

D.2: Theorem 4,

D.3: Theorem 5,

D.4: Theorem 6.

**E – Relative flatness for a uniform bound over general distributions on feature matrices** defines a variant of relative flatness that uniformly bounds feature robustness over all feature matrices.

## A Related Work

It has long been observed that algorithms searching for flat minima of the loss curve lead to better generalization [10, 11]. More recently, an association between flatness and low generalization error has also been validated empirically in deep learning [16, 27, 38]. Here, flatness is measured by the Hessian of the empirical loss evaluated at the model at hand. Indeed, in their recent extensive empirical study of generalization measures, Jiang et al. [14] found that measures based on flatness have the highest correlation with generalization.

For models trained with stochastic gradient descent (SGD), this could present a (partial) explanation for their generalization performance, since the convergence of SGD can be connected to flat local minima by studying SGD as an approximation of a stochastic differential equation [13, 43]. However, while large and small batch methods appear to converge in different basins of attraction, the basins can be connected by a path of low loss, i.e., they can actually converge into the same basin [32]. Moreover, as Dinh et al. [5] remarked, classical flatness measures—which are based only on the Hessian of the loss function—cannot theoretically be related to generalization: For deep neural networks with ReLU activation functions, there are linear reparameterizations that leave the network function unchanged (hence, also the generalization performance), but change any measure derived only from the loss Hessian. Novel measures related to flatness have been proposed that are invariant to linear reparameterizations [21, 31, 37]. Rangamani et al. [31] measure flatness in the quotient space of a suitable equivalence relation, and Liang et al. [21] utilize the Fisher-Rao metric, but the theoretical connection of these two measures to generalization is not well-understood. Neyshabur et al. [26] noted that the reparameterization-issue can in general be resolved by balancing a measure of flatness with a norm on the parameters, which is the way that normalized flatness [37], Fisher-Rao metric [21] and our proposed relative flatness become reparameterization-invariant. However, the solution proposed in Neyshabur et al. [26] necessitates data-dependent priors [7] or related approaches, which "adds non-trivial costs to the generalization bounds" [37].

The question arises in which way the loss Hessian and parameter norm should be combined. A simple scaling of the full Hessian with the squared parameter norm does not provide a reparameterization-invariant measure. Doing so for each layer independently and summing up the results provides a measure that is only invariant under layer-wise reparameterizations. Similarly, only considering a single feature layer yields a measure that is layer-wise reparameterization invariant [30]. While the resulting measure can also be analyzed within our framework to obtain a bound on feature robustness, our proposed measure yields a tighter bound and is also invariant under neuron-wise reparameterizations.

Tsuzuku et al. [37] derive a flatness measure that scales an approximation to the loss Hessian by a parameter-dependent term. Their proposed measure correlates well with generalization and is theoretically connected to it via the PAC-Bayesian framework. However, this connection requires the assumption of Gaussian priors and posteriors and is not informative with respect to conditions under which this connection holds. Moreover the measure is impractical, since computing it requires

solving an optimization problem for every layer that can be numerically unstable. (Tsuzuku et al. [37] propose a solution to the numerical instability at the cost of losing the reparameterization-invariance.) Instead, relative flatness can be computed directly and takes only parameters of a specific layer into account—although combining relative flatness of all layers by simple summation is possible.

A series of recent papers studies flatness by minimizing the loss at local perturbations of the parameters considering  $\min_{\mathbf{a}} \mathcal{E}_{emp}(f(\mathbf{w} + \mathbf{a}), S)$  [8, 34, 39, 44]. Regularization techniques enforcing these notions of flatness during training in classification tasks lead to better generalization. These empirical results follow earlier works by Chaudhari et al. [4] and Izmailov et al. [12] that similarly obtained better generalization by enforcing flatter minima. Their observations are well-explained by our theory: Low error at perturbations  $\mathcal{E}_{emp}(f(\mathbf{w} + \mathbf{a}), S)$  lead to good generalization around training samples. This requires that the underlying distribution has (approximately) locally constant labels (using key equation (1)), which is reasonable for the image classification tasks they consider.

Xu and Mannor [40] propose a notion of robustness over a portion of the input space and derive generalization bounds based on it. However, their notion requires the choice of a partitioning of the input space before seeing any samples. Thus, robustness over the partition can be hard to estimate for a model that depends on a sample set  $S$ . Our notion of feature robustness is measured around a given sample set and thus does not require a uniform data-independent partitioning. Such a sample-dependent notion of robustness is necessary to connect it to the flatness of the loss surface, since flatness is a local property around training points.

Novak et al. [27] find that robustness to input perturbation as measured by the input-output Jacobian correlates well with generalization on classification tasks. This is in line with our findings applied to  $\phi = id_{\mathcal{X}}$  chosen as the identity (for neural networks this means considering the input layer as features): it follows from Equation 1 that robustness to input perturbations directly relates to flatness. Therefore, these findings give additional empirical evidence to the correlation between flatness and generalization. Yao et al. [41] study the Hessian with respect to the input  $x \in \mathcal{X}$  and also find that robust learning tends to converge to minima where the input-output Hessian has small eigenvalues.

## B The effect of local label changes

For classification tasks with one-hot vectors as labels, the assumption of locally constant labels, i.e., locally constant target output function  $y(x)$ , seems reasonable since we would not expect the class label to change under (infinitesimally) small changes. One could nonetheless consider a smooth output function with values encoding class probabilities for classification, which may change locally around the training points. For regression tasks, the assumption of locally constant output function is rather unrealistic or at the very least restrictive.

Taking the term defining feature robustness (4) as a starting point, we investigate its connection to flatness when the output function  $y(x)$  is a smooth function. In the usual setting of machine learning, this information is unknown. We will show that label changes can contribute stronger to the loss in neighborhoods around training samples than (relative) flatness.

To investigate the label dependence, we use the same trick as in (7) to transfer perturbations in the input  $x$  to perturbations in parameter space  $\mathbf{w}$ . To simplify the analysis, we apply feature robustness to the input space (i.e., we only consider  $\phi = id_{\mathcal{X}}$  here). Let  $f(x, \mathbf{w}) = \psi(\mathbf{w}x)$  be a model composed of a matrix multiplication of  $x$  with  $\mathbf{w}$  and a differentiable predictor function  $\psi$ .

$$\begin{aligned}\mathcal{E}_{\bar{F}}(f, S, A) &= \frac{1}{n} \sum_{i=1}^n (\ell(f(x_i + \delta Ax_i, \mathbf{w}), y[x_i + \delta Ax_i]) - \ell(f(x_i, \mathbf{w}), y_i)) \\ &= \frac{1}{n} \sum_{i=1}^n (\ell(f(x_i, \mathbf{w} + \delta \mathbf{w}A), y[x_i + \delta Ax_i]) - \ell(f(x_i, \mathbf{w}), y_i))\end{aligned}$$

Defining a function

$$\gamma_i(\delta) = \ell(f(x_i, \mathbf{w} + \delta \mathbf{w}A), y[x_i + \delta Ax_i]), \quad (8)$$

we have that  $\mathcal{E}_{\bar{F}}(f, S, A) = \frac{1}{n} \sum_{i=1}^n \gamma_i(\delta)$ . For each  $\gamma_i$  we use Taylor approximation in  $\delta$ . In the following, we write  $\ell_{\mathbf{w}}(x_i, \mathbf{w}^*, y_i)$  for the first derivative of the loss with changes in  $\mathbf{w}$  at  $x_i, y_i = y[x_i]$  and  $\mathbf{w}^*$ , and we write  $\ell_y(x_i, \mathbf{w}^*, y_i)$  for the first derivative of the loss with changes of the output  $y$  at  $x_i, y_i = y[x_i]$  and  $\mathbf{w}^*$ . Similarly, we consider second derivatives  $\ell_{\mathbf{w}\mathbf{w}}, \ell_{yy}$  and  $\ell_{\mathbf{w}y}$ . Finally, we denote the derivative of  $y(x)$  with respect to  $x$  by  $y_x$  and the second derivative by  $y_{xx}$ . Then,

$$\gamma_i'(0) = \ell_{\mathbf{w}}(x_i, \mathbf{w}^*, y_i) \cdot (\mathbf{w}^*A) + \ell_y(x_i, \mathbf{w}^*, y_i) \cdot (y_x(x_i) \cdot Ax_i) \quad (9)$$

and

$$\begin{aligned}\gamma_i''(0) &= (\mathbf{w}^*A)^T \ell_{\mathbf{w}\mathbf{w}}(x_i, \mathbf{w}^*, y_i) (\mathbf{w}^*A) + (y_x(x_i) \cdot Ax_i)^T \ell_{yy}(x_i, \mathbf{w}^*, y_i) (y_x(x_i) \cdot Ax_i) \quad (10) \\ &+ \sum_{\text{labels } c} \ell_{y_c}(x_i, \mathbf{w}^*, y_i) \cdot (Ax_i)^T (y_c)_{xx}(x_i) \cdot (Ax_i) + 2(y_x(x_i) \cdot Ax_i)^T \ell_{y\mathbf{w}}(x_i, \mathbf{w}^*, y_i) (\mathbf{w}^*A) \quad (11)\end{aligned}$$

At a critical point we have that  $\sum_i \ell_{\mathbf{w}}(x_i, \mathbf{w}^*, y_i) = 0$ , but since we do not know how the target output function  $y(x)$  changes locally, we do not necessarily<sup>3</sup> enforce that  $\sum_i \ell_y(x_i, \mathbf{w}^*, y_i) = 0$  at a local optimum. In that case,  $\mathcal{E}_{\bar{F}}(f, S, A) = \sum_i \ell_y(x_i, \mathbf{w}^*, y_i) \delta + \mathcal{O}(\delta^2)$  has a non-zero term of first order in  $\delta$  and flatness only contributes as a term of order two. Similarly, other terms in (10) can be nonzero, further reducing the influence of relative flatness to a bound on feature robustness.

As an interesting special case, we note that for one-hot encoded labels in classification and letting the output function  $y(x)$  describe a parameter vector of a conditional label-distribution given  $x$ , we have  $y_x(x_i) = 0$  (recall that we suppose  $y(x_i) = y_i$ ) as each vector component is either 1 or 0 and must be a local extreme point ( $y(x)$  cannot contain values larger than 1 or smaller than 0 by assumption),

We leave a detailed investigation of the consequences of label changes as future work, but identify the implicit assumption of locally constant labels in loss Hessian-based flatness measures as a possible limitation: Flatness can only be descriptive if optimal label changes are approximately locally constant. The fact that a strong correlation between flatness and generalization gap has been often observed points to the fact that distributions in practice satisfy this implicit assumption.

<sup>3</sup>This depends on the loss function in use.



## C Details on the Empirical Validation

Here we provide additional details on the empirical evaluation. Jupyter notebooks containing the experiments are available at <https://github.com/kampmichael/relativeFlatnessGeneralization>, ensuring reproducibility, together with an implementation of the relative flatness measure in pytorch [28].

### C.1 Synthetic Experiments

The experiments on *locally constant labels* and *approximating representativeness* use a synthetic sample in feature space. The schema for both experiments is to

1. create a synthetic dataset in feature space  $S^\phi$  and test set  $T^\phi$ ,
2. create a model  $f = \phi \circ \psi$ ,
3. derive input data as  $S = (\phi^{-1}(S_x^\phi), S_y^\phi)$ ,  $T = (\phi^{-1}(T_x^\phi), T_y^\phi)$
4. compute relative flatness (or other measures) of  $f$  on  $S$ ,
5. and estimate its generalization gap by computing the empirical risk of  $f$  on  $S$ , and computing the test error on the test set  $T$  to estimate the risk.

1) To create  $S^\phi$  with a given class separation  $c$ , we randomly sample 4 cluster centroids  $\theta$  from a hypercube in  $\mathbb{R}^6$  and scale them so that their distance is  $c$ . We then sample a random covariance matrix  $\Sigma$  for each cluster and sample points from a Gaussian  $\mathcal{N}(\theta, \Sigma)$ . Furthermore, we create two redundant features that are a random linear combination of the 6 informative features. We obtain labels by assigning two clusters to class 1 and the other two to class  $-1$ .

2) We create the model  $f$  by first training a linear model  $\psi$  on  $S^\phi$  using ridge regression from scikit-learn [29]. We then sample a random 4-layer MLP (with architecture 784-512-128-16-8, tanh activation, and Glorot initialization [9]) that we use as feature extractor  $\phi$ . With this, we obtain the 5-layer MLP  $f = \phi \circ \psi$  by adding an 8 – 2 layer with weights obtained from  $\psi$ .

3) We obtain input data  $S$  by reverse propagation of samples in feature space  $S_x^\phi$  through the 4-layer MLP  $\phi$ . This is an approximation to the inverse feature extractor  $\phi^{-1}$ . For the output of each layer  $z$ , we first compute  $z' = \tanh^{-1}(z)$ , i.e., the inverse of the activation function. We then solve  $Wz + b = x$ , where  $W, b$  are the weights and bias of that layer, and  $x$  is the corresponding input we want to compute. This yields  $S_x = \phi^{-1}(S_x^\phi)$ . Note that this reverse propagation of samples introduces a small error. To keep experiments realistic, we discard  $S^\phi$  after this step and use only the input dataset  $S$  and model  $f$  in our computations.

4) We compute relative flatness as in Def. 3 (an implementation in pytorch is available on github, see above).

5) We compute the empirical risk of  $f$  on  $S$  and estimate the risk on a test set. For the experiments on locally constant labels, generate 5000 samples, use a training set of size 500, a test set of size 4500 (to ensure an accurate estimate of the risk), and repeat the experiment 100 times for each class separation  $c$ . For the experiment on approximating representativeness, we use a sample of size 600 and perform 3-fold cross-validation.

**Locally constant labels:** For classification, labels are locally constant if in a neighborhood around each point the label does not change. They are approximately locally constant, if this holds for most points. By increasing the distance between the means of the Gaussians, we decrease the likelihood of a point within a neighborhood having a different label. For a finite sample, this means that the likelihood of observing two points close by with different labels decreases. Thus, by increasing the class separation parameter, we increase the degree of locally constant labels.

**Approximating representativeness:** A finite random sample as described in 1) has a higher chance of being representative when the means of the Gaussians have a high distance, because each individual Gaussian can be interpolated easily. Of course, the actual representativeness of a sample at hand can vary. Note that this is a very simple form of generating datasets with varying "difficulty". It will be interesting to further explore the impact of the choice of data distribution on (an approximation to) representativeness.

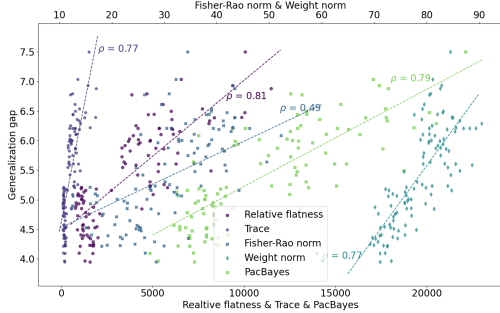


Figure 6: Generalization gap and various flatness measures for 110 local minima as presented in Fig. 5. The generalization gap correlates stronger with relative flatness than standard flatness, Fisher-Rao norm, a PAC-Bayes based measure and the weights norm.

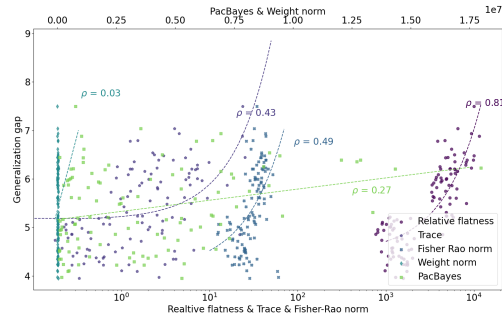


Figure 7: Modifying the local minima in the plot Fig. 6 by reparameterization shows that the proposed relative flatness and the Fisher-Rao norm are invariant to them. It furthermore shows a strong decline in correlation for all other measures.

Experiments on the synthetic datasets are run on a laptop with Intel Core i7 and NVIDIA GeForce GTX 965 M 2 GB GPU. The code of the experiments is provided as a jupyter notebook so that they can be easily reproduced.

## C.2 Relative Flatness Correlates with Generalization

In this experiment, we validate that the proposed relative flatness correlates strongly with generalization in practice. For that, we measure relative flatness (as well as classical flatness measured by the trace of the loss Hessian, the Fisher-Rao norm, a PAC-Bayes based measure <sup>4</sup>, and the weight norm) together with the generalization gap for various local minima.

To obtain model parameters at various local minima, we train networks (LeNet5 [20]) on the CIFAR10 dataset until convergence (measured in terms of achieving a loss of less than 0.1 during an epoch, which has been used as a criteria for convergence in similar experiments [14]) with varying hyperparameters. In accordance to works studying the impact of hyperparameters on generalization [14, 16, 27, 31, 38], we vary learning rate, mini-batch size, initialization, and optimizer. We vary the mini batch size in 64, 128, 256, 512, 1024, and the learning rate in 0.0001, 0.02, 0.05, running 10 randomly initialized training rounds for each setup. We use SGD, ADAM, and RMSProp as optimizers. We only use combinations that lead to convergence. The experiments were conducted on a cluster node with 4 NVIDIA GPU GM200 (GeForce GTX TITAN X). As discussed in Sec. 6, relative flatness has the highest correlation with generalization from all measures we analyzed.

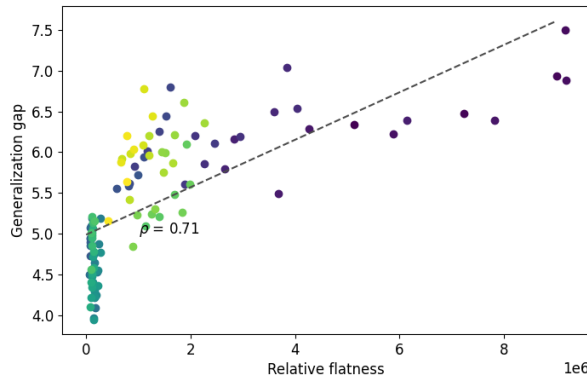


Figure 8: The generalization gap for various local minima correlates with relative flatness measured on the layer different from penultimate layer.

<sup>4</sup>The implementation of PAC-Bayes based flatness measure is taken from <https://github.com/nitarshan/robust-generalization-measures/blob/master/data/generation/measures.py>

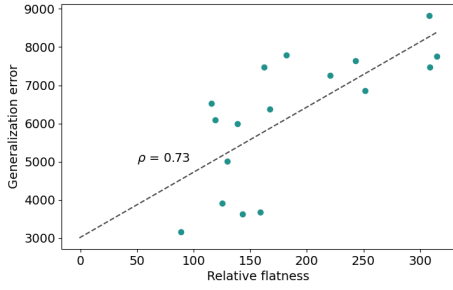


Figure 9: Layer1-based relative flatness for MNIST experiment. Layer1 is the topmost layer.

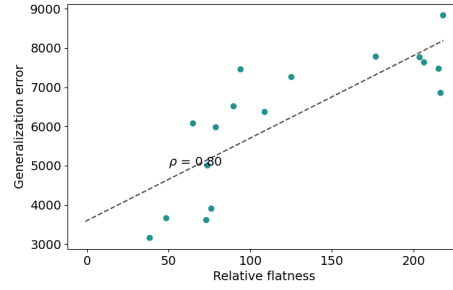


Figure 10: Layer2-based relative flatness for MNIST experiment.

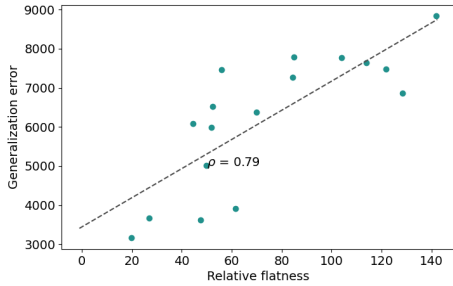


Figure 11: Layer3-based relative flatness for MNIST experiment.

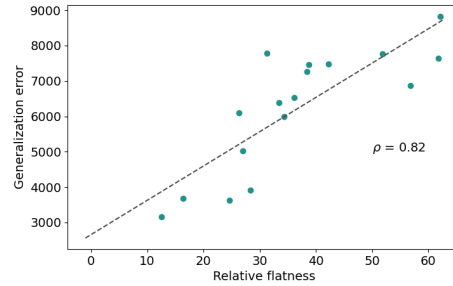


Figure 12: Layer4-based relative flatness for MNIST experiment. Layer4 is penultimate.

To study the effect of reparameterization, we apply layer-wise reparameterizations on the trained network using random factors in the interval  $[5, 25]$  which yields a set of novel local minima. The results in Fig. 7 show that both our proposed relative flatness and the Fisher-Rao norm are invariant to these reparameterization. For all other measures, the correlation with generalization declines substantially. The same would hold for neuron-wise reparameterizations, since both relative flatness and the Fisher-Rao norm are also neuron-wise reparameterization invariant. Relative flatness and the Fisher-Rao norm are also invariant under neuron-wise reparameterizations, which could be used to further break the correlation for the other measures. For future work it would be interesting to investigate further symmetries in neural networks and the impact of reparameterizations along these symmetries on flatness measures.

In addition to the calculation of the relative flatness using the feature space of the penultimate layer, we also performed calculations for another fully-connected layer in the network. The resulting correlation can be seen in Fig. 8. It keeps the high correlation value, but due to less optimal feature space we observe smaller number, than in the previous calculation. Nevertheless, it demonstrates that any  $\phi$ - $\psi$  separation allows to compute relative flatness.

For checking deeper the viability of the relative flatness computed in different feature representations, we ran a similar experiment with a fully-connected  $(784 - 50 - 50 - 50 - 30 - 10)$  neural network trained on MNIST dataset (Fig. 9, 10, 11, 12). We varied parameters of the optimization (batch size in  $1000 - 2000 - 4000 - 8000$  and learning rate in  $0.02 - 0.04 - 0.08 - 0.16$  in order to keep the ratio between them constant) and trained each network with SGD for 500 epochs. Only the networks that achieved training loss lower than 0.07 are used for the plots. The observed correlation with generalization gap is high for each of the representations in the network.

## D Proofs

### D.1 Proof of Proposition 2

*Proof.* Let  $\mathcal{K}_h$  denote probability distribution defined by a rotational-invariant kernel  $k_h$  as in (6) with  $k_h(0, z) = \frac{1}{h^m} \cdot k\left(\frac{\|z\|}{h}\right) \cdot \mathbf{1}_{\|z\| < h}$  and let  $\lambda_i(z) = k_{\delta\|\phi(x_i)\|}(0, z)$ . Let  $\mathcal{L}$  denote a continuous function on  $\mathbb{R}^m$  and  $\mathcal{O}_m$  the set of orthogonal matrices in  $\mathbb{R}^{m \times m}$ . We show that there exists a probability measure  $\kappa$  on a set  $M_\delta$  of matrices of norm smaller than  $\delta$ , defining a probability distribution  $\mathcal{A}_\delta$ , and a probability measure  $\omega$  on the product space  $(0, \delta] \times \mathcal{O}_m$  such that for each  $z \in \mathbb{R}^m \setminus \{0\}$ :

$$\mathbb{E}_{A \sim \mathcal{A}_\delta} [\mathcal{L}(z + Az)] = \mathbb{E}_{(r, O) \sim \omega} [\mathcal{L}(z + rOz)] = \mathbb{E}_{\zeta \sim \mathcal{K}_{\delta\|z\|}} [\mathcal{L}(z + \zeta)] \quad (12)$$

Applying this result for each  $i = 1, \dots, |S|$  to  $\mathcal{L}_i(z) = \ell(\psi(\mathbf{w}, z), y_i[z])$  at  $z = \phi(x_i)$  completes the proof. For all the standard measure-theoretic concepts used in the proof, we refer the reader to [17].

Fix some  $\zeta_0$  in  $\mathbb{R}^m$  with  $\|\zeta_0\| = 1$ . We consider the Haar measure  $\mu$  on the set of orthogonal matrices  $\mathcal{O}_m$ . By [17, Proposition 3.2.1] and the change of variables formula, we have for each  $r \in (0, \delta]$

$$\int_{O \in \mathcal{O}_m} \mathcal{L}(z + r\|z\|O\zeta_0) d\mu(O) = \frac{1}{\text{Vol}(S^{m-1})} \int_{\xi \in S^{m-1}} \mathcal{L}(z + r\|z\|\xi) d\xi$$

where  $S^{m-1}$  is the  $(m-1)$ -sphere. We multiply both sides by  $\frac{\text{Vol}(S^{m-1})}{\delta^m} k\left(\frac{r}{\delta}\right) r^{m-1}$ , integrate over  $r \in (0, \delta]$  to obtain

$$\begin{aligned} & \frac{\text{Vol}(S^{m-1})}{\delta^m} \int_{r=0}^{\delta} \int_{O \in \mathcal{O}_m} \mathcal{L}(z + r\|z\|O\zeta_0) k\left(\frac{r}{\delta}\right) r^{m-1} dr d\mu(O) \\ &= \frac{1}{\delta^m} \int_{r=0}^{\delta} \int_{\xi \in S^{m-1}} \mathcal{L}(z + r\|z\|\xi) k\left(\frac{r}{\delta}\right) r^{m-1} dr d\xi \\ &= \frac{1}{\delta^m} \int_{\|\zeta\| \leq \delta} \mathcal{L}(z + \|\zeta\|\zeta) k\left(\frac{\|\zeta\|}{\delta}\right) d\zeta \\ &= \int_{\|\zeta\| \leq \delta\|z\|} \mathcal{L}(z + \zeta) \frac{1}{(\delta\|z\|)^m} k\left(\frac{\|\zeta\|}{\delta\|z\|}\right) d\zeta \end{aligned}$$

Introducing the product measure  $\omega := \frac{\text{Vol}(S^{m-1})}{\delta^m} \cdot (k\left(\frac{r}{\delta}\right) r^{m-1} dr \times \mu)$  on  $(0, \delta] \times \mathcal{O}_m$ , this implies that

$$\mathbb{E}_{(r, O) \sim \omega} [\mathcal{L}(z + r\|z\|O\zeta_0)] = \mathbb{E}_{\zeta \sim \mathcal{K}_{\delta\|z\|}} [\mathcal{L}(z + \zeta)] \quad (13)$$

The measure  $\omega$  can be pushed forward to a measure on matrices of norm  $\|A\| \leq \delta$ . For this, consider the homeomorphism

$$H : (0, \delta] \times \mathcal{O}_m \rightarrow \{rO \mid r \in (0, \delta], O \in \mathcal{O}_m\} =: M_\delta \subseteq \{A \in \mathbb{R}^{n \times n} \mid \|A\| \leq \delta\}$$

given by  $H(r, O) = rO$ . We use the inverse of  $H$  to push forward the measure  $\omega$  to a measure  $\kappa$  on  $M_\delta$  and obtain from (13) that

$$\mathbb{E}_{A \sim (M_\delta, \kappa)} [\mathcal{L}(z + \|z\|A\zeta_0)] = \mathbb{E}_{\zeta \sim \mathcal{K}_{\delta\|z\|}} [\mathcal{L}(z + \zeta)]$$

Finally, there exists an orthogonal matrix  $O$  such that  $O\|z\|\zeta_0 = z$ . Since  $\kappa(A) = \kappa(AO^{-1})$  by definition of  $\kappa$  and since  $M_\delta O = M_\delta$ , we get for any  $z$  that

$$\begin{aligned} \mathbb{E}_{\zeta \sim \mathcal{K}_{\delta\|z\|}} [\mathcal{L}(z + \zeta)] &= \mathbb{E}_{A \sim (M_\delta, \kappa)} [\mathcal{L}(z + A\|z\|\zeta_0)] \\ &= \mathbb{E}_{A \sim (M_\delta O^{-1}, \kappa)} [\mathcal{L}(z + AO\|z\|\zeta_0)] \\ &= \mathbb{E}_{A \sim (M_\delta, \kappa)} [\mathcal{L}(z + Az)] \end{aligned}$$

Hence, the probability distribution  $\mathcal{A}_\delta$  on matrices with norm bounded by  $\delta$  defined by the probability measure  $\kappa$  with support on  $M_\delta$ , and the space  $(0, \delta] \times \mathcal{O}_m$  equipped with  $\omega = \frac{\text{Vol}(S^{m-1})}{\delta^m} \cdot (k\left(\frac{r}{\delta}\right) r^{m-1} dr \times \mu)$  give the desired probability distributions satisfying (12).  $\square$

## D.2 Proof of Theorem 4

We rephrase Theorem 4 split into Theorem 7 and a subsequent corollary that specify the reparameterizations under consideration. Let  $f = f(\mathbf{w}^1, b^1, \mathbf{w}^2, b^2, \dots, \mathbf{w}^L, b^L)$  denote a ReLU network function parameterized by parameters  $\mathbf{w}^k = w_{s,t}^k$  and bias  $b^k = b_s^k$  of the  $k$ -th layer given by

$$f(x) = \mathbf{w}^L \sigma(\dots \sigma(\mathbf{w}^l \sigma(\mathbf{w}^{l-1} \sigma(\dots \sigma(\mathbf{w}^1 x + b^1)) \dots) + b^{l-1}) + b^l) \dots) + b^L.$$

Recall that we let  $\phi^l(x)$  denote the composition of the first  $(l - 1)$  layers so that we obtain a decomposition  $f(x, \mathbf{w}^l) = g^l(\mathbf{w}^l \phi^l(x))$  of the network. Using (9) we obtain a relative flatness measure  $\kappa_{Tr}^l(\mathbf{w})$  for the chosen layer.

A layer-wise reparameterization multiplies all weights in a layer  $l$  with a positive number  $\lambda$  and divides the weights of another layer  $l' \neq l$  by the same  $\lambda$ . Due to the positive homogeneity of the ReLU activation, this does not change the network function. By a neuron-wise reparameterization, we mean the operation that multiplies all weights into a neuron by some positive  $\lambda$  and divides all outgoing weights of the same neuron by  $\lambda$ . Again, the positive homogeneity of the activation function implies that this operation does not change the network function. A layer-wise reparameterization is simply the parallel application of neuron-wise reparameterization for all neurons of one layer with the same reparameterization parameter  $\lambda > 0$ .

**Theorem 7.** *Let  $f = f(\mathbf{w}^1, b^1, \mathbf{w}^2, b^2, \dots, \mathbf{w}^L, b^L)$  denote a neural network function parameterized by parameters  $\mathbf{w}^k = w_{s,t}^k$  and bias  $b^k = b_s^k$  of the  $k$ -th layer. Suppose there are positive numbers  $\lambda_{s,t}^k$  such that the parameters  $\mathbf{w}_\lambda^k, b_\lambda^k$ , obtained from multiplying  $w_{s,t}^k$  at matrix position  $(s, t)$  in layer  $k$  by  $\lambda_{s,t}^k$  and  $b_s^k$  by  $\lambda_{(s,0)}^k$ , satisfy that  $f(\mathbf{w}^1, b^1, \mathbf{w}^2, b^2, \dots, \mathbf{w}^L, b^L) = f(\mathbf{w}_\lambda^1, b_\lambda^1, \mathbf{w}_\lambda^2, b_\lambda^2, \dots, \mathbf{w}_\lambda^L, b_\lambda^L)$ . If for the layer with index  $l$  it holds that  $\lambda_{(s,t)}^l = \lambda_{(s,t')}^l$  for each  $s, t$  and  $t'$ , then  $\kappa_{Tr}^l(\mathbf{w}) = \kappa_{Tr}^l(\mathbf{w}_\lambda)$  for the notion of relative flatness from Definition 3.*

**Corollary 8.** *Let  $\sigma_i$  denote the variance of the  $i$ -th coordinate of  $\phi^l(x)$  over samples  $x \in S$  and  $V = \text{diag}(\sigma_1, \dots, \sigma_{n_{l-1}})$ . If the relative flatness measure  $\kappa_{Tr}^l$  is applied to the representation  $f = f(\mathbf{w}^1, b^1, \dots, V^{-1} \mathbf{w}^{l-1}, V^{-1} b^{l-1}, \mathbf{w}^l V, b^l, \mathbf{w}^{l+1}, b^{l+1}, \dots, \mathbf{w}^L, b^L)$ , i.e.,*

$$f(x) = \mathbf{w}^L \sigma(\dots \sigma(\mathbf{w}^l V \sigma(V^{-1} \mathbf{w}^{l-1} \sigma(\dots \sigma(\mathbf{w}^1 x + b^1)) \dots) + V^{-1} b^{l-1}) + b^l) \dots) + b^L,$$

then  $\kappa_{Tr}^l$  is invariant under all neuron-wise (and layer-wise) reparameterizations

*Proof.* We are given a neural network function  $f(x; \mathbf{w}^1, b^1, \dots, \mathbf{w}^L, b^L)$  parameterized by parameters  $\mathbf{w}^k$  and bias terms  $b^k$  of the  $k$ -th layer and positive numbers  $\lambda_{(s,t)}^1, \dots, \lambda_{(s,t)}^L$  such that the parameters  $\mathbf{w}_\lambda^k$  obtained from multiplying weight  $w_{(s,t)}^k$  at matrix position  $(s, t)$  in layer  $k$  by  $\lambda_{(s,t)}^k$  and  $b_s^k$  by  $\lambda_{(s,0)}^k$  satisfies that

$$f(x; \mathbf{w}^1, b^1, \mathbf{w}^2, b^2, \dots, \mathbf{w}^L, b^L) = f(x; \mathbf{w}_\lambda^1, b_\lambda^1, \mathbf{w}_\lambda^2, b_\lambda^2, \dots, \mathbf{w}_\lambda^L, b_\lambda^L)$$

for all  $\mathbf{w}^k, b^k$  and all  $x$ .

For fixed layer  $l$ , we denote the  $s$ -th row of  $\mathbf{w}^l$  by  $\mathbf{w}_s^l$  before reparameterization, and we denote the  $s$ -th row of  $\mathbf{w}_\lambda^l$  by  $\mathbf{w}_{\lambda s}^l$  after reparameterization. For simplicity of the notation, we will collect all bias terms in terms  $\mathbf{b}, \mathbf{b}_\lambda$  before and after reparameterization respectively. Let

$$F(\mathbf{u}) := \sum_{i=1}^{|S|} \ell(f(x_i; \mathbf{w}^1, \mathbf{w}^2, \dots, [\mathbf{w}_1^1, \dots, \mathbf{w}_{s-1}^l, \mathbf{u}, \mathbf{w}_{s+1}^l, \dots, \mathbf{w}_d^l], \dots, \mathbf{w}^L, \mathbf{b}), y_i)$$

denote the loss as a function on the parameters of the  $s$ -th neuron in the  $l$ -th layer (encoded in the  $s$ -th row of  $\mathbf{w}^l$ ) before reparameterization and

$$\tilde{F}(\mathbf{u}) := \sum_{i=1}^{|S|} \ell(f(x_i; \mathbf{w}_\lambda^1, \mathbf{w}_\lambda^2, \dots, [\mathbf{w}_{\lambda 1}^1, \dots, \mathbf{w}_{\lambda(s-1)}^l, \mathbf{u}, \mathbf{w}_{\lambda(s+1)}^l, \dots, \mathbf{w}_{\lambda d}^l], \dots, \mathbf{w}_\lambda^L, \mathbf{b}_\lambda), y_i)$$

denote the loss as a function on the parameters into the  $s$ -th neuron in the  $l$ -th layer (encoded in the  $s$ -th row of  $\mathbf{w}^l$ ) after reparameterization.

For the same layer  $l$ , we define a linear function  $\eta_s : \mathbb{R}^m \rightarrow \mathbb{R}^m$  by

$$\eta_s(\mathbf{u}) = \eta_s(u_1, u_2, \dots, u_m) = (u_1 \lambda_{(s,1)}^l, u_2 \lambda_{(s,2)}^l, \dots, u_m \lambda_{(s,m)}^l).$$

By assumption, we have that  $\tilde{F}(\eta_s(\mathbf{w}_s^l)) = F(\mathbf{w}_s^l)$  for all  $\mathbf{w}_s^l$ . By the chain rule, we compute for any coordinate  $u_t$  of  $\mathbf{u}$ ,

$$\begin{aligned} \frac{\partial F(\mathbf{u})}{\partial u_t} \Big|_{\mathbf{u}=\mathbf{w}_s^l} &= \frac{\partial \tilde{F}(\eta_s(\mathbf{u}))}{\partial u_t} \Big|_{\mathbf{u}=\mathbf{w}_s^l} \\ &= \sum_k \frac{\partial \tilde{F}(\eta_s(\mathbf{u}))}{\partial (\eta_s(\mathbf{u})_k)} \Big|_{\eta_s(\mathbf{u})=\eta_s(\mathbf{w}_s^l)} \cdot \frac{\partial (\eta_s(\mathbf{u})_k)}{\partial u_t} \Big|_{\eta_s(\mathbf{u})=\eta_s(\mathbf{w}_s^l)} \\ &= \frac{\partial \tilde{F}(\mathbf{v})}{\partial v_t} \Big|_{\mathbf{v}=\mathbf{w}_{\lambda_s}^l} \cdot \lambda_{(s,t)}^l. \end{aligned}$$

Similarly, for

$$G(\mathbf{u}, \mathbf{u}') := \sum_{i=1}^{|S|} \ell(f(x_i; \mathbf{w}^1, \mathbf{w}^2, \dots, [\mathbf{w}_1^l, \dots, \mathbf{w}_{s-1}^l, \mathbf{u}, \mathbf{w}_{s+1}^l, \dots, \mathbf{w}_{s'-1}^l, \mathbf{u}', \mathbf{w}_{s'+1}^l, \dots, \mathbf{w}_d^l], \dots, \mathbf{w}^L, \mathbf{b}), y_i)$$

denoting the loss as a function on the parameters of the  $s$ -th and  $s'$ -th neuron in the  $l$ -th layer (encoded in the  $s$ -th and  $s'$ -th row of  $\mathbf{w}^l$ ) before reparameterization and for

$$\begin{aligned} \tilde{G}(\mathbf{u}, \mathbf{u}') &:= \sum_{i=1}^{|S|} \ell(f(x_i; \mathbf{w}_{\lambda}^1, \mathbf{w}_{\lambda}^2, \dots \\ &\dots, [\mathbf{w}_{\lambda 1}^l, \dots, \mathbf{w}_{\lambda(s-1)}^l, \mathbf{u}, \mathbf{w}_{\lambda(s+1)}^l, \dots, \mathbf{w}_{\lambda s'-1}^l, \mathbf{u}', \mathbf{w}_{\lambda s'+1}^l, \dots, \mathbf{w}_{\lambda d}^l], \dots \\ &\dots, \mathbf{w}_{\lambda}^L, \mathbf{b}_{\lambda}), y_i) \end{aligned}$$

we have  $\tilde{G}(\eta_s(\mathbf{w}_s^l), \eta_{s'}(\mathbf{w}_{s'}^l)) = G(\mathbf{w}_s^l, \mathbf{w}_{s'}^l)$ . For all  $s, s', t, t'$  we obtain second derivatives

$$\frac{\partial^2 G(\mathbf{u}, \mathbf{u}')}{\partial u_t \partial u_{t'}} \Big|_{\mathbf{u}=\mathbf{w}_s^l, \mathbf{u}'=\mathbf{w}_{s'}^l} = \lambda_{(s,t)}^l \lambda_{(s',t')}^l \frac{\partial^2 \tilde{G}(\mathbf{u}, \mathbf{u}')}{\partial u_t \partial u_{t'}} \Big|_{\mathbf{u}=\mathbf{w}_{\lambda_s}^l, \mathbf{u}'=\mathbf{w}_{\lambda_{s'}}^l}.$$

Consequently, the Hessian  $H(\mathbf{w}^l, S)$  of the empirical risk before reparameterization and the Hessian  $\tilde{H}(\mathbf{w}_{\lambda}^l, S)$  after reparameterization satisfy at the position corresponding to  $w_{s,t}$  and  $w_{s',t'}$  that

$$H_{s,s'}(\mathbf{w}^l, S)_{(t,t')} = \lambda_{(s,t)}^l \lambda_{(s',t')}^l \cdot \tilde{H}_{s,s'}(\mathbf{w}_{\lambda}^l)_{(t,t')}.$$

Assuming that  $\lambda_s^l := \lambda_{(s,t)}^l = \lambda_{(s,t')}^l$  for all  $s, t$  and  $t'$ , then we get that

$$\begin{aligned} \kappa_{Tr}^l(\mathbf{w}) &= \sum_{s,s'=1}^d \langle \mathbf{w}_s^l, \mathbf{w}_{s'}^l \rangle \cdot Tr(H_{s,s'}(\mathbf{w}^l, S)) \\ &= \sum_{s,s'=1}^d \langle \frac{\mathbf{w}_{\lambda_s}^l}{\lambda_s^l}, \frac{\mathbf{w}_{\lambda_{s'}}^l}{\lambda_{s'}^l} \rangle \cdot Tr(\lambda_s^l \lambda_{s'}^l \tilde{H}_{s,s'}(\mathbf{w}_{\lambda}^l, S)) \\ &= \sum_{s,s'=1}^d \langle \mathbf{w}_{\lambda_s}^l, \mathbf{w}_{\lambda_{s'}}^l \rangle \cdot Tr(\tilde{H}_{s,s'}(\mathbf{w}_{\lambda}^l, S)) \\ &= \kappa_{Tr}^l(\mathbf{w}_{\lambda}) \end{aligned}$$

This proves Theorem 7.

To show the corollary, we first observe that all layer-wise reparameterizations are covered by the theorem. To see this, we only need to check that the condition  $\lambda_{(s,t)}^l = \lambda_{(s,t')}^l$  holds for each  $s, t$  and  $t'$ . For layer-wise reparameterizations, we even have that  $\lambda_{(s,t)}^l = \lambda^l$  for all  $s, t$ , since all weights of one layer are multiplied by the same scalar  $\lambda^l$ , and  $\lambda_{(s,t)}^l = \lambda_{(s,t')}^l$  is easily seen to hold true.

Note further, that any neuron-wise reparameterization given by multiplying all weights into a neuron in a layer  $\iota \neq l - 1$  by  $\lambda > 0$  and dividing all outgoing weights by  $\lambda$  is also covered by the theorem. Hence, the only neuron-wise reparameterization that can change the relative flatness measures is the one multiplying some row of  $\mathbf{w}^{l-1}$  by some  $\lambda > 0$  and dividing the corresponding column of  $\mathbf{w}^l$  by the same  $\lambda$ . However, by multiplying both  $\mathbf{w}^{l-1}$  and  $\mathbf{w}^l$  with  $V^{-1}$  and  $V$  from the left and right respectively, we perform an explicit neuron-wise reparameterization that chooses a unique representative and therefore removes the dependence on such reparameterizations.  $\square$

### D.3 Proof of Theorem 5

In this section, we prove Theorem 5. For clarity, we repeat the assumptions and the statement we prove in this section:

We consider a model  $f(x, \mathbf{w}) = g(\mathbf{w}\phi(x))$ , a loss function  $\ell$  and a sample set  $S$ , and let  $O_m \subset \mathbb{R}^{m \times m}$  denote the set of orthogonal matrices. Let  $\delta$  be a positive (small) real number and  $\mathbf{w} = \omega \in \mathbb{R}^{d \times m}$  denote parameters at a local minimum of the empirical risk on a sample set  $S$ . If the output function satisfies that  $y[\phi_{\delta A}(x_i)] = y[\phi(x_i)] = y_i$  for all  $(x_i, y_i) \in S$  and all matrices  $\|A\| \leq 1$ , then we want to show that  $f(x, \omega)$  is  $((\delta, S, O_m), \epsilon)$ -feature robust on average over  $O_m$  for  $\epsilon = \frac{\delta^2}{2m} \kappa_{Tr}^\phi(\omega) + \mathcal{O}(\delta^3)$ , i.e.,

$$\left| \mathcal{E}_{\mathcal{F}}^\phi(f, S, \alpha A) \right| \leq \frac{\delta^2}{2m} \kappa_{Tr}^\phi(\omega) + \mathcal{O}(\delta^3) \text{ for all } 0 \leq \alpha \leq \delta$$

*Proof.* Writing  $z_i = \phi(x_i)$  and  $\mathcal{E}_{emp}(\mathbf{w}, S) = \mathcal{E}_{emp}(f(\mathbf{w}, x), S)$  and using the assumption that  $y[\phi_{\delta A}(x_i)] = y_i$  for all  $(x_i, y_i) \in S$  and all  $\|A\| \leq 1$ , we have for any  $0 \leq \alpha \leq \delta$ ,

$$\begin{aligned} \mathcal{E}_{\mathcal{F}}^\phi(f, S, \alpha A) + \mathcal{E}_{emp}(\mathbf{w}, S) &= \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\psi[\mathbf{w}, \phi_{\alpha A}(x_i)], y[\phi_{\alpha A}(x_i)]) \\ &= \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\psi(\mathbf{w}, z_i + \alpha A z_i), y_i) \\ &= \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\psi(\mathbf{w} + \alpha \mathbf{w} A, z_i), y_i) \\ &= \mathcal{E}_{emp}(\mathbf{w} + \alpha \mathbf{w} A, S) \end{aligned} \tag{14}$$

The latter is the empirical error  $\mathcal{E}_{emp}(\mathbf{w} + \alpha \mathbf{w} A, S)$  of the model  $f$  on the sample set  $S$  at parameters  $\mathbf{w} + \alpha \mathbf{w} A$ . If  $\delta$  is sufficiently small, then by Taylor expansion around the local minimum  $\mathbf{w} = \omega$ , we have up to order of  $\mathcal{O}(\delta^3)$  that

$$\begin{aligned} \mathcal{E}_{emp}(\omega + \alpha \omega A, S) &= \mathcal{E}_{emp}(\omega, S) + \frac{\alpha^2}{2} \sum_{s,t=1}^d (\omega_s A) \cdot H_{s,t}(\omega, \phi(S)) \cdot (\omega_t A)^T \\ &\leq \mathcal{E}_{emp}(\omega, S) + \frac{\delta^2}{2} \sum_{s,t=1}^d (\omega_s A) \cdot H_{s,t}(\omega, \phi(S)) \cdot (\omega_t A)^T \end{aligned} \tag{15}$$

where  $\omega_s$  denotes the  $s$ -th row of  $\omega$ .

We consider the set of orthogonal matrices  $O_m$  as equipped with the (unique) normalized Haar measure. (For the definition of the Haar measure, see e.g. [17].) We need to show that  $\mathbb{E}_{A \sim O_m} \left[ \mathcal{E}_{\mathcal{F}}^\phi(f, S, \alpha A) \right] \leq \frac{\delta^2}{2m} \sum_{s,t} \langle \omega_s, \omega_t \rangle \cdot Tr(H_{s,t})$  for all  $0 \leq \alpha \leq \delta$  with  $\mathcal{E}_{\mathcal{F}}^\phi(f, S, A)$  defined as in Eq. 4. Using (14) and (15) we get

$$\mathbb{E}_{A \sim O_m} \left[ \mathcal{E}_{\mathcal{F}}^\phi(f, S, \alpha A) \right] \leq \mathbb{E}_{A \sim O_m} \left[ \frac{\delta^2}{2} \sum_{s,t=1}^d (\omega_s A) H_{s,t}(\omega, S) (\omega_t A)^T \right] + \mathcal{O}(\delta^3)$$

Using the unnormalized trace  $Tr([m_{s,t}]) = \sum_s m_{s,s}$  we compute with the help of the so-called Hutchinson's trick:

$$\begin{aligned} Tr(\mathbb{E}_{A \sim O_m} [(\omega_t A)^T (\omega_s A)]) &= \mathbb{E}_{A \sim O_m} [Tr((\omega_t A)^T (\omega_s A))] \\ &= \mathbb{E}_{A \sim O_m} [Tr((\omega_s A) (\omega_t A)^T)] \\ &= \mathbb{E}_{A \sim O_m} [Tr(\omega_s \omega_t^T)] \\ &= \langle \omega_s, \omega_t \rangle \end{aligned}$$



We can interchange two vector coordinates by multiplication of a suitable orthogonal matrix  $B$ . Since the Haar measure is invariant under multiplication of an orthogonal matrix, the diagonal of  $\mathbb{E}_{A \sim O_m} [(\omega_t A)^T (\omega_s A)]$  must contain a constant value. This value along the diagonal must then equal  $\frac{1}{m} \langle \omega_s, \omega_t \rangle$ . Further, we can multiply one vector coordinate by  $(-1)$  via multiplication by an orthogonal matrix, and hence the off-diagonal entries of  $\mathbb{E}_{A \sim O_m} [(\omega_t A)^T (\omega_s A)]$  must be zero, giving that

$$\mathbb{E}_{A \sim O_m} [(\omega_t A)^T (\omega_s A)] = \frac{\langle \omega_s, \omega_t \rangle}{m} \cdot I.$$

Therefore

$$\begin{aligned} \mathbb{E}_{A \sim O_m} [(\omega_s A) H_{s,t} (\omega_t A)^T] &= \text{Tr} (\mathbb{E}_{A \sim O_m} [(\omega_s A) H_{s,t} (\omega_t A)^T]) \\ &= \mathbb{E}_{A \sim O_m} [\text{Tr} ((\omega_s A) H_{s,t} (\omega_t A)^T)] \\ &= \mathbb{E}_{A \sim O_m} [\text{Tr} (H_{s,t} (\omega_t A)^T (\omega_s A))] \\ &= \text{Tr} (H_{s,t} \cdot \mathbb{E}_{A \sim O_m} [(\omega_t A)^T (\omega_s A)]) \\ &= \text{Tr} (H_{s,t} \cdot \frac{\langle \omega_s, \omega_t \rangle}{m} \cdot I) \\ &= \frac{\langle \omega_s, \omega_t \rangle}{m} \text{Tr} (H_{s,t}) \end{aligned}$$

Putting things together, we have for the local optimum  $\mathbf{w} = \omega$  that

$$\begin{aligned} \mathbb{E}_{A \sim O_m} [\mathcal{E}_{\mathcal{F}}^\phi(f, S, \alpha A)] &\leq \frac{\delta^2}{2} \sum_{s,t=1}^d \mathbb{E}_{A \sim O_m} [(\omega_s A) H_{s,t} (\omega_t A)^T] + \mathcal{O}(\delta^3) \\ &= \frac{\delta^2}{2m} \sum_{s,t=1}^d \langle \omega_s, \omega_t \rangle \cdot \text{Tr} (H_{s,t}) + \mathcal{O}(\delta^3) \\ &= \frac{\delta^2}{2m} \kappa_{\text{Tr}}^\phi(\omega) + \mathcal{O}(\delta^3) \end{aligned}$$

□

We can further generalize Theorem 5 to more complex labels by introducing a notion of approximately locally constant labels. The following definition frees us from the strong assumption of locally constant labels, i.e.  $y[\phi_{\delta A}(x_i)] = y[\phi(x_i)] = y_i$  for all  $(x_i, y_i) \in S$  and all matrices  $\|A\| \leq 1$ , while still restricting label changes to be one order smaller than the contribution of flatness.

**Definition 9.** Let  $\mathcal{D}$  be a data distribution on a labeled sample space  $\mathcal{X} \times \mathcal{Y}$  and  $S$  a finite iid sample of  $\mathcal{D}$ . Let  $f = \psi \circ \phi$  be a model composed into a feature extractor  $\phi$  and predictor  $\psi$ . We say that  $\mathcal{D}$  has approximately locally constant labels of order three around the points  $(x, y) \in S$  in feature space  $\phi$ , if there is some constant  $C$  such that

$$\frac{1}{|S|} \sum_{i=1}^{|S|} \left| \ell(\psi(\phi(x_i) + \Delta_i), y[\phi(x_i) + \Delta_i]) - \ell(\psi(\phi(x_i) + \Delta_i), y_i) \right| \leq C\delta^3 \text{ for } \|\Delta_i\| \leq \delta \|\phi(x_i)\|$$

**Corollary 10.** Consider a model  $f(x, \mathbf{w}) = \psi(\mathbf{w}, \phi(x)) = g(\mathbf{w}\phi(x))$  as above, a loss function  $\ell$  and a sample set  $S$ , and let  $O_m \subset \mathbb{R}^{m \times m}$  denote the set of orthogonal matrices. Let  $\delta$  be a positive (small) real number and  $\mathbf{w} = \omega \in \mathbb{R}^{d \times m}$  denote parameters at a local minimum of the empirical risk on a sample set  $S$ . If  $\mathcal{D}$  has approximately locally constant labels of order three around the points  $(x, y) \in S$  in feature space, then  $f(x, \omega)$  is  $((\delta, S, O_m), \epsilon)$ -feature robust on average over  $O_m$  for  $\epsilon = \frac{\delta^2}{2m} \kappa_{\text{Tr}}^\phi(\omega) + \mathcal{O}(\delta^3)$ .

*Proof.* As before, we abbreviate  $\phi(x_i)$  by  $z_i$ . We only need to modify (14) to account for the strictly weaker assumption on the labels. For this, we perform Taylor approximation with respect to the

labels at  $y[\phi(x_i)] = y_i$  to obtain

$$\begin{aligned}
\mathcal{E}_{\mathcal{F}}^{\phi}(f, S, \alpha A) + \mathcal{E}_{emp}(\mathbf{w}, S) &= \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\psi[\phi_{\alpha A}(x_i)], y[\phi_{\alpha A}(x_i)]) \\
&\stackrel{Def\ 9}{\leq} \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\psi(\mathbf{w}, z_i + \alpha A z_i), y_i) + \mathcal{O}(\delta^3) \\
&= \frac{1}{|S|} \sum_{i=1}^{|S|} \ell(\psi(\mathbf{w} + \alpha \mathbf{w} A, z_i), y_i) + \mathcal{O}(\delta^3) \\
&= \mathcal{E}_{emp}(\mathbf{w} + \alpha \mathbf{w} A, S) + \mathcal{O}(\delta^3)
\end{aligned}$$

The rest of the proof follows the arguments used to show Theorem 5.

□

#### D.4 Proof of Theorem 6

To prove Theorem 6, we will require a proposition that bounds  $\epsilon$ -representativeness for  $\lambda_i$  the local densities from Proposition 2. This is achieved in Proposition 12 below uniformly over all distributions  $\mathcal{D}$  that satisfy mild regularity assumptions necessary for a well-defined kernel density estimation. We first compose the proof to Theorem 6 and subsequently show the arguments leading to the required proposition.

The main idea to prove Theorem 6 is that the family of distributions considered in Proposition 2 (a) provides an explicit link between  $\epsilon$ -representativeness and feature robustness (Proposition 2 and Equation 5), (b) allows us to approximately bound feature robustness by relative flatness (Theorem 5), and (c) allows us to apply a kernel density estimation to uniformly bound  $\epsilon$ -representativeness (Proposition 12).

Theorem 6 is the informal counterpart to the following version.

**Theorem 11.** *Consider a model  $f(x, \mathbf{w}) = g(\mathbf{w}\phi(x))$ , a loss function  $\ell$ , a sample set  $S$ , and let  $m$  denote the dimension of the feature space defined by  $\phi$  and let  $\delta$  be a positive (small) real number. Let  $\omega \in \mathbb{R}^{d \times m}$  denote a local minimum of the empirical risk on an iid sample set  $S$ .*

*Suppose that the distribution  $\mathcal{D}$  has a smooth density  $p_{\mathcal{D}}^{\phi}$  on the feature space  $\mathbb{R}^m$  such that  $\int_z |\nabla^2 (p_{\mathcal{D}}^{\phi}(z) \|z\|^2)| dz$  and  $\int_z \frac{p_{\mathcal{D}}^{\phi}(z)}{\|z\|^m} dz$  are well-defined and finite. Then for sufficiently large sample size  $|S|$ , if the distribution has approximately locally constant labels of order three (see Definition 9), then it holds with probability  $1 - \Delta$  over sample sets  $S$  that*

$$\mathcal{E}_{gen}(f(\cdot, \omega), S) \lesssim |S|^{-\frac{2}{4+m}} \left( \frac{\kappa_{Tr}^{\phi}(\omega)}{2m} + C_1(p_{\mathcal{D}}^{\phi}, L) + \frac{C_2(p_{\mathcal{D}}^{\phi}, L)}{\sqrt{\Delta}} \right)$$

*up to higher orders in  $|S|^{-1}$  for constants  $C_1, C_2$  that depend only on the distribution in feature space  $p_{\mathcal{D}}^{\phi}$  induced by  $\phi$  and the chosen  $|S|$ -tuple  $\Lambda_{\delta}$  as in Proposition 2 and the maximal loss  $L$ .*

*Proof.* The proof combines Equation 5 with Proposition 12 and Theorem 5. At first we use Equation 5 to split the generalization gap into  $\mathcal{E}_{gen}(f) = \mathcal{E}_{Rep}^{\phi}(f, S, \Lambda_{\delta}) + \mathcal{E}_{\mathcal{F}}(f, S, \mathcal{A}_{\delta})$ . For the family of distributions  $\Lambda_{\delta}$  from Proposition 2, we have by Proposition 12 that

$$|\mathcal{E}_{Rep}^{\phi}(\psi \circ \phi, S, \Lambda_{\delta}^{\phi})| \leq \left( C_1(p_{\mathcal{D}}^{\phi}, L) + \frac{C_2(p_{\mathcal{D}}^{\phi}, L)}{\sqrt{\Delta}} \right) \cdot |S|^{-\frac{2}{4+m}} + \mathcal{O}(|S|^{-\frac{3}{4+m}})$$

when  $\delta = |S|^{-\frac{1}{4+m}}$ . With  $\mathcal{E}_{\mathcal{F}}(f, S, \mathcal{A}_{\delta}) = \mathbb{E}_{A \sim \mathcal{A}_{\delta}} [\mathcal{E}_{\mathcal{F}}(f, S, A)]$ , we use (the proof to) Proposition 2 to see that this can be written as

$$\mathbb{E}_{A \sim \mathcal{A}_{\delta}} [\mathcal{E}_{\mathcal{F}}(f, S, A)] = \mathbb{E}_{0 \leq \alpha \leq \delta} \left[ \mathbb{E}_{A \sim (\mathcal{O}_m, \mu)} [\mathcal{E}_{\mathcal{F}}(f, S, \alpha A)] \right]$$

where  $\mathcal{O}_m \subset \mathbb{R}^{m \times m}$  denote the set of orthogonal matrices and  $\mu$  the Haar measure on this set.

Finally, Theorem 5 bounds the latter by  $\frac{|S|^{-\frac{2}{4+m}}}{2m} \kappa_{Tr}^{\phi}(\omega)$  up to higher orders in  $|S|^{-1}$ .  $\square$

We finally prove that the bound on  $\epsilon$ -representativeness in the proof to the preceding Theorem indeed holds true.

**Proposition 12.** *Consider a model  $f(x, \mathbf{w}) = \psi(\mathbf{w}, \phi(x))$ , a loss function  $\ell$  and let  $S \subseteq \mathcal{X} \times \mathcal{Y}$  be a finite sample set. With  $x_i \in S$ , let  $\lambda_i(z) = k_{\delta \|\phi(x_i)\|}(0, z)$  define an  $|S|$ -tuple  $\Lambda_{\delta}$  of densities as in Proposition 2 and assume that the loss function is bounded by  $L$ . Suppose that the distribution  $\mathcal{D}$  has a smooth density  $p_{\mathcal{D}}^{\phi}$  on a feature space  $\mathbb{R}^m$  such that  $\int_z \nabla^2 (p_{\mathcal{D}}^{\phi}(z) \|z\|^2) dz$  and  $\int_z \frac{p_{\mathcal{D}}^{\phi}(z)}{\|z\|^m} dz$  are well-defined and finite. Then there exist constants  $C_1(p_{\mathcal{D}}^{\phi}, L), C_2(p_{\mathcal{D}}^{\phi}, L)$  depending on the distribution and the maximal loss such that, with probability  $1 - \Delta$  over possible sample sets  $S$ ,  $\epsilon$ -interpolation is bounded for  $\delta = |S|^{-\frac{1}{4+m}}$  by*

$$|\mathcal{E}_{Rep}^{\phi}(f, S, \Lambda_{\delta})| \leq \left( C_1(p_{\mathcal{D}}^{\phi}, L) + \frac{C_2(p_{\mathcal{D}}^{\phi}, L)}{\sqrt{\Delta}} \right) \cdot |S|^{-\frac{2}{4+m}} + \mathcal{O}(|S|^{-\frac{3}{4+m}})$$

*Proof.* We let

$$\hat{p}(z) = \frac{1}{|S|} \sum_{i=1}^{|S|} k_{\delta|\phi(x_i)|}(\phi(x_i), z)$$

With  $\lambda_i = k_{\delta|\phi(x_i)|}(0, z)$  we have

$$\begin{aligned} \left| \mathcal{E}_{Rep}^{\phi}(f, S, \Lambda_{\delta}) \right| &= \left| \mathcal{E}(f) - \frac{1}{|S|} \sum_{i=1}^{|S|} \mathbb{E}_{\xi \sim \lambda_i} [\ell(\psi(\phi(x_i) + \xi), y[\phi(x_i) + \xi])] \right| \\ &= \left| \int_z p_{\mathcal{D}}^{\phi}(z) \cdot \ell(\psi(z), y(z)) dz - \frac{1}{|S|} \sum_{i=1}^{|S|} \int_z k_{\delta|\phi(x_i)|}(\phi(x_i), z) \cdot \ell(\psi(z), y(z)) dz \right| \quad (16) \\ &\leq \underbrace{\left| \int_z (p_{\mathcal{D}}^{\phi}(z) - \mathbb{E}_S [\hat{p}(z)]) \cdot \ell(\psi(z), y(z)) dz \right|}_{(I)} + \underbrace{\left| \int_z (\mathbb{E}_S [\hat{p}(z)] - \hat{p}(z)) \cdot \ell(\psi(z), y(z)) dz \right|}_{(II)} \end{aligned}$$

For the further analysis, we make use of Jones et al. [15] and combine it with the generalization to the multivariate case in Chp. 4.3.1 in Silverman [33]. A Taylor approximation with respect to the bandwidth of the kernel  $\delta$  yields

$$(I) = \frac{\delta^2}{2} \tau_2 \left| \int_z \nabla^2 \left( p_{\mathcal{D}}^{\phi}(z) \|z\|^2 \right) \ell(\psi(z), y(z)) dz \right| + \mathcal{O}(\delta^3)$$

where

$$\tau_2 = \int_z \|z\|^2 k_1(0, z) dz.$$

For (II) we consider the random variable  $Z = \int_z \hat{p}(z) \ell(\psi(z), y(z)) dz$  as a function on the set of possible sample sets of a fixed size. Applying Chebychef's inequality on  $Z$ , we get that

$$Pr \left( \left| Z - \mathbb{E}_S [Z] \right| > \epsilon_{est} \right) \leq \frac{Var(Z)}{\epsilon_{est}^2} =: \Delta .$$

Solving for  $\epsilon_{est}$  yields that with probability  $1 - \Delta$  we have

$$(II) = \left| Z - \mathbb{E}_S [Z] \right| \leq \frac{\sqrt{Var(Z)}}{\sqrt{\Delta}}$$

Further, the variance of  $Z$  can be bounded by

$$\begin{aligned} Var(Z) &= \mathbb{E}_S \left[ (Z - \mathbb{E}_S [Z])^2 \right] \\ &= \mathbb{E}_S \left[ \left( \int \hat{p}(z) \ell(\psi(z), y(z)) dz - \mathbb{E}_S \left[ \int \hat{p}(z) \ell(\psi(z), y(z)) dz \right] \right)^2 \right] \\ &= \mathbb{E}_S \left[ \left( \int (\hat{p}(z) - \mathbb{E}_S [\hat{p}(z)]) \ell(\psi(z), y(z)) dz \right)^2 \right] \\ &\leq \underbrace{\mathbb{E}_S \left[ \int (\hat{p}(z) - \mathbb{E}_S [\hat{p}(z)])^2 dz \right]}_{(III)} \cdot \underbrace{\left( \int_z \ell(\psi(z), y(z))^2 dz \right)}_{\leq L^2 \text{Vol}(\phi(\mathcal{D}))} \end{aligned}$$

It follows from Eq. (2.3) in Jones et al. [15] together with Eq. 4.10 in Silverman [33] for (III) that for small  $\delta$  and large sample size  $|S|$  the term (III), i.e., the variance of  $\tilde{p}$ , is given by

$$(III) = \beta |S|^{-1} \delta^{-m} \alpha + \mathcal{O}(|S|^{-2}) ,$$

where  $\alpha = \int_z \frac{p_{\mathcal{D}}^{\phi}(z)}{\|z\|^m} dz$  and  $\beta = \int_z k_1(0, z)^2 dz$ . Putting things together gives

$$\begin{aligned} \mathcal{E}_{Rep}^{\phi}(f, S, \Lambda_{\delta}) &\leq L \frac{\delta^2}{2} \tau_2 \left| \int_z \nabla^2 \left( p_{\mathcal{D}}^{\phi}(z) \|z\|^2 \right) dz \right| + \frac{L \sqrt{\alpha \beta}}{\sqrt{\Delta}} \sqrt{\text{Vol}(\phi(\mathcal{D}))} |S|^{-\frac{1}{2}} \delta^{-\frac{m}{2}} \\ &\quad + \mathcal{O}(|S|^{-2}) + \mathcal{O}(\delta^3) . \end{aligned}$$

Choosing the bandwidth as  $\delta = |S|^{-\frac{1}{4+m}}$  gives

$$|\mathcal{E}_{Rep}^\phi(f, S, \Lambda_\delta)| \leq |S|^{-\frac{2}{4+m}} \left( \tau_2 L \left| \int_z \nabla^2 \left( p_{\mathcal{D}}^\phi(z) \|z\|^2 \right) \right| dz + \frac{\sqrt{\alpha\beta}L}{\sqrt{\Delta}} \sqrt{\text{Vol}(\phi(\mathcal{D}))} \right) + \mathcal{O}(|S|^{-\frac{3}{m+4}}) .$$

The result follows from setting

$$C_1 = \tau_2 L \left| \int_z \nabla^2 \left( p_{\mathcal{D}}^\phi(z) \|z\|^2 \right) \right| dz$$

$$C_2 = \sqrt{\alpha\beta}L \sqrt{\text{Vol}(\phi(\mathcal{D}))} .$$

□

## E Relative flatness for a uniform bound over general distributions on feature matrices

This article based its consideration on the specific distribution on feature matrices of Proposition 2, since this distribution allows to use standard results of kernel density estimation in the proof to Theorem 6. However, the decomposition of the risk in Equation 5 holds for any distribution on feature matrices  $\mathcal{A}$  and induced distributions on feature space  $\Lambda_{\mathcal{A}}$ . To allow maximal flexibility in the choice of a distribution  $\mathcal{A}$  on feature matrices of norm  $\|A\| \leq 1$ , we define another version of relative flatness based on the maximal eigenvalues of partial Hessians instead of the trace.

**Definition 13.** For a model  $f(\mathbf{w}, x) = g(\mathbf{w}\phi(x))$  with a twice differentiable function  $g$ , a twice differentiable loss function  $\ell$  and a sample set  $S$  we define maximal relative flatness by

$$\kappa^\phi(\mathbf{w}) := \sum_{s=1}^d \|\mathbf{w}_s\|^2 \cdot \lambda_{max}(H_{s,s}(\mathbf{w}, \phi(S))) \quad (17)$$

where  $\lambda_{max}$  denotes the maximal eigenvalue of a matrix and  $H_{s,s}$  the Hessian matrix as in (8).

The analogue to Theorem 5 for maximal relative flatness shows that maximal flatness bounds feature robustness uniformly over all feature matrices of norm  $\|A\| \leq 1$ .

**Theorem 14.** Consider a model  $f(x, \mathbf{w}) = g(\mathbf{w}\phi(x))$  as above, a loss function  $\ell$  and a sample set  $S$ , and let  $O_m \subset \mathbb{R}^{m \times m}$  denote the set of orthogonal matrices. Let  $\delta$  be a positive (small) real number and  $\mathbf{w} = \omega \in \mathbb{R}^{d \times m}$  denote parameters at a local minimum of the empirical risk on a sample set  $S$ . If the labels satisfy that  $y[\phi_{\delta A}(x_i)] = y[\phi(x_i)] = y_i$  for all  $(x_i, y_i) \in S$  and all  $\|A\| \leq 1$ , then, for each feature selection matrix  $\|A\| \leq 1$  the model  $f(x, \omega)$  is  $((\delta, S, A), \epsilon)$ -feature robust for  $\epsilon = \frac{\delta^2 d}{2} \kappa^\phi(\omega) + \mathcal{O}(\delta^3)$

*Proof.* Writing  $z_i = \phi(x_i)$  and  $\mathcal{E}_{emp}(\mathbf{w}, S) = \mathcal{E}_{emp}(f(\mathbf{w}, x), S)$  and using the assumption that  $y[\phi_{\delta A}(x_i)] = y_i$  for all  $(x_i, y_i) \in S$  and all  $\|A\| \leq 1$ , we have by the first part of the proof of Theorem 5 that for any  $0 \leq \alpha \leq \delta$ ,

$$\mathcal{E}_{\mathcal{F}}^\phi(f, S, \alpha A) + \mathcal{E}_{emp}(\mathbf{w}, S) = \mathcal{E}_{emp}(\mathbf{w} + \alpha \mathbf{w}A, S) \quad (18)$$

and

$$\mathcal{E}_{emp}(\omega + \alpha \omega A, S) \leq \mathcal{E}_{emp}(\omega, S) + \frac{\delta^2}{2} \sum_{s,t=1}^d (\omega_s A) \cdot H_{s,t}(\omega, \phi(S)) \cdot (\omega_t A)^T + \mathcal{O}(\delta^3) \quad (19)$$

at a local minimum  $\omega$ , where  $\omega_s$  denotes the  $s$ -th row of  $\omega$ .

Note that for  $\|A\| \leq 1$  and a row vectors  $\mathbf{w}_s$  it holds that  $\|\mathbf{w}_s A\| \leq \|\mathbf{w}_s\|$ . Further, since the full Hessian matrix  $H(\omega, S) = (H_{s,t}(\omega, S))_{s,t}$  is a positive semidefinite matrix at a local minimum  $\omega$ , it holds for each row vectors  $\mathbf{w}_s, \mathbf{w}_t$  that

$$\mathbf{w}_s H_{s,t}(\omega, S) \mathbf{w}_t^T \leq \frac{1}{2} \left( \mathbf{w}_s H_{s,s}(\omega, S) \mathbf{w}_s^T + \mathbf{w}_t H_{t,t}(\omega, S) \mathbf{w}_t^T \right), \quad (20)$$

We therefore get that for any feature matrix  $A$  with  $\|A\| \leq 1$ ,

$$\begin{aligned}
\mathcal{E}_{\mathcal{F}}^{\phi}(f, S, \delta A) &\leq \max_{\|A\| \leq 1} \mathcal{E}_{\mathcal{F}}(f, S, \delta A) \\
&\stackrel{(18),(19)}{\leq} \max_{\|A\| \leq 1} \frac{\delta^2}{2} \sum_{s,t=1}^d (\omega_s A) \cdot H_{s,t}(\omega, S) \cdot (\omega_t A)^T + \mathcal{O}(\delta^3) \\
&\stackrel{(20)}{\leq} \max_{\|A\| \leq 1} \frac{\delta^2 d}{2} \sum_{s=1}^d (\omega_s A) \cdot H_{s,s}(\omega, S) \cdot (\omega_s A)^T + \mathcal{O}(\delta^3) \\
&\leq \frac{\delta^2 d}{2} \sum_{s=1}^d \max_{\|z\| \leq \|\omega_s\|} z H_{s,s}(\omega, S) z^T + \mathcal{O}(\delta^3) \\
&= \frac{\delta^2 d}{2} \sum_{s=1}^d \max_{\|z\|=1} \|\omega_s\|^2 z H_{s,s}(\omega, S) z^T + \mathcal{O}(\delta^3) \\
&= \frac{\delta^2 d}{2} \sum_{s=1}^d \|\omega_s\|^2 \lambda_{max}(H_{s,s}(\omega, S)) + \mathcal{O}(\delta^3) \\
&= \frac{\delta^2 d}{2} \kappa^{\phi}(\omega) + \mathcal{O}(\delta^3)
\end{aligned} \tag{21}$$

where we used the identity that  $\max_{\|x\|=1} x^T M x = \lambda_{max}(M)$  for any symmetric matrix  $M$ . □

With this, the analogue to Theorem 6 (or its version Theorem 11 in the appendix) allows maximal flexibility to choose  $\mathcal{A}_{\delta}$  (and  $\delta > 0$ ) to bound representativeness. This leads to the following generalization bound.

**Theorem 15.** *Consider a model  $f(x, \mathbf{w}) = g(\mathbf{w}\phi(x))$ , a loss function  $\ell$ , a sample set  $S$ , and let  $m$  denote the dimension of the feature space defined by  $\phi$  and let  $\delta$  be a positive (small) real number. Let  $\omega \in \mathbb{R}^{d \times m}$  denote a local minimum of the empirical risk on an iid sample set  $S$ .*

*Let  $\Upsilon_{\delta}$  be the set of all  $|S|$ -tuple of distributions  $\Lambda_{\mathcal{A}_{\delta}}$  on feature vectors induced by a distribution  $\mathcal{A}_{\delta}$  on feature matrices of norm smaller than  $\delta$  as in Section 3. Then it holds that*

$$\mathcal{E}_{gen}(f(\cdot, \omega), S) \leq \inf_{\mathcal{A}_{\delta} \in \Upsilon_{\delta}} \mathcal{E}_{Rep}^{\phi}(f, S, \Lambda_{\mathcal{A}_{\delta}}) + \frac{\delta^2 d}{2} \sum_{s=1}^d \kappa^{\phi}(\omega) + \mathcal{O}(\delta^3).$$

*Proof.* Part (i) follows from combining Equation 5 with Theorem 14. First, we use (5) to split the generalization gap into  $\mathcal{E}_{gen}(f) = \mathcal{E}_{Rep}^{\phi}(f, S, \Lambda_{\mathcal{A}_{\delta}}) + \mathcal{E}_{\mathcal{F}}(f, S, \mathcal{A}_{\delta})$ . Then, Theorem 14 shows that  $\mathcal{E}_{\mathcal{F}}(f, S, \mathcal{A}_{\delta}) \leq \frac{\delta^2 d}{2} \kappa^{\phi}(\omega) + \mathcal{O}(\delta^3)$  as  $\mathcal{E}_{\mathcal{F}}(f, S, \delta A) \leq \frac{\delta^2 d}{2} \kappa^{\phi}(\omega) + \mathcal{O}(\delta^3)$  for all  $\|A\| \leq 1$ . □

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
  - (b) Did you describe the limitations of your work? [Yes] Yes, we describe limitations and assumptions in the introduction, as well as at the appropriate positions in the text and discuss broader properties of the results in this paper in the discussion.
  - (c) Did you discuss any potential negative societal impacts of your work? [No] This work contributes to the theoretical understanding of machine learning, in particular deep learning. As such, it has no direct ethical or societal impact. However, this understanding contributes to the trustworthiness of machine learning. Thus, it indirectly impacts the acceptance of automatic decisions made with such models and might facilitate their adoption in critical fields, such as medical applications and autonomous driving.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We have carefully read the ethics review guidelines and confirm that this paper conforms to them.
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [Yes] We state all assumptions precisely in the introduction and provide further details, interpretations, and all proofs in the appendix.
  - (b) Did you include complete proofs of all theoretical results? [Yes] The proofs can be found in Appdx D.
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Details are in the supplemental material, an implementation is published in a repository at <https://anonymous.4open.science/r/RelativeFlatnessAndGeneralization-B175> (anonymized).
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Details are specified in the appendix and the published code contains all training details to reproduce the results.
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] In our experiments, we report the individual results of multiple runs and capture variations by reporting the correlation.
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] The details are provided in appendix C.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [Yes] We cite existing assets we used in our experiments (CIFAR-10 dataset, LeNet5 network, PacBayes-measure implementation).
  - (b) Did you mention the license of the assets? [N/A]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [No]
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]



(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]