

---

# Explicit Regularization in Overparametrized Models via Noise Injection

---

**Antonio Orvieto\***

Department of Computer Science  
ETH Zürich, Zürich, Switzerland

**Anant Raj\***

Coordinated Science Laboratory  
University of Illinois Urbana-Champaign  
Inria, Ecole Normale Supérieure  
PSL Research University, Paris, France

**Hans Kersting\***

Inria, Ecole Normale Supérieure  
PSL Research University, Paris, France

**Francis Bach**

Inria, Ecole Normale Supérieure  
PSL Research University, Paris, France

## Abstract

Injecting noise within gradient descent has several desirable features, such as smoothing and regularizing properties. In this paper, we investigate the effects of injecting noise before computing a gradient step. We demonstrate that small perturbations can induce explicit regularization for simple models based on the  $\ell_1$ -norm, group  $\ell_1$ -norms, or nuclear norms. However, when applied to overparametrized neural networks with large widths, we show that the same perturbations can cause variance explosion. To overcome this, we propose using independent layer-wise perturbations, which provably allow for explicit regularization without variance explosion. Our empirical results show that these small perturbations lead to improved generalization performance compared to vanilla gradient descent.

## 1 Introduction

Injecting noise in gradient descent has several desirable properties. In previous works [Jin et al., 2017, Reddi et al., 2018, Staib et al., 2019], noise injection *after the gradient step* has been used to escape saddle points efficiently in non-convex optimization. In several recent works [Zhu et al., 2019, Nguyen et al., 2019], it was also shown that noise injection in gradient descent helps to escape local minima while optimizing non-convex functions. Adding noise *within the gradient*, either on labels or on the iterates

before computing the (stochastic) gradient, has been classically used to smooth the objective function (often in the convex case) [Nesterov and Spokoiny, 2017, Polyak and Tsybakov, 1990, Flaxman et al., 2005, Duchi et al., 2012, d’Aspremont and El Karoui, 2014, Nesterov, 2005]. However, the effect of noise injection on generalization is not as well understood.

Recently, there has been some interests in understanding the geometry of the loss landscape enforced by the induced noise and its connection with generalization. A recent result [Liu et al., 2021] shows that perturbed gradient descent with sufficiently small noise converges to flat minima for non-convex matrix factorization tasks. [Orvieto et al., 2022] also provide some explanation and empirical evidence for better performance of noise induced gradient descent as the noise injection favors flat minima.

In this work, we explore several avenues along this line of research. Our aim is to tackle the following questions:

1. Do small perturbations in the model while training induce an explicit form of regularization for simple models common in machine learning?
2. Do the same perturbations also induce an explicit regularization effect while training overparametrized models such as homogeneous neural networks with large (infinite) widths?
3. Is the induced regularization useful in practice, i.e., does it lead to better generalization performance?

In response to the questions above, we make the following contributions:

1. We show in Section 2 that small perturbations induce explicit regularization for simple models, and that this induced regularization is based on the  $\ell_1$ -norm, group  $\ell_1$ -norm, or nuclear norms. Such explicit formulations

---

\* Equal contribution

were available before for only special cases ( $\ell_1$ -norm), or for special loss functions (logistic regression in the separable case).

2. We show in Section 3 that the same perturbations do not work for overparametrized neural networks with large widths due to variance explosion resulting from the number of perturbed neurons tending to infinity. However, we also show that *independent layer-wise perturbations* (where only one layer is perturbed at each gradient step) allow to avoid the exploding variance term, and explicit regularizers can be obtained, for a fully connected network with a single hidden layer. This is extended in Section 4 to deep ReLU networks.
3. We empirically show in Section 5 that small layer-wise perturbations lead to better generalization performance in practice over vanilla (stochastic) gradient descent training, for a variety of shallow and deep overparametrized models, fully connected or convolutional, with minor adjustments to the training procedure.

## 1.1 Related work

Noise injection has been frequently used for various tasks in statistics, machine learning and signal processing. In this section, we discuss three major directions of research and existing literature associated with them which are closely connected to our work.

**Perturbed gradient descent (PGD).** PGD methods are versions of (stochastic) gradient descent (GD) where noise is injected in one of the following two ways: In the first option, perturbations are added *after* each GD step. Such versions of PGD have been shown to speed up the escape from spurious local minima and saddle points; see Zhou et al. [2019] and Jin et al. [2021], respectively. These methods correspond to discretizations of the continuous-time Langevin dynamics [Li et al., 2017] where the perturbations add to the diffusion coefficient. In the second option, perturbations are added to the iterate *before* the gradient is evaluated. The so-perturbed gradient is then used to update the unperturbed iterate. In the convex case, this is typically used to smooth a non-smooth objective function [Nesterov and Spokoiny, 2017, Polyak and Tsybakov, 1990, Flaxman et al., 2005, Duchi et al., 2012, d’Aspremont and El Karoui, 2014, Nesterov, 2005]. In the non-convex case, such versions of PGD have been shown to exhibit an implicit bias toward flat minima (that generalize better) in the case of nonconvex matrix factorization [Liu et al., 2021] and more general models [Orvieto et al., 2022]. Flat minima are widely believed to have a better generalization property [Keskar et al., 2017, Chaudhari et al., 2019, Jiang et al., 2019], and we follow this line of work, with several generalization guarantees that are based on PAC-Bayesian analysis and also apply to our framework [Tsuzuku et al., 2020, Neyshabur et al., 2018].

The second option of PGD can be interpreted as an instance of the first option where the perturbations are instead anti-correlated. Such methods are therefore also referred to as *Anti-PGD*; see Section 2 from Orvieto et al. [2022]. In our current paper, we determine the *effective loss* that is (on average) used in the second option.

**Explicit and implicit regularization in learning.** It is well understood that regularization plays a crucial role in the generalization performance of a learning model. The most common form of regularization is Tikhonov regularization [Weese, 1993, Golub et al., 1999], which is imposed explicitly in the optimization objective for linear and non-linear models. There are other ways to induce regularization explicitly while training machine learning models [Hanson and Pratt, 1988, Srivastava et al., 2014, Raj and Bach, 2021]. Apart from explicit regularization methods, optimization algorithms can also induce implicit bias on the optimal solution, which was observed by Soudry et al. [2018], Gunasekar et al. [2018], Kubo et al. [2019]. In a closely related line of work, Camuto et al. [2020] study the regularization induced in neural networks by Gaussian noise injections to all network activations, and studies its regularization effect in the Fourier domain. However, while they obtain explicit regularizers similar to ours, we study the effect of noise injection to the model parameter instead of network activations, and hence our analysis is applicable to general machine learning models. Moreover, we deal with infinite widths through layer-wise perturbations.

**Robustness.** Apart from inducing regularization and finding flatter minima, noise injection is also helpful in getting robust models. Recent works [Salman et al., 2019, Cohen et al., 2019, Lecuyer et al., 2019, Li et al., 2018] have shown that randomized smoothing is a scalable way of building provably robust neural network based classifiers. However, these approaches rely on input perturbation instead of model perturbation. Wu et al. [2020] showed that adversarial weight perturbation helps in flattening the weight loss landscape, which improves the robust generalization gap. Zheng et al. [2021] also employ adversarial model perturbation to favor flat local minima of the empirical risk. However, a systematic study of the influence of noise injection in the model is lacking.

## 2 Finite-dimensional models

In this section, we consider a loss function  $L : \mathbb{R}^n \rightarrow \mathbb{R}$ , as well as a predictor  $\Phi : \mathbb{R}^m \rightarrow \mathbb{R}^n$ , and we aim to minimize

$$R(w) = L(\Phi(w)) \quad (1)$$

with respect to  $w \in \mathbb{R}^m$ . In machine learning contexts,  $\Phi(w)$  represents the  $n$  real-valued predictions over  $n$  inputs, and  $L$  is the associated loss function, typically,  $L(\varphi) = \frac{1}{2n} \sum_{i=1}^n (y_i - \varphi_i)^2$  for least-squares regression, or  $L(\varphi) = \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i \varphi_i))$  for logistic regression, where  $y \in \mathbb{R}^n$  is the vector of output labels. The

framework applies equally well to multi-dimensional outputs (such as multi-category classification and multinomial loss), where  $\mathbb{R}^n$  is replaced by  $\mathbb{R}^{n \times k}$  (this is in fact considered in Sections 2.4, 3 and 4 for neural networks).

In this section, we only consider the asymptotics with respect to  $\sigma$ , assuming that the number of parameters  $m$  remains fixed (hence the denomination “finite-dimensional”). We analyze the effect of strong overparametrization with the number of parameters tending to infinity in Sec. 3 and 4.

## 2.1 Gaussian perturbations and Taylor expansions

Following previous work on PGD (see Section 1.1), we consider perturbing the iterate by a standard Gaussian vector  $\varepsilon$  (with zero mean and identity covariance matrix), and consider the function

$$R_\sigma(w) = \mathbb{E}[L(\Phi(w + \sigma\varepsilon))]. \quad (2)$$

We denote by  $D\Phi(w) \in \mathbb{R}^{n \times m}$  the Jacobian (matrix of first-order derivatives) of  $\Phi$ , and  $D^2\Phi(w) \in \mathbb{R}^{n \times m \times m}$  the tensor of second-order derivatives.

**Taylor expansions.** Throughout this section, we assume for simplicity of exposition that  $L$  and  $\Phi$  are three times differentiable functions with bounded third derivatives. In Appendix A we discuss the case of functions where this is satisfied only piece-wise (to cover ReLU activations).

The Taylor expansion of  $\Phi$  at  $w$  can be written:

$$\begin{aligned} \Phi(w + \sigma\varepsilon) &= \Phi(w) + \sigma D\Phi(w)\varepsilon \\ &\quad + \frac{\sigma^2}{2} D^2\Phi(w)[\varepsilon\varepsilon^\top] + O(\sigma^3\|\varepsilon\|^3), \end{aligned} \quad (3)$$

where  $D^2\Phi(w)[M] \in \mathbb{R}^n$  is defined as  $D^2\Phi(w)[M]_j = \sum_{a,b=1}^m D^2\Phi(w)_{jab} M_{ab}$ .

We also need a Taylor expansion of  $L$  around  $\Phi(w)$ , which can be written as

$$\begin{aligned} L(\Phi(w) + \Delta) &= L(\Phi(w)) + DL(\Phi(w))\Delta \\ &\quad + \frac{1}{2} D^2L(\Phi(w))[\Delta\Delta^\top] + O(\|\Delta\|^3), \end{aligned} \quad (4)$$

where, for  $\varphi \in \mathbb{R}^n$ ,  $DL(\varphi) \in \mathbb{R}^{1 \times n}$  is the row-vector of first-order partial derivatives, and  $D^2L(\varphi) \in \mathbb{R}^{n \times n}$  the matrix of second-order partial derivatives, with the notation  $D^2L(\varphi)[M] = \sum_{a,b=1}^n D^2L(\varphi)_{ab} M_{ab} \in \mathbb{R}$ .

To compute an expansion of  $R_\sigma$ , as defined in Eq. (2), we compose the two expansions and get, with  $\Delta = \sigma D\Phi(w)\varepsilon + \frac{\sigma^2}{2} D^2\Phi(w)[\varepsilon\varepsilon^\top] + O(\sigma^3\|\varepsilon\|^3)$  the expansion

$$\begin{aligned} &L(\Phi(w + \sigma\varepsilon)) \\ &= L(\Phi(w) + \sigma D\Phi(w)\varepsilon + \frac{\sigma^2}{2} D^2\Phi(w)[\varepsilon\varepsilon^\top] + O(\sigma^3\|\varepsilon\|^3)) \\ &= L(\Phi(w)) + DL(\Phi(w))(\sigma D\Phi(w)\varepsilon + \frac{\sigma^2}{2} D^2\Phi(w)[\varepsilon\varepsilon^\top] \\ &\quad + O(\sigma^3\|\varepsilon\|^3)) + \frac{1}{2} D^2L(\Phi(w))[\sigma D\Phi(w)\varepsilon(\sigma D\Phi(w)\varepsilon)^\top \\ &\quad + O(\sigma^3\|\varepsilon\|^3)] + O(\sigma^3\|\varepsilon\|^3). \end{aligned}$$

Taking expectations, using  $\mathbb{E}[\varepsilon] = 0$  and  $\mathbb{E}[\varepsilon\varepsilon^\top] = I$ :

$$\begin{aligned} R_\sigma(w) &= L(\Phi(w)) + \frac{\sigma^2}{2} DL(\Phi(w))D^2\Phi(w)[I] \\ &\quad + \frac{\sigma^2}{2} D^2L(\Phi(w))[D\Phi(w)D\Phi(w)^\top] + O(\sigma^3). \end{aligned} \quad (5)$$

(The careful reader will have noticed that we assumed that  $\Phi$  is in  $C^3$ ; in Appendix A we present an extension to the case where  $\Phi$  is only piecewise in  $C^3$  — e.g. ReLU.) Note that the added terms on top of  $R(w) = L(\Phi(w))$  can be derived as well from the trace of the Hessian of  $R$ , as obtained by Orvieto et al. [2022].

**Minimizing  $R_\sigma$  with stochastic gradient descent.** As discussed by Orvieto et al. [2022], in order to minimize  $R_\sigma$  defined as an expectation in Eq. (2), we do not need to compute the expectation: we simply use a single perturbation (i.e., a single Gaussian vector  $\varepsilon$ ), and compute the gradient with respect to parameters, which leads to an unbiased estimate that can be used within a stochastic gradient descent algorithm. Since the loss function  $L$  has a finite sum structure, we can also use stochastic gradient for over observations, as traditionally done in machine learning.

## 2.2 Asymptotically equivalent objective functions

In several situations, we will show that we obtain the following “asymptotically equivalent” cost function  $R_\sigma^{(\text{eff})}(w)$ , defined as

$$R_\sigma^{(\text{eff})}(w) = R(w) + \frac{\sigma^2}{2} D^2L(\Phi(w))[D\Phi(w)D\Phi(w)^\top], \quad (6)$$

in the sense that finding a global minimizer of  $R_\sigma(w) = \mathbb{E}[L(\Phi(w + \sigma\varepsilon))]$  is essentially equivalent to finding a global minimizer of  $R_\sigma^{(\text{eff})}(w)$  as  $\sigma \rightarrow 0$ . In other words, the term  $\frac{\sigma^2}{2} DL(\Phi(w))D^2\Phi(w)[I]$  in Eq. (5) has no impact. This will be done by showing that the order of approximation between  $R$  and  $R_\sigma$  is of order  $\sigma^2$ , while the one between  $R_\sigma$  and  $R_\sigma^{(\text{eff})}$  is of order  $\sigma^3$ , in two different cases.

**Hessians with no cross-products.** One simple sufficient condition for the term  $\frac{\sigma^2}{2} DL(\Phi(w))D^2\Phi(w)[I]$  to have no impact is that  $D^2\Phi(w)[I] = 0$  for all  $w \in \mathbb{R}^m$ . This is satisfied for certain models such as neural networks, where the second-order derivatives have no cross-product terms, or equivalently when  $\text{tr}(\Phi_j(w)) = 0$  for all  $j = 1, \dots, m$ ; this is e.g. the case when the ReLU activation is used, see Sections 2.4, 3, and 4 for more such examples. The following theorem makes it precise for sufficiently regular functions. See proof and more general results in App. A (already sketched above).

**Theorem 1.** *Assume that (a)  $\Phi$  and  $L$  are three-times continuously differentiable with uniformly bounded third derivatives, (b)  $D^2\Phi(w)[I] = 0$  for all  $w \in \mathbb{R}^m$ . Then, there exist constants  $C$  and  $C'$  (independent of  $w$ ) such that  $\forall w \in \mathbb{R}^m$ ,  $|R_\sigma(w) - R(w)| \leq C(1 + \|w\|^2)\sigma^2$  while  $|R_\sigma(w) - R_\sigma^{(\text{eff})}(w)| \leq C'\sigma^3$ .*

**Overparametrized models.** A more general condition is related to overparametrization (when the model is sufficiently rich to lead to the minimizer of  $L$ ), and we make a formal statement below with simplified assumptions. See Appendix A for the proof and more refined conditions.

**Theorem 2.** *Assume that (a)  $\Phi$  is three-times continuously differentiable with uniformly bounded second and third derivatives, (b)  $L$  is strongly convex with uniformly bounded second and third derivatives, with unique global minimizer  $\varphi^*$ , (c) there exists a (non-unique)  $w_* \in \mathbb{R}^m$  such that  $\Phi(w_*) = \varphi^*$ , and that (d) there exist minimizers  $w_*^\sigma$  and  $w_*^{\sigma,(\text{eff})}$  of  $R_\sigma$  and  $R_\sigma^{(\text{eff})}$  that lie in a compact set  $\Omega \subset \mathbb{R}^m$ . Then, if  $w_*^\sigma$  is a minimizer of  $R_\sigma$ , and  $w_*^{\sigma,(\text{eff})}$  is a minimizer of  $R_\sigma^{(\text{eff})}$ , we have  $\|\Phi(w_*^\sigma) - \varphi^*\|_2^2 = O(\sigma^2)$  while  $\|\Phi(w_*^{\sigma,(\text{eff})}) - \Phi(w_*^\sigma)\|_2^2 = O(\sigma^3)$ .*

Note that (a) we characterize the prediction function  $\Phi$  taken at the various minimizers to characterize asymptotic equivalence (but it is not possible to characterize a distance between parameters because in overparametrized models,  $\Phi$  cannot be injective), and (b) when  $\sigma$  tends to zero the minimizer should converge to the interpolator  $\Phi(w) = \varphi_*$  with minimal  $D^2L(\Phi(w)) [D\Phi(w)D\Phi(w)^\top]$ .

Below, we consider specific cases for the least square loss and the logistic loss in the separable case, before considering specific simple models for  $\Phi$ .

**Least-squares regression.** We consider  $L(\varphi) = \frac{1}{2n} \|y - \varphi\|_2^2$ , which is a quadratic function with constant Hessian  $\frac{1}{n}I$ . We therefore get an extra regularizer on top of  $R(w)$  equal to  $\frac{\sigma^2}{n} \text{tr} [DL(\Phi(w))DL(\Phi(w))^\top]$ .

Note that our regularizer is then related to the neural tangent kernel [Jacot et al., 2018], i.e., this is the trace of the corresponding kernel matrix. While this could lead to interesting further connections, we note that our work applies beyond the lazy training regime (that is, our parameters  $w$  can vary significantly) [Chizat et al., 2019].

**Separable logistic regression.** If we consider the logistic loss with an overparametrized model, then the infimum of  $L$  is attained at infinity, and we should expect  $\Phi(w)$  to diverge along a specific direction, like in other similar models [Soudry et al., 2018, Chizat and Bach, 2020].

For homogeneous models like considered by [Lyu and Li, 2019], we can expect  $w$  to diverge in some direction, that is,  $w = \lambda\Delta$  for  $\lambda \rightarrow +\infty$ . The term  $D\Phi(w)D(\Phi(w))^\top$  grows in  $\lambda$ , while  $D^2L(\Phi(w))$  converges to zero exponentially fast. Overall, we conjecture (and empirically observe) that we get divergence, and that the direction  $\Delta$  ends up being proportional to the minimizer of  $\|D\Phi(\Delta)\|_F^2$  such that  $y \circ \Phi(\Delta) \geq 1$ , where  $y \in \{-1, 1\}^n$  is the vector of labels. See Appendix B for more details.

### 2.3 Lasso and L1-norm

Following [Woodworth et al., 2020], [Vaskevicius et al., 2019], [Pesme et al., 2021], we consider ‘‘diagonal net-

works’’. All of the three mentioned works achieve implicit sparsity regularization under sufficiently small initialization of the model parameter. Here we show that noise injection in the model also induces an  $\ell_1$ -regularization in the problem. The main advantage we have over the discussed work is that our approach is applicable to more complex models as we show below.

We consider  $w = (w_1, w_2) \in \mathbb{R}^{2d}$ , and  $\Phi(w) = X(w_1 \circ w_1 - w_2 \circ w_2)$  for  $X \in \mathbb{R}^{n \times d}$ ,  $L(\varphi) = \frac{1}{2n} \|y - \varphi\|_2^2$ , with  $m = 2d$ . When the model is overparametrized, that is  $X$  has rank  $n$ , we can apply Theorem 2, and we then get an equivalent risk:

$$R_\sigma^{(\text{eff})}(w) = \frac{1}{2n} \|y - X(w_1 \circ w_1 - w_2 \circ w_2)\|_2^2 + 2\sigma^2 \text{diag}(X^\top X/n)^\top (w_1 \circ w_1 + w_2 \circ w_2), \quad (7)$$

which is exactly the Lasso once considering  $\beta = w_1 \circ w_1 - w_2 \circ w_2$ , that is, minimizing

$$\frac{1}{2n} \|y - X\beta\|_2^2 + 2\sigma^2 \text{diag}(X^\top X/n)^\top |\beta|, \quad (8)$$

where  $\beta = w_1 \circ w_1 - w_2 \circ w_2$  (an explicit derivation of Eq. (7) is given in App. C). However, there exist many efficient algorithms dealing with  $\ell_1$ -regularization [see Bach et al., 2012, and references therein], and one should probably not use our reduction to optimize Eq. (8).

### 2.4 Nuclear norm (linear networks)

Following [Baldi and Hornik, 1995], [Arora et al., 2019], [Saxe et al., 2019], [Gidel et al., 2019], we consider  $w = (W_1, W_2)$ , with  $W_1 \in \mathbb{R}^{d_1 \times d_0}$  and  $W_2 \in \mathbb{R}^{d_2 \times d_1}$ , and  $\Phi(w) = W_2 W_1 X^\top$  with input data  $X \in \mathbb{R}^{n \times d_0}$ . We consider the square loss for simplicity, that is,  $L(\varphi) = \frac{1}{2n} \|Y^\top - \varphi\|_F^2$  for a response  $Y \in \mathbb{R}^{n \times d_2}$  and  $\varphi \in \mathbb{R}^{d_2 \times n}$ . We can then apply Thm. 2, where the Hessian has no diagonal term, and we get (see detailed computations in App. C):

$$R_\sigma^{(\text{eff})}(W_1, W_2) = \frac{1}{2n} \|Y^\top - W_2 W_1 X^\top\|_F^2 + \frac{\sigma^2}{2n} [d_2 \|W_1 X^\top\|_F^2 + \|W_2\|_F^2 \cdot \|X\|_F^2]. \quad (9)$$

Given the matrix  $M = W_2 W_1 \in \mathbb{R}^{d_2 \times d_0}$ , we can optimize  $\frac{1}{2} [d_2 \|W_1 X^\top\|_F^2 + \|W_2\|_F^2 \cdot \|X\|_F^2]$  with respect to compatible matrices  $W_1$  and  $W_2$ , leading to the penalty  $\sqrt{d_2} \|X\|_F \cdot \|MX^\top\|_*$ , where  $\|\cdot\|_*$  is the nuclear norm (sum of singular values), which favors low-rank matrices. Thus, minimizing  $R_\sigma^{(\text{eff})}$  above is equivalent to minimizing

$$\frac{1}{2n} \|Y^\top - MX^\top\|_F^2 + \frac{\sigma^2}{n} \sqrt{d_2} \|X\|_F \cdot \|MX^\top\|_*. \quad (10)$$

We obtain a nuclear norm penalty, thus providing an explicit regularizer in situations where the classical techniques for diagonal networks do not extend. See an empirical illustration in Section 5.

We would also like to mention here a conjecture from [Gunasekaran et al., 2017], which claims that with small enough



step sizes and initialization close enough to the origin, gradient descent on a full dimensional factorization converges to the minimum nuclear norm solution, which we obtain explicitly using our noise injection approach without making those assumptions. Like for the Lasso, we are not advocating for gradient descent to be a particularly good algorithm for solving the problem above, although a detailed analysis is worthwhile performing as future work.

## 2.5 Extensions

Other models could be considered as well, with relationships to other explicit regularizers.

**Group Lasso.** It is traditional to recover the group Lasso as a special case of nuclear norm minimization [see, e.g., Bach 2008]. This is detailed in Appendix C and can also be applied to the Lasso with a different formulation from Section 2.3 (and no need for overparametrization).

**Optimizing over the PSD cone.** By considering  $w \in \mathbb{R}^{d \times r}$ , with  $r \geq d$  and  $\Phi(w) = ww^\top \in \mathbb{R}^{d \times d}$ , then for a cost function  $L$  over symmetric matrices, we would get a trace penalty over positive semi-definite matrices.

**Beyond quadratic models.** Thus far, we have considered only functions  $\Phi$  that were quadratic. This can be extended to higher order polynomials, like done for linear networks below, but we could imagine applications to tensor decomposition beyond the matrix decomposition above [Kolda 2001, Kolda and Bader 2009, Sidiropoulos et al. 2017].

## 3 One-hidden-layer neural networks with infinite widths

So far we have discussed the effect of noise injection on simple machine learning models. In this section, we will focus our attention on more complex models like neural networks with potentially large widths (where the limit of large widths is taken *before* the limit of small  $\sigma$ ). In this section, we start with a single hidden layer, and extend to deeper networks in Section 4.

We first start with linear networks, before considering ReLU activation functions in Sec. 3.3. For simplicity, we only consider the (multivariate) square loss in our analysis.

### 3.1 Exploding variance for linear networks

As already shown in Section 2.4, we consider  $W_1 \in \mathbb{R}^{d_1 \times d_0}$  and  $W_2 \in \mathbb{R}^{d_2 \times d_1}$ , and  $\Phi(W_1, W_2) = W_2 W_1 X^\top$  with input data  $X \in \mathbb{R}^{n \times d_0}$ . We consider  $L(\varphi) = \frac{1}{2n} \|Y^\top - \varphi\|_F^2$  for  $Y \in \mathbb{R}^{n \times d_2}$ , the matrix of labels (we allow here for multi-dimensional outputs).

We here consider the overparametrized limit  $d_1 \rightarrow +\infty$  using initialization with random  $\square$  weights of order  $(W_1)_{ij} \sim$

<sup>1</sup>The same analysis holds of course for Glorot initialization [Glorot and Bengio 2010], with similar results.




Figure 1: Numerical illustration of the analysis in Sec. 3.1 & 3.3 on a linear network with 1 hidden layer of dimension  $d_1$  (ReLU case and deep networks studied in the coming sections). If noise with standard deviation  $\sigma$  is added to all weights (**left**), then the regularized loss  $R_\sigma$  explodes as  $d_1 \rightarrow \infty$  due to the variance term  $\sigma^4 d_2 d_1 \|X\|_F^4$ . Instead, perturbing  $W_1$  and  $W_2$  in alternation (**right**) with standard deviation  $\sqrt{2}\sigma$  provides mathematically the same expected regularization but avoids the variance term and therefore provides a much more reliable regularization, as clear also from the experimental section. Runs are repeated 100 times and shown is the average.

$\frac{1}{\sqrt{d_1 d_0}}$  for all  $i, j$ , that is  $\|W_1\|_F^2$  not exploding with  $d_1$ , and  $(W_2)_{ij} \sim \frac{1}{\sqrt{d_2 d_1}}$ , that is  $\|W_2\|_F^2$  not exploding with  $d_1$ .

We have, with Gaussian perturbations  $E_1$  and  $E_2$ , an explicit exact expansion for linear networks:

$$\begin{aligned} \Phi(w + \sigma\varepsilon) &= (W_2 + \sigma E_2)(W_1 + \sigma E_1)X^\top \\ &= W_2 W_1 X^\top + \sigma(W_2 E_1 + E_2 W_1)X^\top + \sigma^2 E_2 E_1 X^\top. \end{aligned}$$

Taking expectations and using that  $E_1, E_2$  have zero mean and are independent, and such that,  $\mathbb{E}[E_i M E_i^\top] = \text{tr}(M)I$  for  $i = 1, 2$ , and  $M$  any symmetric matrix of compatible size, we can get  $\mathbb{E}[\Phi(w + \sigma\varepsilon)] = \Phi(w)$ , and:

$$\begin{aligned} \mathbb{E}[\|\Phi(w + \sigma\varepsilon)\|_F^2] &= \|\Phi(w)\|_F^2 + \sigma^2 [\|W_2\|_F^2 \|X\|_F^2 \\ &\quad + d_2 \|W_1 X^\top\|_F^2] + \sigma^4 d_2 d_1 \|X\|_F^2. \end{aligned}$$

We can now compute  $R_\sigma$  as:

$$\begin{aligned} R_\sigma(W_1, W_2) &= R(W_1, W_2) + \frac{\sigma^2}{2n} [\|W_2\|_F^2 \|X\|_F^2 \\ &\quad + d_2 \|X W_1^\top\|_F^2] + \frac{\sigma^4}{2n} d_1 d_2 \|X\|_F^2. \end{aligned}$$

We recover the expression from Section 2.4, but we have an extra term  $\frac{\sigma^4}{2n} d_2 d_1 \|X\|_F^2$ , which is of superior order in  $\sigma$ , but problematic when  $d_1 \rightarrow +\infty$ .

Note that  $\frac{\sigma^2}{2n} \text{tr}(W_2^\top W_2) \text{tr}(X^\top X)$  scales as  $\frac{\sigma^2}{n} \|X\|_F^2$ , while the term  $\frac{\sigma^2}{n} d_2 \text{tr}(W_1 X^\top X W_1^\top)$  scales as  $d_2 \frac{\sigma^2}{n} \|X\|_F^2$ , with thus no explosion in  $d_1$ . However, the term  $\frac{\sigma^4}{2n} d_2 d_1 \|X\|_F^2$  explodes when  $d_1$  goes to infinity. *This is the exploding variance problem of perturbing all layers simultaneously* — see discussion in the experiment section as well as Figures 2 and 4. We provide a fix by perturbing one layer at a time, which we now present.

### 3.2 Layer-wise perturbations for linear networks

We consider the selection of one layer  $j \in \{1, 2\}$  out of the two layers at random, and add  $\sqrt{2}E_j$  to  $W_j$  before computing the gradient. This corresponds to minimizing an objective function  $\tilde{R}_\sigma(W_1, W_2)$ , which is an expectation (both with respect to  $j$  and  $E_j$ ) equal to

$$\begin{aligned} \tilde{R}_\sigma(W_1, W_2) &= \frac{1}{2}\mathbb{E}[R(W_1 + \sigma\sqrt{2}E_1, W_2)] \\ &\quad + \frac{1}{2}\mathbb{E}[R(W_1, W_2 + \sigma\sqrt{2}E_2)]. \end{aligned}$$

Hence

$$\begin{aligned} \tilde{R}_\sigma(W_1, W_2) &= R(W_1, W_2) + \frac{1}{4n}\mathbb{E}[\|\sigma\sqrt{2}W_2E_1X^\top\|_F^2] \\ &\quad + \frac{1}{4n}\mathbb{E}[\|\sigma\sqrt{2}E_2W_1X^\top\|_F^2] \\ &= R(W_1, W_2) + \frac{\sigma^2}{2n}[\text{tr}(W_2^\top W_2)\text{tr}(X^\top X) \\ &\quad + d_2\text{tr}(W_1X^\top XW_1^\top)], \end{aligned}$$

that is, the exploding variance term is not there anymore, and we arrive at the same nuclear norm interpretation as in Section 2.4. This generalizes to non-linear activation functions as we show next.

### 3.3 One-hidden layer with non-linear activations

We now consider the same set-up as above, but with ReLU activations, that is,  $\Phi(w) = W_2(W_1X^\top)_+$ , with  $X \in \mathbb{R}^{n \times d_0}$ , and  $L(\varphi) = \frac{1}{2n}\|Y^\top - \varphi\|_F^2$ . We consider the square loss for simplicity, but this applies more generally.

We also consider weights  $W_1$  and  $W_2$  initialized with respective scales  $1/\sqrt{d_1d_0}$  and  $1/\sqrt{d_2d_1}$ . After rescaling, it corresponds exactly to the mean-field limit that allows representation learning [Chizat and Bach, 2018, Mei et al., 2018, Wojtowysch, 2020, Sirignano and Spiliopoulos, 2022], for which the scaling s.t.  $\|W_2\|_F^2$  and  $\|W_1\|_F^2$  do not explode is preserved throughout optimization.

Here, we consider asymptotics explicitly with respect  $d_1$  (going to infinity), and then with respect to  $\sigma$  (going to zero). We show in App. D that the same variance explosion as in Sec. 3.1, when perturbing all layers simultaneously, still occurs in the limit of infinite widths and non-zero  $\sigma$ . We thus only focus here on the layer-wise perturbation.

Locally, we obtain  $\Phi(W_1 + \sigma\sqrt{2}E_1, W_2) = \Phi(W_1, W_2) + \sigma\sqrt{2}W_2[(W_1X^\top)_+^0 \circ E_1X^\top]$  and  $\Phi(W_1, W_2 + \sigma\sqrt{2}E_2) = \Phi(W_1, W_2) + \sigma\sqrt{2}E_2(W_1X^\top)_+$ ; see the Appendix D for a formal derivation. We can therefore compute  $\tilde{R}_\sigma$  as:

$$\begin{aligned} \tilde{R}_\sigma(W_1, W_2) &= R(W_1, W_2) + \frac{\sigma^2}{2n}\|(W_1X^\top)_+\|_F^2 \\ &\quad + \frac{\sigma^2}{2n}\sum_{j=1}^{d_1}\sum_{i=1}^n\|(W_2)_{\cdot j}\|_2^2 \times |((W_1X^\top)_+^0)_{ji}|^2. \end{aligned}$$

Note that here our regularizer is 2-homogeneous, like considered by [Chizat and Bach, 2020], but that the penalty is not any more separable in  $W_1$  and  $W_2$ . Like done by [Chizat and Bach, 2020], we can take the number  $d_1$  of hidden neurons to infinity and study the associated function space norm, which has similar adaptivity properties (see the Appendix D for details).

## 4 Deeper networks

In this section, we show that we still get the explicit effect of perturbing each layer selected randomly with standard deviation  $\sigma\sqrt{M}$ , where  $M$  is the number of layers, and the extra  $\sqrt{M}$  term is added to take into account that the expectation with respect to the choice of the layer adds a multiplicative factor  $1/M$ . The main purpose of this section is to show that the expression of  $R_\sigma^{(\text{eff})}$  found for finite-dimensional models in Section 2 is still valid, that is:

$$R_\sigma^{(\text{eff})}(w) = R(w) + \frac{\sigma^2}{2}\text{tr}(\nabla^2 L(\Phi(w))D\Phi(w)D\Phi(w)^\top). \quad (11)$$

Crucially, note that with full perturbations, and with a number of neurons going to infinity, we would get a variance explosion (see discussion in Appendix E). This finding is also supported by our experimental findings in the next section.

In this section, to simplify the exposition and avoid technicalities, we only consider linear networks (see Appendix E for non-linear ones).

**Linear networks.** Consider a linear neural network with  $w = (W_1, \dots, W_M)$  with  $W_i \in \mathbb{R}^{d_i \times d_{i-1}}$ , and  $\Phi(w) = W_M W_{M-1} \dots W_1 X^\top \in \mathbb{R}^{d_M \times n}$  with the input data  $X \in \mathbb{R}^{n \times d_0}$ , with  $L(\varphi) = \frac{1}{2n}\|Y^\top - \varphi\|_F^2$  for output data  $Y \in \mathbb{R}^{n \times d_M}$  and  $\varphi \in \mathbb{R}^{d_M \times n}$ . We now analyze the effect of independent layer-wise perturbations.

We consider the selection of one layer  $j$  out of the  $M$  layers at random, and add a random Gaussian matrix  $\sqrt{M}E_j$  to  $W_j$ . We then have an expectation (both with respect to  $j$  and  $E_j$ ), equal to

$$\begin{aligned} \tilde{R}_\sigma(W_1, \dots, W_M) &= \frac{1}{M}\sum_{j=1}^M\mathbb{E}[R(W_1, \dots, W_j, W_j + \sqrt{M}\sigma E_j, W_{j+1}, \dots, W_M)] \\ &= \frac{1}{M}\sum_{j=1}^M\mathbb{E}[L(W_M W_{M-1} \dots W_1 X^\top \\ &\quad + \sqrt{M}\sigma W_M \dots W_{j+1} E_j W_{j-1} \dots W_1 X^\top)]. \end{aligned}$$

We can then compute expectations and get

$$\begin{aligned} \tilde{R}_\sigma(W_1, \dots, W_M) &= L(\Phi(W_1, \dots, W_M)) \\ &\quad + \frac{\sigma^2}{2n}\sum_{j=1}^M\|W_M \dots W_{j+1}\|_F^2\|W_{j-1} \dots W_1 X^\top\|_F^2, \end{aligned}$$

which is exactly Eq. (11). We thus have a non-exploding explicit regularizer, based on the gradient of the prediction




Figure 2: MLP with 1 hidden layer and linear activations on synthetic data (see the Appendix for a full description and loss curves). For 50 hidden neurons (**left panels**) both injecting noise layer-wise and to all weights lead to minimization of the regularized loss. If we have 300 hidden neurons (**right panels**), then only injecting noise layer-wise takes us close to a minimizer for the regularized loss. For all methods the learning rate is set to 0.1.

function *with respect to the parameters*. We show in Section 5 that this extra regularization does improve test accuracy, and leave for future work the detailed theoretical study of this new regularizer.

## 5 Experiments

The goal of this section is to provide experimental evidence<sup>2</sup> to back-up the results of the last sections. In particular, we compare gradient descent (GD) with the perturbed variants studied in this paper: (1) noise injection at each weight in the model, for each iteration (named “GD+noise”), and (2) layer-wise perturbations (injection layer is sampled at random) at each iteration (named “GD+noise (layer)”). All perturbations are performed before the gradient computations. As seen in Section 3.2, if the variance of layer-specific weight perturbations is  $\sigma^2 M$ , where  $M$  is the number of layers, then the two noise injection methods minimize a similar regularized loss, in expectation. However, as the degree of overparametrization (e.g., number of parameters) increases, the theory suggests that layer-wise perturbations are preferable since they overcome variance explosion (see Sections 3.1 and 3.2).

**Minimization of the regularized loss.** For a start, we consider a one hidden layer network with linear activations and either 50 or 300 hidden neurons on a randomly generated sparse<sup>3</sup> synthetic regression data set with inputs in  $\mathbb{R}^{10}$  and outputs in  $\mathbb{R}$ . In Figure 2, we show the dynamics of the original and regularized loss derived in this paper. As the theory predicts, minimization of the regularized loss in the strongly overparametrized setting is only achievable with layer-wise perturbations (see Section 3.2). As discussed, to keep the same explicit regularization, the variance of the noise in the layer-wise approach is doubled compared to the vanilla approach. This depth scaling is adopted for all

further experiments.

**Effect of tuning.** The findings of the last paragraph are reported for one specific value of  $\sigma$ . We test the influence of  $\sigma$  on a slightly more complex model: an MLP with 2 hidden layers (5000 neurons each) and ReLU activations on Fashion MNIST [Xiao et al., 2017] (classification). Given that the data is relatively easy to fit, we train on a subset of 1024 data points — to induce heavy overparametrization. We run full-batch gradient descent with learning rate 0.005, and plot the test accuracy evolution (computed using 10K data points) for different values of  $\sigma$ . Figure 3 shows that, for this wide model, layer-wise perturbation acts as an effective regularizer that is able to increase the test accuracy. This is in contrast to standard noise injection (cf. Section 3.1). A similar result also holds true for CIFAR10 on ResNet18 (see Figure 6).




Figure 3: Test accuracy for an MLP on FMNIST (2 HL with 5000 neurons). Comparison of perturbation effects.

**Deep MLPs.** Next, we test how the findings carry over to deeper networks. In the same data setting as the last paragraph, we consider now a ReLU MLP with 5 hidden layers and either 1000 (narrow) or 5000 (wide) hidden neurons. In Figure 4 we test our methods ( $\sigma = 0.05$ ) against full-batch GD, with a step size of 0.005 in the narrow setting and 0.001 in the wide setting, to account for the more complex landscape. As can be seen, again layer-wise perturbations yield the best result. This is also reflected in the size of the regularizer (trace of Hessian), which is minimized by the best-performing method in terms of test error.

<sup>2</sup>[https://github.com/aorvieto/noise\\_injection\\_overparam](https://github.com/aorvieto/noise_injection_overparam)

<sup>3</sup>We consider 40 data points sampled from a Gaussian in 10 dimensions. The solution is sparse and is planted as the result of a linear prediction from the first two dimensions.




Figure 4: MLP with 5 hidden Layers and ReLU activations on a subset of Fashion MNIST. The hidden layers each have 1000 neurons (narrow, **left** panels) or 5000 neurons (wide, **right** panels). We run full-batch gradient descent. The results show that injecting noise separately in each layer results in improved regularization and test accuracy, even in the overparametrized setting.




Figure 5: Behavior of noise injection in convolutional networks with different degrees of overparametrization. On the two **left** plots we show the performance using a toy PyTorch CNN+MLP (5 channels) while on the two **right** plots we perform the same task on a fully convolutional network with 5 layers and 128 channels on 3 layers. All details are reported in the Appendix F. We run SGD with batch size 1024 and learning rate 0.005 (0.001 for the PyTorch CNN+MLP) and cosine annealing [Loshchilov and Hutter, 2016]. We plot the performance for  $\sigma = 0.005$  of the two noise injection methods. Plots of the loss dynamics are provided in Appendix F.

**Small Convolutional Networks without normalization.**

In Fig. 5 we go a bit beyond our theoretical setting and test the application to convolutional networks on CIFAR10 [Krizhevsky and Hinton, 2009], with full data and stochastic gradients (batch-size 1024). In this setting, we use a Toy CNN (2 conv layers with 5 channels + MLP, details in the App), and compare against a wider fully-convolutional network with moderate width (4 conv layers with 128 channels + one linear layer). The pattern we discussed above can be observed in this setting, albeit less pronounced due to the narrow nature of convolutional networks with reasonable number of channels. Note that, in the fully convolutional wide network, the trace of the Hessian in the noise injection setting is damped but oscillates around zero. This is quite different from the MLP setting.

**Deep Residual Networks.**

To conclude, we test noise injection (layer-wise or in all layers simultaneously) on a ResNet18 (around 11M parameters) [He et al., 2016] with batch normalization. We use for this the baseline SGD configuration in the popular repository <https://github.com/kuangliu/pytorch-cifar>. On this baseline, which reaches around 94.4% test accuracy, we simply add noise<sup>4</sup> injection at every step. In Figure 6 we tested different noise injection standard deviations  $\sigma$  and plotted mean and standard error of the mean (3 runs) for the final test accuracy and Hessian trace after 150 epochs. We stopped noise injection at epoch 135 to allow the networks to converge and use cosine annealing, batch-size 128 and learning

rate 0.01 in all methods. The results clearly show that injecting noise in all parameters is able to only regularize the objective and improve test accuracy for very small  $\sigma$ , but for  $\sigma > 0.1$  it hurts performance. Instead, injecting layer-wise noise provides a much more consistent regularization, and is able to improve performance by +0.3%, which is a sizable margin given the strong SGD baseline. The poor performance of standard noise injection is predicted by the theory in Section 3, which explains the sharp increase of the Hessian trace (for SGD + noise) observed in the second panel at  $\sigma = 0.1$ .




Figure 6: **Final Test accuracy (left)** and Hessian trace (**right**) for SGD on a CIFAR10 ResNet18 with batch normalization. Both plots show mean and standard error of the mean (3 runs). SGD is plotted at  $\sigma = 1e - 3$  instead of  $\sigma = 0$  for better visualization. Layer-wise noise injection is able to boost performance by explicit regularization with a high  $\sigma$ . Instead, for such high  $\sigma$ , injecting noise simultaneously in all parameters results in instabilities. We remind that the total noise injection variance in normalized in the two methods, as detailed in this section.

<sup>4</sup>Here by “layers” we mean each network parameter group. That is, noise is also injected in the batch-norm parameters.



## 6 Conclusion

In this paper, we showed how explicit regularizers could be obtained from small random perturbations of parameters, both for simple machine learning models, and deep learning models with potentially a large number of neurons. This regularization is obtained by independent layer-wise perturbations, can be easily added to existing training codes, and lead to improved accuracies in our experiments.

Several avenues for future research naturally open: (a) obtain non-asymptotic optimization convergence results for the models in Section 2, (b) a detailed theoretical study of the regularization properties of our new regularizer for deep networks, and (c) performance ablations on other architectures such as transformers [Vaswani et al., 2017]

## Acknowledgements

Anant Raj is supported by the Marie Skłodowska-Curie Fellowship (project NN-OVEROPT 101030817). Hans Kersting and Francis Bach thank the European Research Council for support through ERC grant 724063.

## References

- Sanjeev Arora, Noah Golowich, Nadav Cohen, and Wei Hu. A convergence analysis of gradient descent for deep linear neural networks. In *International Conference on Learning Representations*, 2019.
- Francis Bach. Consistency of trace norm minimization. *Journal of Machine Learning Research*, 9:1019–1048, 2008.
- Francis Bach. Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research*, 18(1):629–681, 2017.
- Francis Bach, Rodolphe Jenatton, Julien Mairal, and Guillaume Obozinski. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 4(1):1–106, 2012.
- Pierre F. Baldi and Kurt Hornik. Learning in linear neural networks: A survey. *IEEE Transactions on Neural Networks*, 6(4):837–858, 1995.
- Alexander Camuto, Matthew Willetts, Umut Simsekli, Stephen J. Roberts, and Chris C. Holmes. Explicit regularisation in Gaussian noise injections. *Advances in Neural Information Processing Systems*, 33:16603–16614, 2020.
- Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-SGD: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124018, 2019.
- Lenaic Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in Neural Information Processing Systems*, 31, 2018.
- Lenaic Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pages 1305–1338, 2020.
- Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. *Advances in Neural Information Processing Systems*, 32, 2019.
- Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, pages 1310–1320, 2019.
- Alexandre d’Aspremont and Nouredine El Karoui. A stochastic smoothing algorithm for semidefinite programming. *SIAM Journal on Optimization*, 24(3):1138–1177, 2014.
- John C. Duchi, Peter L. Bartlett, and Martin J. Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.
- Abraham D. Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the Symposium on Discrete algorithms*, pages 385–394, 2005.
- Gauthier Gidel, Francis Bach, and Simon Lacoste-Julien. Implicit regularization of discrete gradient dynamics in linear neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pages 249–256, 2010.
- Gene H. Golub, Per Christian Hansen, and Dianne P. O’Leary. Tikhonov regularization and total least squares. *SIAM Journal on Matrix Analysis and Applications*, 21(1):185–194, 1999.
- Suriya Gunasekar, Blake E. Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. *Advances in Neural Information Processing Systems*, 30, 2017.
- Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. In *International Conference on Machine Learning*, pages 1832–1841, 2018.
- Stephen Hanson and Lorien Pratt. Comparing biases for minimal network construction with back-propagation. *Advances in Neural Information Processing Systems*, 1, 1988.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in Neural Information Processing Systems*, 31, 2018.
- Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. In *International Conference on Learning Representations*, 2019.
- Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M. Kakade, and Michael I. Jordan. How to escape saddle points efficiently. In *International Conference on Machine Learning*, pages 1724–1732, 2017.
- Chi Jin, Praneeth Netrapalli, Rong Ge, Sham M. Kakade, and Michael I. Jordan. On nonconvex optimization for machine learning: Gradients, stochasticity, and saddle points. *J. ACM*, 68(2), 2021.
- Nitish Shirish Keskar, Jorge Nocedal, Ping Tak Peter Tang, Dheevatsa Mudigere, and Mikhail Smelyanskiy. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2017.
- Tamara G. Kolda. Orthogonal tensor decompositions. *SIAM Journal on Matrix Analysis and Applications*, 23(1):243–255, 2001.
- Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images, 2009.
- Masayoshi Kubo, Ryotaro Banno, Hidetaka Manabe, and Masataka Minoji. Implicit regularization in over-parameterized neural networks. *arXiv preprint arXiv:1903.01997*, 2019.
- V. Kurkova and M. Sanguinetti. Bounds on rates of variable-basis and neural-network approximation. *IEEE Transactions on Information Theory*, 47(6):2659–2665, Sep 2001.
- Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. In *Symposium on Security and Privacy (SP)*, pages 656–672. IEEE, 2019.
- Bai Li, Changyou Chen, Wenlin Wang, and Lawrence Carin. Second-order adversarial attack and certifiable robustness, 2018.
- Qianxiao Li, Cheng Tai, and Weinan E. Stochastic modified equations and adaptive stochastic gradient algorithms. In *Proceedings of the International Conference on Machine Learning*, volume 70, pages 2101–2110, 2017.
- Tianyi Liu, Yan Li, Song Wei, Enlu Zhou, and Tuo Zhao. Noisy gradient descent converges to flat minima for non-convex matrix factorization. In *International Conference on Artificial Intelligence and Statistics*, pages 1891–1899, 2021.
- Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. In *International Conference on Learning Representations*, 2019.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- Yu Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566, 2017.
- Behnam Neyshabur, Srinadh Bhojanapalli, and Nathan Srebro. A PAC-Bayesian approach to spectrally-normalized margin bounds for neural networks. In *International Conference on Learning Representations*, 2018.
- Thanh Huy Nguyen, Umut Simsekli, Mert Gurbuzbalaban, and Gaël Richard. First exit time analysis of stochastic gradient descent under heavy-tailed gradient noise. *Advances in Neural Information Processing Systems*, 32, 2019.
- Antonio Orvieto, Hans Kersting, Frank Proske, Francis Bach, and Aurelien Lucchi. Anticorrelated noise injection for improved generalization. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 17094–17116. PMLR, 2022.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch, 2017.
- Scott Pesme, Loucas Pillaud-Vivien, and Nicolas Flammarion. Implicit bias of SGD for diagonal linear networks: a provable benefit of stochasticity. *Advances in Neural Information Processing Systems*, 34, 2021.
- Boris Teodorovich Polyak and Aleksandr Borisovich Tsybakov. Optimal order of accuracy of search algorithms in stochastic optimization. *Problemy Peredachi Informatsii*, 26(2):45–53, 1990.

- Anant Raj and Francis Bach. Explicit regularization of stochastic gradient methods through duality. In *International Conference on Artificial Intelligence and Statistics*, pages 1882–1890, 2021.
- Sashank Reddi, Manzil Zaheer, Suvrit Sra, Barnabas Poczos, Francis Bach, Ruslan Salakhutdinov, and Alex Smola. A generic approach for escaping saddle points. In *International Conference on Artificial Intelligence and Statistics*, pages 1233–1242, 2018.
- Hadi Salman, Jerry Li, Ilya Razenshteyn, Pengchuan Zhang, Huan Zhang, Sebastien Bubeck, and Greg Yang. Provably robust deep learning via adversarially trained smoothed classifiers. *Advances in Neural Information Processing Systems*, 32, 2019.
- Andrew M. Saxe, James L. McClelland, and Surya Ganguli. A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences*, 116(23):11537–11546, 2019.
- Nicholas D. Sidiropoulos, Lieven De Lathauwer, Xiao Fu, Kejun Huang, Evangelos E. Papalexakis, and Christos Faloutsos. Tensor decomposition for signal processing and machine learning. *IEEE Transactions on Signal Processing*, 65(13):3551–3582, 2017.
- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of deep neural networks. *Mathematics of Operations Research*, 47(1):120–152, 2022.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Matthew Staib, Sashank Reddi, Satyen Kale, Sanjiv Kumar, and Suvrit Sra. Escaping saddle points with adaptive gradient methods. In *International Conference on Machine Learning*, pages 5956–5965, 2019.
- 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 *International Conference on Machine Learning*, pages 9636–9647, 2020.
- Tomas Vaskevicius, Varun Kanade, and Patrick Rebeschini. Implicit regularization for optimal sparse recovery. *Advances in Neural Information Processing Systems*, 32, 2019.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Jürgen Weese. A regularization method for nonlinear ill-posed problems. *Computer Physics Communications*, 77(3):429–440, 1993.
- Stephan Wojtowytsch. On the convergence of gradient descent training for two-layer relu-networks in the mean field regime. *arXiv preprint arXiv:2005.13530*, 2020.
- Blake Woodworth, Suriya Gunasekar, Jason D. Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory*, pages 3635–3673, 2020.
- Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. *Advances in Neural Information Processing Systems*, 33: 2958–2969, 2020.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Yaowei Zheng, Richong Zhang, and Yongyi Mao. Regularizing neural networks via adversarial model perturbation. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, pages 8156–8165, 2021.
- Mo Zhou, Tianyi Liu, Yan Li, Dachao Lin, Enlu Zhou, and Tuo Zhao. Toward understanding the importance of noise in training neural networks. In *International Conference on Machine Learning*, pages 7594–7602, 2019.
- Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. The anisotropic noise in stochastic gradient descent: Its behavior of escaping from sharp minima and regularization effects. In *International Conference on Machine Learning*, pages 7654–7663, 2019.

# Explicit Regularization in Overparametrized Models via Noise Injection: Supplementary Materials

## A Proofs of equivalences

Note that in finite-dimensional models, all norms are equivalent, so we use any of them, unless otherwise stated.

### A.1 Proof of Theorem 1

*Proof.* We first observe from Eq. (2) that

$$|R_\sigma(w) - R(w)| = \frac{\sigma^2}{2} DL(\Phi(w)) \underbrace{D^2\Phi(w)[I]}_{=0} + \frac{\sigma^2}{2} D^2L(\Phi(w)) [D\Phi(w)D\Phi(w)^\top] + O(\sigma^3). \quad (12)$$

By using the inequality (due to bounded second derivative of  $\Phi$ )

$$D\Phi(w) \leq D\Phi(0) + C\|w\| \leq C(1 + \|w\|), \quad (13)$$

we obtain the first desired bound

$$|R_\sigma(w) - R(w)| = \left| \frac{\sigma^2}{2} \underbrace{D^2L(\Phi(w))}_{\text{uniformly bounded by Ass.}} [D\Phi(w)D\Phi(w)^\top] \right| \leq C(1 + \|w\|^2)\sigma^2. \quad (14)$$

For the second desired bound, we observe from Eqs. (5) and (6) that

$$|R_\sigma(w) - R_\sigma^{\text{eff}}(w)| = \left| \frac{\sigma^2}{2} DL(\Phi(w)) \underbrace{D^2\Phi(w)[I]}_{=0} + O(\sigma^3) \right|, \quad (15)$$

which concludes the proof. (Note that in this proof the precise value of the constant  $C > 0$  is changing from line to line. It is not necessarily the same as the  $C$  in the statement of the theorem.)  $\square$

### A.2 Proof of Theorem 2

*Proof.* As proved in Eq. (13) under the same assumptions, we again have

$$\left| \frac{\sigma^2}{2} D^2L(\Phi(w)) [D\Phi(w)D\Phi(w)^\top] \right| \leq C(1 + \|w\|_2^2)\sigma^2. \quad (16)$$

(Here, and in the rest of the proof,  $C > 0$  is a constant independent of  $\sigma$  which may change from line to line.) Similarly, by both the boundedness of second derivatives of  $L$  and  $\Phi$ , we have that

$$\left| \frac{\sigma^2}{2} DL(\Phi(w)) D^2\Phi(w)[I] \right| \leq C(\sigma^2(1 + \|\Phi(w)\|_2)). \quad (17)$$

Thus, by putting both bounds together, there exists a constant  $C > 0$  such that for all  $w$ ,

$$\begin{aligned} |R_\sigma(w) - R(w)| &\leq \frac{\sigma^2}{2} D^2L(\Phi(w)) [D\Phi(w)D\Phi(w)^\top] + \frac{\sigma^2}{2} DL(\Phi(w)) D^2\Phi(w)[I] \\ &\leq C\sigma^2(1 + \|\Phi(w)\|_2 + \|w\|_2^2). \end{aligned} \quad (18)$$

Hence, if  $w_*^\sigma$  is a minimizer of  $R_\sigma$ , we have the chain of inequalities

$$\begin{aligned} R(w_*^\sigma) - C\sigma^2(1 + \|\Phi(w_*^\sigma)\|_2 + \|w_*^\sigma\|_2^2) &\leq R_\sigma(w_*^\sigma) \leq R_\sigma(w_*) \\ R_\sigma(w_*) &\leq R(w_*) + C\sigma^2(1 + \|\Phi(w_*)\|_2 + \|w_*\|_2^2). \end{aligned}$$



This leads to

$$\frac{\mu}{2} \|\Phi(w_*^\sigma) - \varphi_*\|_2^2 \leq R(w_*^\sigma) - R(w_*) \leq 2C\sigma^2(2 + \|\Phi(w_*)\|_2 + \|\Phi(w_*^\sigma)\|_2 + \|w_*^\sigma\|_2^2 + \|w_*\|_2^2),$$

by  $\mu$ -strong-convexity of  $L$ . Thus, since  $w_*$  and  $w_*^\sigma$  are in a predefined compact set  $\Omega$  for  $\sigma$  small enough and since  $\Phi(w)$  is bounded on  $\Omega$  due to boundedness of its second derivative, we arrive at:  $\|\Phi(w_*^\sigma) - \varphi_*\|_2^2 = O(\sigma^2)$ . This shows the first statement.

We also consider a minimizer  $w_*^{\sigma,(\text{eff})}$  of  $R_\sigma^{(\text{eff})}$ , which satisfies the same bound as above (exact same reasoning), that is,

$$\|\Phi(w_*^{\sigma,(\text{eff})}) - \varphi_*\|_2^2 \leq C''\sigma^2. \quad (19)$$

Since

$$\|\text{DL}(\Phi(w_*^{\sigma,(\text{eff})}))\| \leq C,$$

due to  $w_*^{\sigma,(\text{eff})} \in \Omega$  (compact set) and

$$\|\text{D}^2\Phi(w_*^{\sigma,(\text{eff})})\| = \|\text{D}^2\Phi(w_*^{\sigma,(\text{eff})}) - \text{D}^2\Phi(\varphi_*)\| \leq C\|\Phi(w_*^{\sigma,(\text{eff})}) - \varphi_*\| \leq C''\sigma^2,$$

due to Eq. (19), we get

$$\frac{\sigma^2}{2} \text{DL}(\Phi(w_*^{\sigma,(\text{eff})})) \text{D}^2\Phi(w_*^{\sigma,(\text{eff})}) [I] \leq C\sigma^3.$$

This implies moreover that

$$|R_\sigma^{(\text{eff})}(w_*^{\sigma,(\text{eff})}) - R_\sigma(w_*^{\sigma,(\text{eff})})| \leq C\sigma^3.$$

Moreover, since  $w_*^\sigma$  minimizes  $R_\sigma$ , the differential of  $R_\sigma$  at  $w_*^\sigma$  is equal to zero, and using a similar expansion than Theorem 1 for the differential, we get  $0 = \text{DL}(\Phi(w_*^\sigma)) + O(\sigma)$ . This leads to  $|R_\sigma^{(\text{eff})}(w_*^\sigma) - R(w_*^\sigma)| \leq C\sigma^3$  due to  $\text{DL}(\Phi(w_*^\sigma)) \leq C\sigma$ . We can thus use the chain of inequalities

$$R(w_*^{\sigma,(\text{eff})}) - C\sigma^3 \leq R_\sigma^{(\text{eff})}(w_*^{\sigma,(\text{eff})}) \leq R_\sigma^{(\text{eff})}(w_*^\sigma) \leq R(w_*^\sigma) + C\sigma^3,$$

and follow the same reasoning as above to obtain the desired bound  $\|\Phi(w_*^{\sigma,(\text{eff})}) - \Phi(w_*^\sigma)\|_2^2 = O(\sigma^3)$ .  $\square$

### A.3 Extension of asymptotic theory from Section 2 to piecewise regularity

In this section, we show how the derivation of  $R_\sigma(w)$ , Eq. (5), and Theorem 1 can be extended to deal with a more general predictor  $\Phi$ . More precisely, we extend it to the case where globally  $\Phi$  is only assumed to be continuous and the additional regularity of  $\Phi$  is only given *piecewise*, i.e., where there exists a finite number of open convex sets  $\{\Omega_i, i = 1, \dots, M\}$  with  $\bigcup_{i \in I} \overline{\Omega}_i = \mathbb{R}^m$  such that  $\Phi$  restricted to  $\Omega_i$  is in  $C^3$  for all  $i$ . This is an important extension because it takes care of locally non-differentiable activation functions, such as ReLU.

**Extension of Eq. (5)** Let  $w \in \Omega_i$  for a fixed  $i = 1, \dots, M$ . Under the above piecewise assumption, the Taylor expansion of  $\Phi$  at  $w$  can now be written:

$$\Phi(w + \sigma\varepsilon) = \mathbb{1}_{w+\sigma\varepsilon \notin \Omega_i} [\Phi(w + \sigma\varepsilon)] \quad (20)$$

$$+ \mathbb{1}_{w+\sigma\varepsilon \in \Omega_i} [\Phi(w) + \sigma \text{D}\Phi(w)\varepsilon + \frac{\sigma^2}{2} \text{D}^2\Phi(w)[\varepsilon\varepsilon^\top] + O(\sigma^3\|\varepsilon\|^3)]. \quad (21)$$

Thus, as in the main text, we have

$$\begin{aligned} L(\Phi(w + \sigma\varepsilon)) &= \mathbb{1}_{w+\sigma\varepsilon \notin \Omega_i} [L(\Phi(w + \sigma\varepsilon))] \\ &+ \mathbb{1}_{w+\sigma\varepsilon \in \Omega_i} \left[ L(\Phi(w)) + \text{DL}(\Phi(w))(\sigma \text{D}\Phi(w)\varepsilon + \frac{\sigma^2}{2} \text{D}^2\Phi(w)[\varepsilon\varepsilon^\top] + O(\sigma^3\|\varepsilon\|^3)) \right. \\ &\quad \left. + \frac{1}{2} \text{D}^2L(\Phi(w))[\sigma \text{D}\Phi(w)\varepsilon(\sigma \text{D}\Phi(w)\varepsilon)^\top + O(\sigma^3\|\varepsilon\|^3)] + O(\sigma^3\|\varepsilon\|^3) \right]. \end{aligned}$$

Now, taking the expectation we get

$$\begin{aligned} R_\sigma(w) &= \mathbb{E}[L(\Phi(w + \sigma\varepsilon))] \\ &= \int_{\{\varepsilon: w + \sigma\varepsilon \notin \Omega_i\}} L(\Phi(w + \sigma\varepsilon)) dN(\varepsilon; 0, I) + \int_{\{\varepsilon: w + \sigma\varepsilon \in \Omega_i\}} L(\Phi(w + \sigma\varepsilon)) dN(\varepsilon; 0, I) \\ &=: J_1(\sigma) + J_2(\sigma), \end{aligned}$$

with the integrands from the above cases of  $L(\Phi(w + \sigma\varepsilon))$ . Now,

$$\lim_{\sigma \rightarrow 0} J_2(\sigma) = L(\Phi(w)) + \frac{\sigma^2}{2} DL(\Phi(w))D^2\Phi(w)[I] + \frac{\sigma^2}{2} D^2L(\Phi(w))[D\Phi(w)D\Phi(w)^\top] + O(\sigma^3)$$

analogously to the main text, and

$$\lim_{\sigma \rightarrow 0} J_1(\sigma) = 0$$

because the probability mass of  $\{\varepsilon : w + \sigma\varepsilon \notin \Omega_i\}$  under  $N(\varepsilon; 0, I)$  goes to zero exponentially fast, while  $L(\Phi)$  can diverge at most polynomially. Thus, we arrive at Eq. (5) from the main text.

**Extension of Theorems 1 and 2** Above, we saw how our Taylor expansions and averaging carry over to the case where  $\Phi$  is only piece-wise in  $C^3$ . Holding this in mind, we can now verify that the above proofs of Theorems 1 and 2 also carry through in this case:

Since both Theorems are formulated only for sufficiently small  $\sigma > 0$ , this can be easily verified. For Theorem 1 this is trivial, one just has to emphasize that  $\sigma > 0$  is small enough after each step. For Theorem 2 one additionally has to make sure that the minimizers  $w_*^\sigma$  and  $w_*^{\sigma,(\text{eff})}$  of  $R(w) = L(\Phi(w))$  and  $R_\sigma^{\text{(eff)}}$  are also in  $\Omega_i$ . But this will hold true for all sufficiently small  $\sigma$  for the following reason: First, for  $w_*^\sigma$  this follows from the fact that  $J_2(\sigma)$  goes to 0 exponentially fast, as  $\sigma \rightarrow 0$ . Hence, only values of  $L(\Phi(w))$  for  $w \in \Omega_i$  matter to determine the minimizer  $w_*^\sigma$ , which will thus lie in  $\Omega_i$ . Second, under the assumptions of Theorem 2, the difference between  $R_\sigma$  and  $R_\sigma^{\text{(eff)}}$  is in  $O(\sigma^3)$ . Thus, for  $\sigma$  small enough,  $w_*^{\sigma,(\text{eff})}$  will lie in the same open convex set as  $w_*^\sigma$ , i.e., in  $\Omega_i$ .

## B Logistic regression

We consider a 2-homogeneous model such that  $\Phi(\lambda w) = \lambda^2\Phi(w)$  for any  $w$  and  $\lambda > 0$ . We aim to minimize

$$\frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i\Phi(w)_i)),$$

where  $y \in \{-1, 1\}^n$ . We assume that the model is overparameterized, so that there exists  $w$  such that  $y_i\Phi(w)_i > 0$  for all  $i$ .

Following Lyu and Li [2019], we expect  $w$  to diverge in some direction, that is,  $w = \lambda\Delta$  for  $\lambda \rightarrow +\infty$  and  $\|\Delta\|_2 = 1$ . We assume that this is the case, and we derive here an informal argument highlighting what the limit direction  $\Delta$  should be.

In the function

$$R_\sigma^{\text{(eff)}}(w) = R(w) + \frac{\sigma^2}{2} D^2L(\Phi(w))[D\Phi(w)D\Phi(w)^\top]$$

from Equation (6), the term  $D\Phi(w)D(\Phi(w))^\top$  grows in  $\lambda$  as  $\lambda^2 D\Phi(\Delta)D(\Phi(\Delta))^\top$ , while  $D^2L(\Phi(w))$  is diagonal and proportional to  $\exp(-\lambda^2|\Phi(\Delta)_i|)$ , with the same scaling as the loss function. We thus get an asymptotic approximate cost function equal to, since  $|\Phi(\Delta)_i| = y_i\Phi(\Delta)_i$  (because we have perfect predictions):

$$\frac{1}{n} \sum_{i=1}^n e^{-y_i\lambda^2\Phi(\Delta)_i} \left(1 + \frac{\sigma^2}{2} \lambda^2 \|D(\Phi(\Delta))_i\|_2^2\right) \sim \frac{1}{n} \sum_{i=1}^n e^{-y_i\Phi(w)_i} \left(1 + \frac{\sigma^2}{2} \|D(\Phi(w))_i\|_2^2\right).$$

By taking the gradient with respect to  $w$ , we get that

$$\sum_{i=1}^n e^{-y_i\Phi(w)_i} \left( \frac{\sigma^2}{2} \frac{\partial}{\partial w} \|D(\Phi(w))_i\|_2^2 - y_i \left(1 + \frac{\sigma^2}{2} \|D(\Phi(w))_i\|_2^2\right) \frac{\partial}{\partial w} \Phi(w)_i \right),$$

which is asymptotically equivalent to

$$\sum_{i=1}^n e^{-y_i \Phi(w)_i} \left( \frac{\sigma^2}{2} \frac{\partial}{\partial w} \|\mathbf{D}(\Phi(w))_i\|_2^2 - y_i \frac{\partial}{\partial w} \Phi(w)_i \right) \propto -w.$$

We conjecture that this is equivalent to the optimality conditions of the problem

$$\max_{\|\Delta\|_2 \leq 1} \min_{i \in \{1, \dots, n\}} \left\{ y_i \Phi(w)_i - \frac{\sigma^2}{2} \|\mathbf{D}(\Phi(w))_i\|_2^2 \right\},$$

where for  $\sigma = 0$ , this is the result of [Lyu and Li \[2019\]](#). We leave a formalization of such a result for future work.

## C Direct derivations of formulas from Sections 2.3 and 2.4

This section contains some explicit derivation of formulas that were asserted in the main text.

### C.1 Lasso

Here, we present the derivation of Equation (7). First we note that  $\mathbf{D}^2 L(\Phi(w)) = \frac{1}{n} I$ , i.e. the identity matrix, and that

$$\mathbf{D}\Phi(w) = [2X \text{diag}(w_1), 2X \text{diag}(w_2)]. \quad (22)$$

Insertion of these two identities into Eq. (6) yields

$$R_\sigma^{\text{(eff)}}(w) = R(w) + \frac{\sigma^2}{2} \mathbf{D}^2 L(\Phi(w)) [\mathbf{D}\Phi(w) \mathbf{D}\Phi(w)^\top] \quad (23)$$

$$= R(w) + 2\sigma^2 \frac{1}{n} I \underbrace{\left[ \begin{array}{cc} X \text{diag}(w_1), & X \text{diag}(w_2) \\ X \text{diag}(w_1) & X \text{diag}(w_2) \end{array} \right]}_{=: J}. \quad (24)$$

By our above notation  $\mathbf{D}^2 L(\varphi)[M] = \sum_{a,b=1}^n \mathbf{D}^2 L(\varphi)_{ab} M_{ab} \in \mathbb{R}$ , we have

$$J = \frac{1}{n} \sum_{a=1}^d [X \text{diag}(w_1^2) X^\top / n]_{aa} + \sum_{a=1}^d [X \text{diag}(w_2^2) X^\top / n]_{aa} := J_1 + J_2. \quad (25)$$

We keep computing:

$$J_1 = \sum_{a=1}^n \sum_{i=1}^d \frac{1}{n} X_{ai} w_1^2 X_{ia}^\top = \sum_{i=1}^d \underbrace{\left( \sum_{a=1}^n \frac{1}{n} X_{ia}^\top X_{ai} \right)}_{[X^\top X]_{ii}} w_1^2 = \text{diag}(X^\top X / n)^\top [w_1 \circ w_1], \quad (26)$$

and analogously  $J_2 = \text{diag}(X^\top X / n)^\top [w_2 \circ w_2]$ . Thus,  $J = \text{diag}(X^\top X / n)^\top [w_1 \circ w_1] + \text{diag}(X^\top X / n)^\top [w_2 \circ w_2]$  which we insert back into Eq. (23). This yields the desired Eq. (7).

### C.2 Nuclear norm (linear networks)

For the main text, only the derivation of Eq. (9) is missing. To this end, we first observe that

$$[W_2 W_1 X^\top]_{ij} = \sum_{k=1}^{d_0} \sum_{l=1}^{d-1} [W_2]_{il} [W_1]_{lk} X_{jk}. \quad (27)$$

Moreover,

$$\mathbf{D}^2 L(\Phi(w)) = \frac{1}{2n} I. \quad (28)$$

Hence,

$$D^2 L(\Phi(w)) [D\Phi(w)D\Phi(w)^\top] = \frac{1}{2n} \sum_{a=1}^{nd_2} [D\Phi(w)D\Phi(w)^\top]_{aa} \quad (29)$$

$$= \frac{1}{2n} \sum_{i=1}^{d_2} \sum_{j=1}^n \left[ \left\| \nabla_w \sum_{k=1}^{d_0} \sum_{l=1}^{d_1} [W_2]_{il} [W_1]_{lk} X_{jk} \right\|_2^2 \right] \quad (30)$$

$$= \frac{1}{2n} \sum_{i=1}^{d_2} \sum_{j=1}^n \left[ \sum_{k=1}^{d_0} \sum_{l=1}^{d_1} ([W_2]_{il} X_{kj}^\top)^2 + \sum_{l=1}^{d_1} \left( \sum_{k=1}^{d_0} [W_1]_{lk} X_{kj}^\top \right)^2 \right] \quad (31)$$

$$= \frac{1}{2n} (J_1 + J_2), \quad (32)$$

with

$$J_1 := \sum_{i=1}^{d_2} \sum_{j=1}^n \sum_{k=1}^{d_0} \sum_{l=1}^{d_1} [W_2]_{il}^2 X_{jk}^2 \quad (33)$$

$$= \left[ \sum_{i=1}^{d_2} \sum_{l=1}^{d_1} [W_2]_{il}^2 \right]^2 \left[ \sum_{j=1}^n \sum_{k=1}^{d_0} X_{jk}^2 \right] \quad (34)$$

$$= \|W_2\|_F^2 \|X\|_F^2 \quad (35)$$

and

$$J_2 := \sum_{i=1}^{d_2} \sum_{j=1}^n \sum_{l=1}^{d_1} [W_2 X^\top]_{lj}^2 = d_2 \|W_1 X^\top\|_F^2. \quad (36)$$

Putting all equations together concludes the derivation of Eq. (9).

### C.3 Group Lasso

**Extension to “group Lasso”.** It is traditional to recover the group Lasso as a special of nuclear norm minimization [see, e.g., Bach, 2008]. We thus consider  $w = (v_1, w_1, \dots, v_d, w_d) \in \mathbb{R}^{d(k+1)}$ , with  $v_j \in \mathbb{R}$  and  $w_j \in \mathbb{R}^k$ , and  $\Phi(w) = v_1 X_1 w_1 + \dots + v_d X_d w_d$ , for  $X_1, \dots, X_d \in \mathbb{R}^{n \times k}$ , and  $L(\varphi) = \frac{1}{2n} \|y - \varphi\|_2^2$ . This corresponds exactly to a linear network defined above, with  $X = (X_1, \dots, X_n)$ ,  $W_1$  defined by blocks and block-diagonal with blocks  $w_j^\top$ , and  $W_2 = v^\top$ .

Regardless whether the model is overparametrized or not, we can apply Theorem 1, and we get:

$$R_\sigma^{\text{(eff)}}(w) = \frac{1}{2n} \|y - v_1 X_1 w_1 - \dots - v_d X_d w_d\|_2^2 + \frac{\sigma^2}{2n} \left[ \|X_1 w_1\|_2^2 + v_1^2 \|X_1\|_F^2 + \dots + \|X_d w_d\|_2^2 + v_d^2 \|X_d\|_F^2 \right]. \quad (37)$$

See the Appendix for a detailed derivation of Eq. (37). Optimizing over the “invariance of scale” (we can multiply  $v_j$  by  $\alpha_j$  and divide  $w_j$  by the same  $\alpha_j$ ), we get the equivalent problem of minimizing

$$\frac{1}{2n} \|y - X_1 \beta_1 - \dots - X_d \beta_d\|_2^2 + \frac{\sigma^2}{n} \left[ \|X_1\|_F \cdot \|X_1 \beta_1\|_2 + \dots + \|X_d\|_F \cdot \|X_d \beta_d\|_2 \right], \quad (38)$$

where  $\beta_j = |v_j| w_j$ , which is a form of group Lasso. In particular if all  $\beta_j$ 's have dimension one, we recover the Lasso with a different formulation from Section 2.3 (and no need for overparametrization).

**Derivation of Equation (37)** First we note that  $D^2 L(\Phi(w)) = \frac{1}{n} I$ , i.e. the identity matrix, and that

$$\Phi(w) = \sum_{i=1}^d \Phi_i(w), \quad \text{where } \Phi_i(v_i, w_i) := v_i X_i w_i. \quad (39)$$



Under this notation, we get the Jacobian

$$D\Phi(w) = \begin{bmatrix} D\Phi_1(w) & & \\ & \ddots & \\ & & D\Phi_d(w) \end{bmatrix}, \quad (40)$$

with

$$D\Phi_i(v_i, w_i) = [X_i w_i, v_i X_i] \in \mathbb{R}^{n \times (k+1)}. \quad (41)$$

Now,

$$D\Phi_i(v_i, w_i) D\Phi_i(v_i, w_i)^\top = [X_i w_i, v_i X_i] \begin{bmatrix} w_i^\top X_i^\top \\ v_i X_i^\top \end{bmatrix} = [X_i w_i w_i^\top X_i^\top, v_i^2 X_i X_i^\top], \quad (42)$$

which implies that

$$D\Phi(w) \cdot D\Phi(w)^\top = \begin{bmatrix} D\Phi_1(w) D\Phi_1(w)^\top & & \\ & \ddots & \\ & & D\Phi_d(w) D\Phi_d(w)^\top \end{bmatrix} \quad (43)$$

$$= \begin{bmatrix} [X_1 w_1 w_1^\top X_1^\top, v_1^2 X_1 X_1^\top] & & \\ & \ddots & \\ & & [X_d w_d w_d^\top X_d^\top, v_d^2 X_d X_d^\top] \end{bmatrix}. \quad (44)$$

By our above notation  $D^2 L(\varphi)[M] = \sum_{a,b=1}^n D^2 L(\varphi)_{ab} M_{ab} \in \mathbb{R}$ , we now get the desired formula:

$$R_\sigma^{\text{eff}}(w) = R(w) + \frac{\sigma^2}{2} D^2 L(\Phi(w)) [D\Phi(w) D\Phi(w)^\top] \quad (45)$$

$$= R(w) + \frac{\sigma^2}{2n} I [D\Phi(w) \cdot D\Phi(w)^\top] \quad (46)$$

$$= R(w) + \frac{\sigma^2}{2n} [\|X_1 w_1\|_2^2 + v_1^2 \|X_1\|_F^2 + \dots + \|X_d w_d\|_2^2 + v_d^2 \|X_d\|_F^2]. \quad (47)$$

## D Appendix for Section 3 (one-hidden layer)

### D.1 Explosion of full perturbations

We have discussed the exploding variance phenomena in one hidden layer linear networks in Section 3.1. We here discuss the variance explosion in one hidden layer neural network with ReLU activations. Similarly, we here consider the overparametrized limit  $d_1 \rightarrow +\infty$  with initialization which corresponds to having weights of order  $(W_1)_{ij} \sim \frac{1}{\sqrt{d_1 d_0}}$  for all  $i, j$ , that is  $\|W_1\|_F^2$  not exploding with  $d_1$ , and  $(W_2)_{ij} \sim \frac{1}{\sqrt{d_2 d_1}}$ , that is  $\|W_2\|_F^2$  not exploding with  $d_1$ . This initialization corresponds to the one in [Glorot and Bengio \[2010\]](#) in the case where the network was constant width. We have

$$\Phi(w) = W_2(W_1 X^\top)_+,$$

where the positive part is taken element-wise.

We have, with Gaussian perturbations  $E_1$  and  $E_2$ , an explicit exact expansion for one hidden layer ReLU networks for small  $\sigma$  (so that  $\Phi$  is locally quadratic):

$$\begin{aligned} & \Phi(w + \sigma \varepsilon) \\ &= \Phi(W_2 + \sigma E_2, W_1 + \sigma E_1) = (W_2 + \sigma E_2) [(W_1 + \sigma E_1) X^\top]_+ \\ &= (W_2 + \sigma E_2) (W_1 X^\top)_+ + \sigma (W_2 + \sigma E_2) [(W_1 X^\top)_+^0 \circ E_1 X^\top] \\ &= W_2 (W_1 X^\top)_+ + \sigma E_2 (W_1 X^\top)_+ + \sigma W_2 [(W_1 X^\top)_+^0 \circ E_1 X^\top] + \sigma^2 E_2 [(W_1 X^\top)_+^0 \circ E_1 X^\top] \\ &= \Phi(w) + \sigma (E_2 (W_1 X^\top)_+ + W_2 [(W_1 X^\top)_+^0 \circ E_1 X^\top]) + \sigma^2 E_2 [(W_1 X^\top)_+^0 \circ E_1 X^\top]. \end{aligned}$$

Taking expectations and using that  $E_1, E_2$  have zero mean and are independent, and such that,  $\mathbb{E}[E_i M E_i^\top] = \text{tr}(M) I$  for  $i = 1, 2$ , and  $M$  any symmetric matrix of compatible size, we can get  $\mathbb{E}[\Phi(w + \sigma \varepsilon)] = \Phi(w)$ , and:

$$\begin{aligned} & \mathbb{E}[\|\Phi(w + \sigma\varepsilon)\|_F^2] \\ &= \|\Phi(w)\|_F^2 + \sigma^2\|(W_1X^\top)_+\|_F^2 + \sigma^2\sum_{j=1}^{d_1}\sum_{i=1}^n\|(W_2)_{\cdot j}\|_2^2 \times |((W_1X^\top)_+^0)_{ji}|^2 \times \|X_i\|_2^2 \\ & \quad + \sigma^4\mathbb{E}\langle E_2[(W_1X^\top)_+^0 \circ E_1X^\top], E_2[(W_1X^\top)_+^0 \circ E_1X^\top] \rangle. \end{aligned}$$

Consider a case when  $(W_1X^\top)_+^0 = 1$  at all of its entries. In that case, We can now compute  $R_\sigma$  as:

$$\begin{aligned} R_\sigma(W_1, W_2) &= R(W_1, W_2) + \frac{\sigma^2}{2n} [\|(W_1X^\top)_+\|_F^2 \\ & \quad + \frac{\sigma^2}{2n} \sum_{j=1}^{d_1} \sum_{i=1}^n \|(W_2)_{\cdot j}\|_2^2 \times |((W_1X^\top)_+^0)_{ji}|^2 \times \|X_i\|_2^2 + \frac{\sigma^4}{2n} d_1 d_2 \|X\|_F^2]. \end{aligned}$$

The extra term  $\frac{\sigma^4}{2n} d_2 d_1 \|X\|_F^2$  is of superior order in  $\sigma$ , but problematic when  $d_1 \rightarrow +\infty$  using same argument as that for linear net with one hidden layer.

## D.2 Layer-wise perturbation

Now, we show that layer wise perturbation helps in variance control. Consider the following two cases :

$$\begin{aligned} \Phi(W_2, W_1 + \sqrt{2}\sigma E_1) &= W_2[(W_1 + \sqrt{2}\sigma E_1)X^\top]_+ = W_2(W_1X^\top)_+ + \sqrt{2}\sigma W_2[(W_1X^\top)_+^0 \circ E_1X^\top] \\ &= \Phi(W_2, W_1) + \sqrt{2}\sigma W_2[(W_1X^\top)_+^0 \circ E_1X^\top]. \end{aligned}$$

Similarly,

$$\begin{aligned} \Phi(W_2 + \sqrt{2}\sigma E_2, W_1) &= (W_2 + \sqrt{2}\sigma E_2)(W_1X^\top)_+ = W_2(W_1X^\top)_+ + \sqrt{2}\sigma E_2(W_1X^\top)_+ \\ &= \Phi(W_2, W_1) + \sqrt{2}\sigma E_2(W_1X^\top)_+. \end{aligned}$$

Hence,

$$\mathbb{E}\|\Phi(W_2, W_1 + \sqrt{2}\sigma E_1)\|_F^2 = \|\Phi(W_2, W_1)\|_F^2 + 2\sigma^2\sum_{j=1}^{d_1}\sum_{i=1}^n\|(W_2)_{\cdot j}\|_2^2 \times |((W_1X^\top)_+^0)_{ji}|^2 \times \|X_i\|_2^2].$$

Similarly,

$$\mathbb{E}\|\Phi(W_2 + \sqrt{2}\sigma E_2, W_1)\|_F^2 = \|\Phi(W_2, W_1)\|_F^2 + 2\sigma^2\|(W_1X^\top)_+\|_F^2.$$

If we choose uniformly at random the layer to be perturbed, then one can then see that<sup>5</sup>

$$\begin{aligned} R_\sigma(W_1, W_2) &= R(W_1, W_2) + \frac{\sigma^2}{2n} [\|(W_1X^\top)_+\|_F^2 \\ & \quad + \frac{\sigma^2}{2n} \sum_{j=1}^{d_1} \sum_{i=1}^n \|(W_2)_{\cdot j}\|_2^2 \times |((W_1X^\top)_+^0)_{ji}|^2 \times \|X_i\|_2^2]. \end{aligned}$$

Due to same argument as that for one hidden layer linear net,  $R_\sigma(W_1, W_2)$  is well behaved.

## D.3 Function space

Here we change slightly the notations to be closer to notations from [Chizat and Bach \[2020\]](#).

Given some data  $x_1, \dots, x_n$ , we consider the function  $\psi(a, b) = a(b^\top x_i)_+$ , with  $(a, b) \in \mathbb{R}^{d+1}$ , and  $\Phi(a_1, b_1, \dots, a_m, b_m) = \frac{1}{m} \sum_{i=1}^n \psi(a_i, b_i)$ , with all weights with unit scale. This corresponds exactly to the one-hidden layer neural network with Glorot initialization, where  $m = d_1$ .

<sup>5</sup>Note that in the main paper, there is a typo, and the term  $\|X_i\|_2^2$  was missing.

We consider  $L(\varphi) = \frac{1}{2n} \|y - \varphi\|_2^2$ , and the equivalent cost function that we obtained is exactly

$$\frac{1}{2n} \sum_{i=1}^n \left( y_i - \frac{1}{m} \sum_{j=1}^m a_j (b_j^\top x_i)_+ \right)^2 + \frac{\sigma^2}{2nm} \sum_{i=1}^n \sum_{j=1}^m \left\{ (b_j^\top x_i)_+^2 + a_j^2 \|x_i\|_2^2 (b_j^\top x_i)_+^0 \right\}.$$

We can optimize over the scale  $\lambda_j > 0$  so that  $a_j \rightarrow a_j \lambda$  and  $b_j \rightarrow b_j / \lambda$ , leading to

$$\frac{1}{2n} \sum_{i=1}^n \left( y_i - \frac{1}{m} \sum_{j=1}^m a_j (b_j^\top x_i)_+ \right)^2 + \frac{\sigma^2}{m} \sum_{j=1}^m \sqrt{\frac{1}{n} \sum_{i=1}^n |a_j|^2 \|x_i\|_2^2 (b_j^\top x_i)_+^0} \sqrt{\frac{1}{n} \sum_{i=1}^n (b_j^\top x_i)_+^2},$$

that is,

$$\frac{1}{2} \mathbb{E} \left( y - \int a (b^\top x)_+ d\mu(a, b) \right)^2 + \sigma^2 \int |a| \sqrt{\mathbb{E} \|x\|_2^2 (b^\top x)_+^0} \sqrt{\mathbb{E} (b^\top x)_+^2} d\mu(a, b),$$

where  $\mathbb{E}$  is the empirical expectation over the data and  $d\mu(a, b) = \frac{1}{m} \sum_{j=1}^m \delta_{(a_j, b_j)}$ . Denoting  $d\nu(b) = a d\mu(a, b)$ , we get the following cost function with respect to  $\nu$ :

$$F_\sigma(\nu) = \frac{1}{2} \mathbb{E} \left( y - \int (b^\top x)_+ d\nu(b) \right)^2 + \sigma^2 \int \sqrt{\mathbb{E} \|x\|_2^2 (b^\top x)_+^0} \sqrt{\mathbb{E} (b^\top x)_+^2} |d\nu(b)|.$$

We thus get a specific  $\ell_1$ -penalty, where the difference with the variation norm of [Kurkova and Sanguineti \[2001\]](#), [Bach \[2017\]](#) is the presence of the data-dependent multiplier  $\sqrt{\mathbb{E} \|x\|_2^2 (b^\top x)_+^0} \sqrt{\mathbb{E} (b^\top x)_+^2}$ , rather than simply having a constant. With sufficiently many data points, these expectations are bounded from below and above, so the regularization properties are the same, with the same adaptivity to linear substructures highlighted by [Chizat and Bach \[2020\]](#).

## E Appendix for Section 4 (deep networks)

### E.1 Explosion of full perturbations

#### E.1.1 Linear Network

Consider a linear neural network with  $w = (W_1, \dots, W_M)$  with  $W_i \in \mathbb{R}^{d_i \times d_{i-1}}$ , and  $\Phi(w) = W_M W_{M-1} \dots W_1 X^\top \in \mathbb{R}^{d_M \times n}$  with the input data  $X \in \mathbb{R}^{n \times d_0}$ , with  $L(\varphi) = \frac{1}{2n} \|Y^\top - \varphi\|_F^2$  for output data  $Y \in \mathbb{R}^{n \times d_M}$  and  $\varphi \in \mathbb{R}^{d_M \times n}$ . Let us denote the noise with  $\varepsilon = (E_1, E_2, \dots, E_M)$ . Let us now compute,

$$\begin{aligned} \Phi(w + \sigma\varepsilon) &= \Phi(W_M + \sigma E_M, \dots, W_1 + \sigma E_1) = (W_M + \sigma E_M) \dots (W_1 + \sigma E_1) X^\top \\ &= W_M \dots W_1 X^\top + \sigma^M E_M \dots E_1 X^\top + \sum_{i=1}^{M-1} \sigma^i C_i \\ &= \Phi(W_M, \dots, W_1) + \sigma^M E_M \dots E_1 X^\top + \sum_{i=1}^{M-1} \sigma^i C_i, \end{aligned}$$

where  $C_i$  depends on  $W_M, \dots, W_1$  and  $E_M, \dots, E_1$ . It is easy to check that  $\mathbb{E}[\Phi(w + \sigma\varepsilon)] = \Phi(w)$ . To show that  $R_\sigma$  explodes we will show that  $\|\Phi(w + \sigma\varepsilon)\|_F^2$  explodes. Because of independence of  $E_1, \dots, E_M$ , we have

$$\mathbb{E} \|\Phi(w + \sigma\varepsilon)\|_F^2 = \|\Phi(W_M, \dots, W_1)\|_F^2 + \sigma^{2M} \mathbb{E} \|E_M \dots E_1 X^\top\|_F^2 + \sum_{i=1}^{M-1} \sigma^{2i} \mathbb{E} \|C_i\|_F^2.$$

Consider the term

$$\sigma^{2M} \mathbb{E} \|E_M \dots E_1 X^\top\|_F^2 = \sigma^{2M} d_1 d_2 \dots d_M \|X\|_F^2.$$

In the above we used the fact that  $\mathbb{E}[E_i M E_i^\top] = \text{tr}(M)I$ . As  $d_i \rightarrow \infty$  for  $i \in \{1, \dots, M-1\}$ ,  $\sigma^{2M} \mathbb{E} \|E_M \dots E_1 X^\top\|_F^2$  diverges and hence  $R_\sigma$  explodes as well. Using the same argument, we can show that all coefficients  $\sigma^{2i}$  for all  $i > 1$  explode.

### E.1.2 ReLU Network

Consider a deep neural network with  $w = (W_1, \dots, W_M)$  with  $W_i \in \mathbb{R}^{d_i \times d_{i-1}}$ , and  $\Phi(w) = W_M W_{M-1} \dots W_1 X^\top \in \mathbb{R}^{d_M \times n}$  with the input data  $X \in \mathbb{R}^{n \times d_0}$ , with  $L(\varphi) = \frac{1}{2n} \|Y^\top - \varphi\|_F^2$  for output data  $Y \in \mathbb{R}^{n \times d_M}$  and  $\varphi \in \mathbb{R}^{d_M \times n}$ . We use ReLU activation here. Let us denote the noise with  $\varepsilon = (E_1, E_2, \dots, E_M)$ . We have,

$$\Phi(W_M, \dots, W_1) = W_M(W_{M-1} \dots (W_1 X^\top)_+ \dots)_+$$

Now, let us compute

$$\begin{aligned} \Phi(W_M + \sigma E_M, \dots, W_1 + \sigma E_1) &= (W_M + \sigma E_M)((W_{M-1} + \sigma E_{M-1}) \dots ((W_1 + \sigma E_1) X^\top)_+ \dots)_+ \\ &= W_M(W_{M-1} \dots (W_1 X^\top)_+ \dots)_+ + \sum_{i=1}^{M-1} \sigma^i C_i \\ &\quad + \sigma^M E_M((W_{M-1} \dots (W_1 X^\top)_+ \dots)_+^0 \circ E_{M-1}(\dots (W_1 X^\top)_+^0 \circ E_1 X^\top)), \end{aligned}$$

where  $C_i$  depends on  $W_M, \dots, W_1$  and  $E_M, \dots, E_1$ . Using similar argument as for linear network,

$$\begin{aligned} \mathbb{E} \|\Phi(w + \sigma \varepsilon)\|_F^2 &= \|\Phi(w)\|_F^2 + \sigma^{2M} \mathbb{E} \|E_M((W_{M-1} \dots (W_1 X^\top)_+ \dots)_+^0 \circ E_{M-1}(\dots (W_1 X^\top)_+^0 \circ E_1 X^\top))\|_F^2 \\ &\quad + \sum_{i=1}^{M-1} \sigma^{2i} \mathbb{E} \|C_i\|_F^2. \end{aligned}$$

Let us again consider a special case and assume that  $w = (W_M, \dots, W_1)$  are such that  $W_1 X^\top > 0$  at all entries and  $W_2, \dots, W_M$  has all positive entries. In such scenario, for small enough  $\sigma$

$$\begin{aligned} \sigma^{2M} \mathbb{E} \|E_M((W_{M-1} \dots (W_1 X^\top)_+ \dots)_+^0 \circ E_{M-1}(\dots (W_1 X^\top)_+^0 \circ E_1 X^\top))\|_F^2 \\ = \sigma^{2M} \mathbb{E} \|E_M \dots E_1 X^\top\|_F^2 = \sigma^{2M} d_1 \dots d_M \|X\|_F^2. \end{aligned}$$

In the above we used the fact that  $\mathbb{E}[E_i M E_i^\top] = \text{tr}(M)I$ . As  $d_i \rightarrow \infty$  for  $i \in \{1, \dots, M-1\}$ ,  $\sigma^{2M} \mathbb{E} \|E_M \dots E_1 X^\top\|_F^2$  diverges and hence  $R_\sigma$  explodes as well. We can use similar constructions to show that all coefficients  $\sigma^{2i}$  for  $i > 1$  explode.

## E.2 Layer-wise perturbation

In this section, we will show that how layer wise perturbations leads to the desired result. We then proceed to show that the term is not exploding.

The perturbation argument follows exactly the same line as Section [D.2](#) using that neural networks with ReLU activations or without activations are positively homogeneous with respect to each layer.

We thus simply need to show that the extra term is not exploding. For linear networks, the regularization term is

$$\frac{\sigma^2}{2n} \sum_{j=1}^M \|W_M \dots W_{j+1}\|_F^2 \|W_{j-1} \dots W_1 X^\top\|_F^2,$$

and with scaling described in the paper, we get that each element of  $W_M \dots W_{j+1}$  is the product of  $d_{M-1} \dots d_{j+1}$  terms, each of scale [6](#)  $\frac{1}{\sqrt{d_M d_{M-1}}} \dots \frac{1}{\sqrt{d_{j+1} d_j}}$ , thus we an overall scale of

$$d_{M-1} \dots d_{j+1} \times \frac{1}{\sqrt{d_M d_{M-1}}} \dots \frac{1}{\sqrt{d_{j+1} d_j}} = \frac{1}{\sqrt{d_M d_j}}.$$

Then, the Frobenius norm of  $W_M \dots W_{j+1}$  is of order  $d_M d_j \times \left(\frac{1}{\sqrt{d_M d_j}}\right)^2 = 1$ . The reasoning is the same for  $\|W_{j-1} \dots W_1 X^\top\|_F^2$  and leads to no explosion.

With ReLU activation functions, since the derivative of the ReLU activation is either 0 or 1, this does not change the scalings.

<sup>6</sup>We recall that this initialization, i.e. each weight between layer  $i$  and  $i-1$  with variance  $1/\sqrt{d_i d_{i-1}}$ , which stabilizes the expected Frobenius norm of the weight matrices.






Figure 7: Fashion MNIST MLP 1: MLP with 4 Layers (3 hidden), width 500. Full Batch.

## F Further experimental details

We provide a thorough overview of the experimental results presented in this work. Some of the plots of the main paper are also presented again here for ease of comparison.

**Datasets and networks.** We use the Fashion MNIST [Xiao et al., 2017] and CIFAR10 [Krizhevsky and Hinton, 2009] datasets and train on different neural network models (6 experiments in total).

**Setup.** We repeat all experiments 3 times (different network initializations and injected noise) and show mean and 1 standard deviation. All experiments are conducted in PyTorch [Paszke et al., 2017] on up to 8 Tesla V100 GPUs with 32 GB memory.

**Noise injection.** We perform noise injection in two different ways, as described in the main paper. In *GD (SGD) +noise* at each iteration we perturb all the network weights with Gaussian noise with standard deviation  $\sigma/\sqrt{M}$ , where  $M$  is the number of layers. Instead, in *GD (SGD) +noise (layer)*, at each iteration, a specific layer<sup>7</sup> is picked and all weights are perturbed with Gaussian noise with standard deviation  $\sigma$ . As explained in the main text (see Section 3.2&4), we expect the layer-wise perturbation to provide a more stable approach for the minimization of the regularized loss, as width increases. In the results below, we show that this is indeed the case and provide detailed comparisons.

### F.1 Experiments on Fashion MNIST MLPs

**Batch-size, learning rate and noise injection strength.** For the simulations on Fashion MNIST, we use a random subsample of the dataset (1024 samples), to introduce a sizable train-test gap. This also allows us to train with full-batch gradient descent on this subset. Performance of stochastic gradient descent (i.e. mini-batch case) is evaluated for CNNs in the next subsection. We selected for all experiments a sizeable learning rate of  $5e-3$ , except for the Deep and Wide case (MLP 4), where a smaller learning rate is necessary to avoid instabilities. Since we have no access to the problem Lipschitz constant, to allow methods to converge to potentially sharp regions, we use cosine annealing [Loshchilov and Hutter, 2016] for the learning rate. Regarding  $\sigma$ , we select the value that showcases best the implications of our theory: picking a value around 0.05 works in all settings. A plot showing the effect of tuning  $\sigma$  on both noise injection schemes can be found in Section 5 in the main paper. Test performance is evaluated on the full 10K testing dataset.

**Fashion MNIST MLP 1 (Shallow Narrow).** The input ( $28 \times 28 = 784$ -dimensional) is processed with the following 4 narrow layers with ReLU activations (except last layer) to an output with dimension 10:

$$784 \rightarrow 500 \rightarrow 500 \rightarrow 500 \rightarrow 10. \quad (48)$$

The final output is then processed using a log-softmax and signal is backpropagated using a cross-entropy loss. We run full-batch gradient descent on the 1024 datapoints (subset of the data) and train with learning rate  $5e-3$  and  $\sigma = 5e-2$ . Figure 7 shows that, since the network is not too wide, both the noise injection schemes are able to decrease the Hessian trace and improve test accuracy. Yet, interestingly, layer-wise noise injection outperforms standard noise injection in test accuracy.

**Fashion MNIST MLP 2 (Shallow Wide).** The 784-dimensional input is processed with the following 4 wide layers with ReLU activations (except last layer) to an output with dimension 10:

$$784 \rightarrow 5000 \rightarrow 5000 \rightarrow 5000 \rightarrow 10. \quad (49)$$

<sup>7</sup>In the implementation we injected noise sequentially through the network, and we actually perturb separately each parameter group (i.e. half iterations noise is injected on specific layer biases) and pick  $M$  to be the number of parameter groups.




Figure 8: Fashion MNIST MLP 2: MLP with 4 Layers (3 hidden), width 5000. Full Batch.




Figure 9: Fashion MNIST MLP 3: MLP with 6 Layers (5 hidden), width 1000. Full Batch.

The final output is then processed as before, and we select a stepsize  $5e - 3$  and  $\sigma = 1e - 1$ . Figure 8 shows that, since the network is wide, only layer-wise injection provides a successful regularization. These findings are complemented by Figure 3 in the main paper.

**Fashion MNIST MLP 3 (Deep Narrow).** 784-dimensional input is processed with the following 6 not extremely wide layers with ReLU activations (except last layer) to an output with dimension 10:

$$784 \rightarrow 1000 \rightarrow 1000 \rightarrow 1000 \rightarrow 1000 \rightarrow 1000 \rightarrow 10. \tag{50}$$

The final output is then processed as before, and we select a stepsize  $5e - 3$  and  $\sigma = 1e - 1$ . Figure 9 shows that the behavior is similar to the shallow case (Figure 7).

**Fashion MNIST MLP 4 (Deep Wide).** 784-dimensional input is processed with the following 6 wide layers with ReLU activations (except last layer) to an output with dimension 10:

$$784 \rightarrow 5000 \rightarrow 5000 \rightarrow 5000 \rightarrow 5000 \rightarrow 5000 \rightarrow 10. \tag{51}$$

The final output is then processed as before, and we select a stepsize  $1e - 3$  and  $\sigma = 5e - 2$ . Figure 10 shows that the behavior is similar to the narrow case (Figure 8): as the width increases, only layer-wise injection provides an improvement in test accuracy, while standard noise injection hurts convergence at the same value of  $\sigma$  (scaled properly).

## F.2 Experiments on CIFAR10 CNNs

**Batch-size and learning rate, noise injection strength.** For the simulations on CIFAR10, we use the full training dataset (50K samples) and train with SGD with batch-size 1024 and cosine annealing on a tuned learning rate.  $\sigma$  is picked




Figure 10: Fashion MNIST MLP 4: MLP with 6 Layers (5 hidden), width 5000. Full Batch.

to be 0.05 in all settings and layer-wise noise injection is performed with  $\sqrt{M}$  scaling compared to the standard case, where  $M$  is the number of parameter groups in the network. Test performance is evaluated on the full 10K testing dataset.

**CIFAR10 Toy Narrow CNN+MLP.** The  $28 \times 28 \times 3$  image (RGB) is processed by 2 convolutions with 5 channels followed by an MLP. Both Conv and Lin layers have ReLU activations (except last layer) and max pooling is applied after each convolution after the ReLU. Flattening of the tensor is performed before the MLP. The layers are reported<sup>8</sup> next:

$$\text{Conv}(3, 6, 5), \text{Conv}(6, 16, 5), \text{Lin}(16 \times 5 \times 5, 120), \text{Lin}(120, 84), \text{Lin}(84, 10). \quad (52)$$

Backpropagated gradients are computed using a cross-entropy loss. For this network, we train with SGD with batch-size 1024, using a stepsize of  $1e-3$  with cosine annealing. We select here  $\sigma = 5e-2$ . In Figure 11 we see that, while the test performance of SGD with batch size 1024 degrades overtime (known effect of linear layers), the performance of all noise injection schemes leads to better and stable generalization. This is reflected in the regularized Hessian trace. We note that here the performances of the noise injection methods are comparable since the network is not very wide.

**CIFAR10 Wider fully convolutional.** The  $28 \times 28 \times 3$  image (RGB) is processed by a fully convolutional network with up to 128 channels (i.e. wide) with max pooling and ReLU activations (except last layer). The last layer acts on the average output of each channel. The network operations are as follows:

$$\text{Conv}(3, 32, 3), \text{Conv}(32, 64, 3), \text{Conv}(64, 128, 3), \text{Conv}(128, 128, 3), \text{Lin}(128, 10) \quad (53)$$

For this network, we train with SGD with batch-size 1024, using a stepsize of  $5e-3$  with cosine annealing. We select here  $\sigma = 5e-2$ . Figure 12 shows that, in contrast to the toy narrow CNN we inspected above, here layer-wise noise injection provides an improved performance in terms of test loss over standard noise injection. Yet, compared to the MLP case, here injecting noise to all weights does not lead to a quick performance degradation — to observe this, one has to design a network with unrealistically big number of channels. Regarding the Hessian trace, we observe that injecting noise dampens the oscillations in the dynamics of this quantity, which is in line with our result on minimization of a regularized loss (see Section 3.2), even though this was derived in the MLP case.




Figure 11: CIFAR10, PyTorch Toy CNN+MLP. Batch size 1024.




Figure 12: CIFAR10, Wide CNN (128 channels). Batch size 1024.

<sup>8</sup>by  $\text{Conv}(C_{in}, C_{out}, k)$  we mean a convolutional layer with kernel size  $k$ , number of input channels  $C_{in}$  and number output channels  $C_{out}$