

Mathematics of Data: From Theory to Computation

Prof. Volkan Cevher
volkan.cevher@epfl.ch

Lecture 8: Double descent curves and overparametrization

Laboratory for Information and Inference Systems (LIONS)
École Polytechnique Fédérale de Lausanne (EPFL)

EE-556 (Fall 2021)



License Information for Mathematics of Data Slides

- ▶ This work is released under a [Creative Commons License](#) with the following terms:
- ▶ **Attribution**
 - ▶ The licensor permits others to copy, distribute, display, and perform the work. In return, licensees must give the original authors credit.
- ▶ **Non-Commercial**
 - ▶ The licensor permits others to copy, distribute, display, and perform the work. In return, licensees may not use the work for commercial purposes – unless they get the licensor's permission.
- ▶ **Share Alike**
 - ▶ The licensor permits others to distribute derivative works only under a license identical to the one that governs the licensor's work.
- ▶ [Full Text of the License](#)

Outline


This lecture :

- ▶ The classical trade-off between model complexity and risk
- ▶ The generalization mystery in deep learning
- ▶ Implicit regularization of optimization algorithms
- ▶ Double descent curves
- ▶ Generalization bounds based on algorithmic stability

Next lecture :

- ▶ Optimization in Deep Learning

Understanding the trade-off between model complexity and expected risk



Models

Let $[\mathcal{X}_i : i = 1, \dots]$ be a nested sequence of parameter domain, i.e., $\mathcal{X}_i \subseteq \mathcal{X}_{i+1}$. For example, let \mathcal{X}_i = neural networks with i neurons.

1. $R_n(\mathbf{x}_i^*) = \min_{\mathbf{x} \in \mathcal{X}_i} R_n(\mathbf{x})$: ERM solution over \mathcal{X}_i
2. $R(\mathbf{x}_i^*)$: True risk of the ERM solution over \mathcal{X}_i
3. $\sup_{\mathbf{x} \in \mathcal{X}_i} |R(\mathbf{x}) - R_n(\mathbf{x})|$: Worst-case Generalization error of \mathcal{X}_i

Practical performance of the ERM estimator

$$R(\mathbf{x}_i^*) \leq \min_{\mathbf{x} \in \mathcal{X}_i} R_n(\mathbf{x}) + \sup_{\mathbf{x} \in \mathcal{X}_i} |R(\mathbf{x}) - R_n(\mathbf{x})| \quad (1)$$

As we increase the index $i \rightarrow i + 1$ of the parameter domain, i.e., we choose a larger (more complex) model

1. The minimum empirical risk decreases $\min_{\mathbf{x} \in \mathcal{X}_i} R_n(\mathbf{x}) \geq \min_{\mathbf{x} \in \mathcal{X}_{i+1}} R_n(\mathbf{x})$.
2. The generalization error increases. $\sup_{\mathbf{x} \in \mathcal{X}_i} |R(\mathbf{x}) - R_n(\mathbf{x})| \leq \sup_{\mathbf{x} \in \mathcal{X}_{i+1}} |R(\mathbf{x}) - R_n(\mathbf{x})|$.
3. What happens with the true risk $R(\mathbf{x}_i^*)$?

The classical trade-off between model complexity and risk




Figure: Bias-variance trade-off [9].

Occam's Razor: Simple is better than complex.

The dangers of complex function classes: sévère (cevher) overfitting

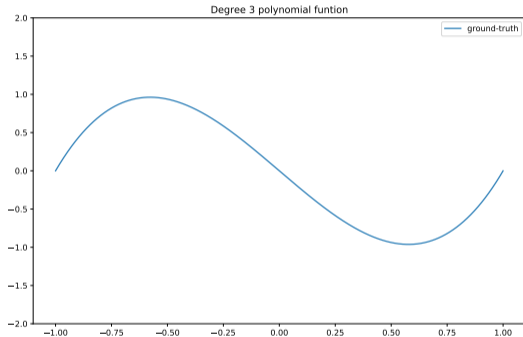
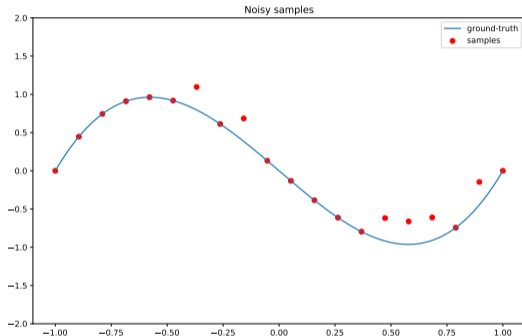


Figure: Training over a complex function class can lead to overfitting.

The dangers of complex function classes: sévère (cevher) overfitting




$$\min_{\mathbf{x} \in \mathcal{X}} R_n(\mathbf{x}) \nearrow$$

$$\sup_{\mathbf{x} \in \mathcal{X}} |R(\mathbf{x}) - R_n(\mathbf{x})| \searrow$$

Figure: Training over a complex function class can lead to overfitting.

The dangers of complex function classes: sévère (cevher) overfitting



$$\min_{\mathbf{x} \in \mathcal{X}} R_n(\mathbf{x}) \searrow$$

$$\sup_{\mathbf{x} \in \mathcal{X}} |R(\mathbf{x}) - R_n(\mathbf{x})| \nearrow$$

Figure: Training over a complex function class can lead to overfitting.

The complexity vs risk trade-off in practice (I)




Figure: Training (empirical) and test (true) error for one-hidden-layer networks of increasing width, trained with SGD [20].

Empirical error becomes zero for a wide enough network. What should happen for even wider networks?

The complexity vs risk trade-off in practice (II)

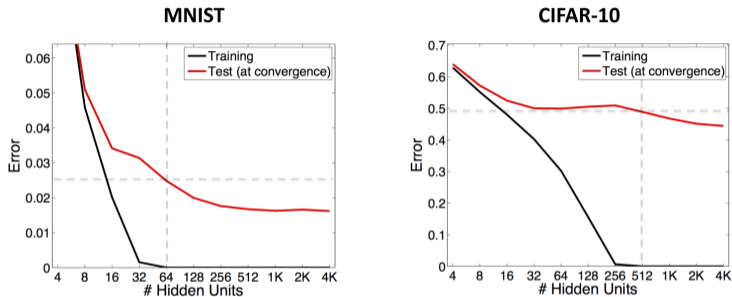



Figure: Training (empirical) and test (true) error for one-hidden-layer networks of increasing width, trained with SGD [20].

Test error continues to go down even if we keep increasing the complexity of the model!

The benefits of overparametrization



$$\min_{\mathbf{x} \in \mathcal{X}} R_n(\mathbf{x}) \searrow$$

$$\sup_{\mathbf{x} \in \mathcal{X}} |R(\mathbf{x}) - R_n(\mathbf{x})| \searrow$$

Figure: Overparametrization leads to benign overfitting.

The generalization mystery in deep learning

UNDERSTANDING DEEP LEARNING REQUIRES RE-THINKING GENERALIZATION

Chiyuan Zhang*

Massachusetts Institute of Technology
chiyuan@mit.edu

Samy Bengio

Google Brain
bengio@google.com

Moritz Hardt

Google Brain
mrtz@google.com

Benjamin Recht†

University of California, Berkeley
brecht@berkeley.edu

Oriol Vinyals

Google DeepMind
vinyals@google.com

ABSTRACT


Despite their massive size, successful deep artificial neural networks can exhibit a remarkably small difference between training and test performance. Conventional wisdom attributes small generalization error either to properties of the model family, or to the regularization techniques used during training.

Through extensive systematic experiments, we show how these traditional approaches fail to explain why large neural networks generalize well in practice. Specifically, our experiments establish that state-of-the-art convolutional networks for image classification trained with stochastic gradient methods easily fit a random labeling of the training data. This phenomenon is qualitatively unaffected by explicit regularization, and occurs even if we replace the true images by completely unstructured random noise. We corroborate these experimental findings with a theoretical construction showing that simple depth two neural networks al-

A gap between theory and practice


- In practice, simple algorithms like SGD can train neural networks to zero error *and* achieve low test error.
- This happens even for large and complex neural network architectures.
- Complexity measures like the Rademacher complexity suggest the opposite behaviour (overfitting)

Multiple global minimizers of the empirical risk




- The global minimum is R_n^*

Multiple global minimizers of the empirical risk



- The global minimum is R_n^* , but many parameters can attain such value.


Multiple global minimizers of the empirical risk



- The global minimum is R_n^* , but many parameters can attain such value.
- Each minimizer of the **empirical risk** might have a different **true risk**.

Not all global minimizers are the same

- Consider a simple $2D$ classification task, and train a neural network with fixed step-size SGD.
- The plots below correspond to two different global minimizers:



SGD never lands on the global minimum on the right! Why?

Understanding the implicit bias of optimization algorithms

- SGD seems to be *biased* towards *good* global minimizers (low true risk).
- Some optimization algorithms have an implicit bias towards certain kinds of global minimizers.
- Can we characterize this implicit bias?

Understanding the implicit bias of optimization algorithms

- SGD seems to be *biased* towards *good* global minimizers (low true risk).
- Some optimization algorithms have an implicit bias towards certain kinds of global minimizers.
- Can we characterize this implicit bias?

Definition (Algorithm)

We will refer to a function (deterministic or randomized) $\mathcal{A} : \mathcal{Z} \rightarrow \mathcal{X}$, mapping $Z \mapsto \mathcal{A}_Z$ as an *algorithm* with input $Z \in \mathcal{Z}$ and output $\mathcal{A}_Z \in \mathcal{X}$.

Example: Gradient Descent Algorithm

We denote $\text{GD}_{(T, \alpha, \mathbf{x}^0, \nabla f)} := T$ -steps of GD with stepsize α , starting from \mathbf{x}^0 , using gradient ∇f .

What is implicit regularization?

Definition (Implicit Regularization of a Deterministic Algorithm)

Consider a minimization problem

$$F^* = \min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x})$$

and let \mathcal{A} be a deterministic algorithm with input $Z \in \mathcal{Z}$ and output $\mathcal{A}_Z \in \mathcal{X}$.

We say that \mathcal{A} solves problem (2) and has *implicit regularization* $H : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$ if

$$\mathcal{A}_Z \in \arg \min_{F(\mathbf{x})=F^*} H(\mathbf{x}, Z).$$

Given the input $Z \in \mathcal{Z}$, the algorithm outputs a global minimizer of F that, **additionally**, minimizes $H(\cdot, Z)$.

Implicit bias of gradient descent for linear regression

- Consider for example an underdetermined linear system

$$\mathbf{Ax} = \mathbf{b}, \quad \text{with } \mathbf{A} \in \mathbb{R}^{n \times p}, \quad n < p$$

- If a solution exists (i.e., $\mathbf{b} \in \text{colspan}(\mathbf{A})$), then there is an *infinite number of solutions* to this system.

Finding a solution

To find a valid \mathbf{x} , we could apply one of the optimization algorithms seen in class to the convex problem


$$\arg \min_{\mathbf{x} \in \mathbb{R}^p} \frac{1}{2} \|\mathbf{Ax} - \mathbf{b}\|_2^2$$

Among all the possible solutions, which one will the algorithm converge to ?

Same problem and same initialization vs different algorithms and different solutions


Consider the following simple 2D example :

$$\begin{bmatrix} 1 & 2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = 5$$



Different Solutions

Gradient Descent and AdaGrad converge to *different* points on the line.



Implicit bias of gradient descent for linear regression

- Gradient descent seems to converge to the closest one in terms of ℓ_2 -norm.

Theorem (Implicit bias of Gradient Descent [10])

For the underdetermined, realizable linear system


$$F^* = \min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$$

the gradient descent algorithm $GD_{(T, \alpha, \mathbf{x}^0, \nabla F)}$, for $T = \infty$ and for any $\mathbf{x}^0 \in \mathbb{R}^p$, and valid step-size α , has implicit bias $H(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}^0\|_2$, i.e.,

$$GD_{(T=\infty, \alpha, \mathbf{x}^0, \nabla F)} = \arg \min_{F(\mathbf{x})=F^*} \|\mathbf{x} - \mathbf{x}^0\|_2.$$

- Remark:**
- The theorem also holds for stochastic gradient descent, see [2].

Same problem and same initialization vs different algorithms and different solutions



Proof : For simplicity, take $x_0 = 0$.

- ▶ The gradient of F is $\mathbf{A}^T(\mathbf{A}\mathbf{x} - b)$.
- ▶ This implies that $\forall \mathbf{x}, \nabla f(\mathbf{x}) \in \text{colspan}(\mathbf{A}^T)$.

GD iterates stay in the rowspan


Gradient Descent is therefore constrained to the space

$$\text{colspan}(\mathbf{A}^T) = \text{rowspan}(\mathbf{A})$$

So its limit point at $T = \infty$ is in $\text{rowspan}(\mathbf{A})$.

- ▶ Note that because of the preconditioning, AdaGrad can get out of the $\text{rowspan}(\mathbf{A})$.

Same problem and same initialization vs different algorithms and different solutions



Proof (continued):

- ▶ The minimum norm solution

$$\hat{\mathbf{x}}_{\text{candidate}} = \arg \min_{\mathbf{x}: \mathbf{A}\mathbf{x}=\mathbf{b}} \|\mathbf{x}\|_2^2$$

is also in $\text{rowspan}(\mathbf{A})$.

- ▶ So both $\hat{\mathbf{x}}_{\text{candidate}}$ and the limit point of GD are solutions of $\mathbf{A}\mathbf{x} = \mathbf{b}$ that are in the $\text{rowspan}(\mathbf{A})$
- ▶ Since $\text{null}(\mathbf{A}) \cap \text{rowspan}(\mathbf{A}) = \{0\}$, there can only be one solution in the $\text{rowspan}(\mathbf{A})$, so

$$\mathbf{x}_{\text{GD}}^* = \hat{\mathbf{x}}_{\text{candidate}}$$

Implicit bias for linear models

- We can extend this analysis to linear models:

$$\arg \min_{\mathbf{x} \in \mathbb{R}^p} F(\mathbf{x}) := \sum_{i=1}^n L(\langle \mathbf{x}, \mathbf{a}_i \rangle, b_i).$$

- If the observations are realizable and there are many global minima $\mathbf{Glob} = \{\mathbf{x} : F(\mathbf{x}) = 0\}$, then

Theorem (Implicit Bias of Gradient Descent [10])


If the loss L is convex and has a unique (attained) minimum, then the iterates \mathbf{x}^t of Gradient Descent converge to the global minimum that is *closest to initialization* \mathbf{x}_0 in the ℓ_2 -distance :

$$\mathbf{x}^t \xrightarrow[t \rightarrow \infty]{} \arg \min_{\mathbf{x} \in \mathbf{Glob}} \|\mathbf{x} - \mathbf{x}_0\|_2$$

Proof : (Sketch) The assumption on L implies the problem reduces to a linear system: If \mathbf{x} is a global minimum, we must have $\langle \mathbf{x}, \mathbf{a}_i \rangle = b_i$ for all $i \in \{1, \dots, n\}$. We can recycle the results we have just seen.

Implicit bias for linearly separable datasets

- For linearly separable datasets, we know of an algorithm capable of finding a separating hyperplane.
- It maximizes the *margin* (i.e., distance between the boundary and the nearest training-data point).



Hard-margin Support Vector Machines

The hard margin Support Vector Machine solves the following optimization problem :

$$\arg \min_{\mathbf{x} \in \mathbb{R}^P} \|\mathbf{x}\|_2 \quad \text{subject to } y_i \langle \mathbf{x}, \mathbf{a}_i \rangle \geq 1.$$

It finds a hyperplane that maximizes the margin. It does so *by design*.

Implicit bias for linearly separable datasets

- What happens if we do not explicitly enforce margin maximization?

Theorem (Implicit Bias of Gradient Descent on Separable Data [23, 10])

For the logistic loss (and some other strictly monotonically decreasing losses) and for linearly separable datasets, the direction of the iterates \mathbf{x}^t of Gradient Descent *for any initialization* converges to the hard-margin SVM direction:

$$\frac{\mathbf{x}^t}{\|\mathbf{x}^t\|_2} \xrightarrow{t \rightarrow \infty} \frac{\mathbf{x}_{SVM}^*}{\|\mathbf{x}_{SVM}^*\|_2} \quad \text{where } \mathbf{x}_{SVM}^* = \left\{ \arg \min_{\mathbf{x} \in \mathbb{R}^p} \|\mathbf{x}\|_2 \quad \text{subject to } y_i \langle \mathbf{x}, \mathbf{a}_i \rangle \geq 1 \right\}$$

Remarks:

- Here, without explicit instructions, gradient descent maximizes the margin.
- The rate of this convergence is $O\left(\frac{1}{\log t}\right)$.

Implicit bias for linearly separable datasets

- A similar result can be established for stochastic gradient descent for the logistic loss on separable datasets.

Theorem (Implicit Bias of *Stochastic* Gradient Descent on Separable Data [18])

The direction of the iterates \mathbf{x}^t of Stochastic Gradient Descent *for any initialization* and for a small enough *fixed step-size*, converges almost surely to the hard-margin SVM direction:

$$\left\| \frac{\mathbf{x}^t}{\|\mathbf{x}^t\|_2} - \frac{\mathbf{x}_{SVM}^*}{\|\mathbf{x}_{SVM}^*\|_2} \right\|_2 = O\left(\frac{1}{\log t}\right)$$

- Remarks:**
- This result is particularly interesting as it establishes convergence of fixed step-size SGD.
 - Both SGD and GD have the same implicit bias towards maximizing margins.

Implicit bias for non-convex objectives

- Characterizing implicit bias of stochastic gradient descent for non-convex objectives is an active research area.
- Some papers study deep matrix factorization as a first step towards getting results for neural networks.

Deep Matrix Factorization

Deep matrix factorization consists of parametrizing a matrix \mathbf{M} as a product of N matrices:

$$\mathbf{M} = \mathbf{X}_N \mathbf{X}_{N-1} \dots \mathbf{X}_1$$

which can be understood as parametrizing \mathbf{M} by a depth N “*linear neural network*,” i.e., a neural network with no activations and with weight matrices \mathbf{X} .

Implicit bias for deep matrix completion

- The matrix completion problem consists of filling the missing entries of a partially observed matrix.
- The deep matrix factorization approach consists of solving the following problem with gradient descent:

$$\arg \min_{\mathbf{X}_N, \mathbf{X}_{N-1}, \dots, \mathbf{X}_1} \sum_{(i,j) \in \Omega} ([\mathbf{X}_N \mathbf{X}_{N-1} \dots \mathbf{X}_1]_{i,j} - b_{i,j})^2.$$

- It was conjectured in 2017 [11] that gradient descent was biased towards solutions with small nuclear norm.

Theorem (Implicit Regularization May Not Be Explainable by Norms (2020) [22])

*For deep matrix completion the implicit bias **can not** be expressed as a function of a norm or semi-norm.*

Double descent

- o A failure of conventional wisdom





Figure: The classical U-shaped risk curve vs. double-descent risk curve. source: [7].

- ▶ classical large-sample limit setting: $n \rightarrow \infty$ under fixed p
- ▶ high dimensional setting: n and p comparably large


Double descent curve in practice (I)

Typical examples:

- ▶ linear/nonlinear regression [13]
- ▶ random features, random forest, and shallow neural networks [7]



(a) Random features model



(b) A fully connected neural network

Figure: Experiments on MNIST. Source: [7].

Double descent curve in practice (II)




Figure: Left: Train and test error as a function of model size, for ResNet18s of varying width on CIFAR-10 with 15% label noise. Right: Test error, shown for varying train epochs. source: [19].

Underparametrized regime




Figure: Low generalization but high empirical error




Figure: Sweet spot for the model complexity

Interpolation threshold




Figure: The unique degree 19 polynomial that can fit 20 samples.

Benign overfitting in the over-parametrized regime




Figure: A degree 200 polynomial that can harmlessly fits noisy 20 points.




Figure: Double descent for polynomial fits

Benign Overfitting [6]: good prediction with zero training error for regression loss

- ▶ Statistical wisdom: a predictor should not fit too well.
- ▶ deep networks fit perfectly on noisy data and generalize well on test data.

A simple case: linear regression with Gaussian data

Problem setting

- linear model:
 - ▶ training data $\{\mathbf{a}_i\}_{i=1}^n$ with $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ such that $\mathbf{A} = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n]^\top$
 - ▶ label $b_i = \langle \mathbf{x}^\dagger, \mathbf{a}_i \rangle + w_i$ with the target vector $\mathbf{x}^\dagger \sim \mathcal{N}(\mathbf{0}, \frac{1}{p} \mathbf{I}_p)$, noise $w_i \sim \mathcal{N}(0, \sigma^2)$
- min-norm solution:
 - ▶ $\mathbf{x}^* = \arg \min_{\mathbf{x}} \{\|\mathbf{x}\|_2 : \mathbf{A}\mathbf{x} = \mathbf{b}\} = (\mathbf{A}^\top \mathbf{A})^\dagger \mathbf{A}^\top \mathbf{b}$
- excess risk: for a test point \mathbf{a}
 - ▶ $R(\mathbf{x}^*; \mathbf{x}^\dagger) = \mathbb{E}[(\langle \mathbf{a}, \mathbf{x}^* \rangle - \langle \mathbf{a}, \mathbf{x}^\dagger \rangle)^2 | \mathbf{A}]$

Theorem [13]

Under the above problem setting, assume that $\|\mathbf{x}^\dagger\|_2^2 = r^2$, as $n, p \rightarrow \infty$, and $p/n \rightarrow \gamma$, then we have


$$R(\mathbf{x}^*, \mathbf{x}^\dagger) \rightarrow \begin{cases} \sigma^2 \frac{\gamma}{1-\gamma}, & \text{for } \gamma < 1 \\ r^2 \left(1 - \frac{1}{\gamma}\right) + \sigma^2 \frac{1}{\gamma-1}, & \text{for } \gamma > 1. \end{cases}$$

Remark: The asymptotic risk curves depend on γ and the SNR r^2/σ^2 .

Alternatives to complexity-based generalization bounds

- So far we have seen that complexity based generalization bounds:
 - ▶ characterize **worst-case scenario**
 - ▶ not tight in practice
 - ▶ disregard the effect of the optimization algorithm


Can we understand generalization as a property of an optimization algorithm?



Alternatives to complexity-based generalization bounds

- So far we have seen that complexity based generalization bounds:
 - ▶ characterize **worst-case scenario**
 - ▶ not tight in practice
 - ▶ disregard the effect of the optimization algorithm

Can we understand generalization as a property of an optimization algorithm? YES!



Formal definition of stability (I)

Definition (Uniform Stability [12])

Let $\mathcal{A} : \mathcal{Z} \rightarrow \mathcal{H}$ be a randomized algorithm with input a finite sample S , and output a function $\mathcal{A}_S \in \mathcal{H}$.

The algorithm \mathcal{A} has uniform stability $(\beta_n)_{n \geq 1}$ with respect to the loss function L if for all subsets $S, S' \subseteq \mathcal{A} \times \mathcal{B}$ such that $|S| = |S'| = n$ and S and S' differ in at most one sample:

$$\sup_{(\mathbf{a}, b) \in \mathcal{A} \times \mathcal{B}} \mathbb{E} |L(\mathcal{A}_S(\mathbf{a}), b) - L(\mathcal{A}_{S'}(\mathbf{a}), b)| \leq \beta_n$$

The expectation is taken with respect to the randomness in the algorithm \mathcal{A} .

Misnomer: Lower stability (small values of β_n) means the difference in the output of the algorithm is smaller.

Formal definition of stability (II)

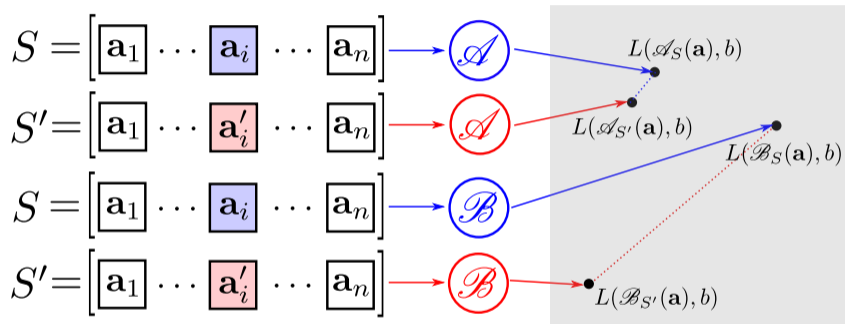


Figure: Algorithm \mathcal{B} is less stable than algorithm \mathcal{A} .

Generalization bounds based on uniform stability — definitions

Definition (Empirical Risk on a set)

Let $S := [(\mathbf{a}_1, b_1), \dots, (\mathbf{a}_n, b_n)]$ be an i.i.d. sample drawn from a distribution on $\mathcal{A} \times \mathcal{B}$. Let $L : \mathcal{B} \times \mathcal{B} \rightarrow \mathbb{R}$ be a loss function and \mathcal{H} be a class of functions $h : \mathcal{A} \rightarrow \mathcal{B}$. The empirical risk of $h \in \mathcal{H}$ on the set S is defined as:

$$R_S(h) := \frac{1}{n} \sum_{i=1}^n L(h(\mathbf{a}_i), b_i)$$

(Almost) same definition as before. Makes explicit the dependence on the set S .

Definition (Expected Generalization Error)

Let $\mathcal{A} : \mathcal{Z} \rightarrow \mathcal{H}$ be a randomized algorithm that takes as input a finite sample S of arbitrary size, and outputs a function $\mathcal{A}_S \in \mathcal{H}$. Suppose that $S = [(\mathbf{a}_1, b_1), \dots, (\mathbf{a}_n, b_n)]$ is an i.i.d. sample from probability distribution on $\mathcal{A} \times \mathcal{B}$. The expected generalization error on a sample of size n is the value

$$\mathbb{E}[R_S(\mathcal{A}_S) - R(\mathcal{A}_S)]$$

the expectation is taken with respect to the draw of the sample S and the randomness of \mathcal{A} .

Generalization bounds based on uniform stability — Fundamental theorem (I)

Theorem (Hardt et al. 2016 [12])

Let A be uniformly stable with stability $(\beta_n)_{n \geq 1}$, then for a random i.i.d. sample S of size n , the expected generalization error is bounded as follows

$$\mathbb{E}[|R_S(\mathcal{A}_S) - R(\mathcal{A}_S)|] \leq \beta_n$$

Proof.

Let $S = [(\mathbf{a}_1, b_1), \dots, (\mathbf{a}_n, b_n)]$ and $S' = [(\mathbf{a}'_1, b'_1), \dots, (\mathbf{a}'_n, b'_n)]$ be two i.i.d. samples of size n . Denote

$$S^{(i)} := [(\mathbf{a}_1, b_1), \dots, (\mathbf{a}_{i-1}, b_{i-1}), (\mathbf{a}'_i, b'_i), (\mathbf{a}_{i+1}, b_{i+1}), \dots, (\mathbf{a}_n, b_n)]$$

the sample that results from replacing (\mathbf{a}_i, b_i) by (\mathbf{a}'_i, b'_i) in S .

$$\begin{aligned} \mathbb{E}[R_S(\mathcal{A}_S)] &= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_S(\mathbf{a}_i), b_i) \right] = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_{S^{(i)}}(\mathbf{a}'_i), b'_i) \right] \\ &= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_{S^{(i)}}(\mathbf{a}'_i), b'_i) - \frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_S(\mathbf{a}'_i), b'_i) \right] + \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_S(\mathbf{a}'_i), b'_i) \right] \end{aligned}$$

Generalization bounds based on uniform stability — Fundamental theorem (II)

Proof. (continued).

We have

$$\mathbb{E}[R_S(\mathcal{A}_S)] = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_{S^{(i)}}(\mathbf{a}'_i), b'_i) - \frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_S(\mathbf{a}'_i), b'_i) \right] + \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_S(\mathbf{a}'_i), b'_i) \right]$$

Note that S and $S^{(i)}$ only differ in one sample: uniform stability allows bounding the first term as:

$$= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_{S^{(i)}}(\mathbf{a}'_i), b'_i) - \frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_S(\mathbf{a}'_i), b'_i) \right] = \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[L(\mathcal{A}_{S^{(i)}}(\mathbf{a}'_i), b'_i) - L(\mathcal{A}_S(\mathbf{a}'_i), b'_i) \right] \leq \beta_n$$

Finally note that because the samples (\mathbf{a}_i, b_i) are independent of S we have:

$$\mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n L(\mathcal{A}_S(\mathbf{a}'_i), b'_i) \right] = R(\mathcal{A}_S)$$

analogously we can show $\mathbb{E}[R(\mathcal{A}_S) - R_S(\mathcal{A}_S)] \leq \beta_n$. □

The stability of SGD

- Let $h_{\mathbf{x}} \in \mathcal{H}_{\mathcal{X}}$ be an element of a parametric function class. Consider the ERM optimization objective:

$$f(\mathbf{x}) := \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x}), \quad f_i(\mathbf{x}) := L(h_{\mathbf{x}}(\mathbf{a}_i), b_i).$$

- The SGD iterates for $t = 0, \dots, T$ are $\mathbf{x}_{t+1} = \mathbf{x}_t - \alpha_t \nabla_{\mathbf{x}} f_i(\mathbf{x}_t)$, for $i \sim \text{Unif}[n]$.

Algorithm	Assumptions on f_i	Stability
SGD	convex, L -smooth, β -Lipschitz, $\alpha_t \leq 2/L$	$\frac{\beta^2}{n} \sum_{t=0}^T \alpha_t$
SGD	μ -str convex, L -smooth, β -Lipschitz, $\alpha_t \leq 2/L$	$\frac{\beta^2}{n\mu}$
SGD	μ -str convex, L -smooth, β -Lipschitz, $\alpha_t = \frac{1}{\mu t}$	$\frac{\beta^2 + L\rho}{n\mu}$
SGD avg. iterate	convex, L -smooth, β -Lipschitz	$\frac{\beta^2 T}{n\mu}$
SGD	non-convex, L -smooth, β -Lipschitz, $\alpha_t = 1/t$	$\frac{1 + 1/\beta}{n} \beta^{\frac{2}{L+1}} T^{\frac{L}{L+1}}$

Table: Summary of stability upper bounds for different assumptions on the objective function [12]

Effect of the number of iterations on the stability of SGD and the generalization error

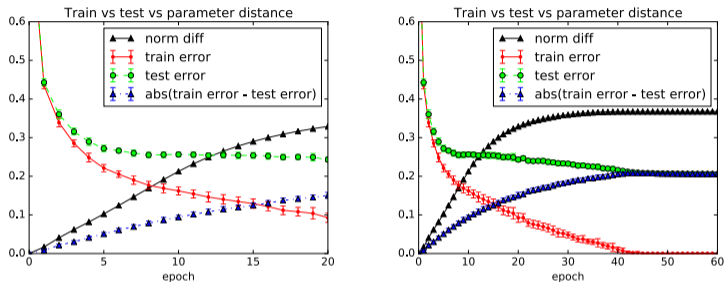


Figure: Normalized parameter distance between two networks trained on two datasets S, S' differing only in one sample, training error, test error and generalization error (0-1 loss) on CIFAR10 [12].

- Parameter distance is a stronger notion than stability.
- More iterations \Rightarrow Parameter distance increases (we expect stability to increase).
- Generalization error follows the same behavior as the parameter distance (proxy for stability).

Wrap up!

- The visualizations can be deceiving to understand the high-dimensional behavior
- Are we really in the interpolation regime in machine learning?

Theorem (Probability of interpolation [5])

Given a p -dimensional dataset $\mathcal{A}_n = \{\mathbf{a}_1, \dots, \mathbf{a}_n\}$ with i.i.d. samples, where $\mathbf{a}_i \sim \mathcal{N}(0, I)$ for all $i = 1, \dots, n$, the probability that a new sample $\mathbf{a} \sim \mathcal{N}(0, I)$ is in the interpolation regime (i.e., within the convex hull of \mathcal{A}_n) has the following limiting behavior

$$\lim_{p \rightarrow \infty} p(\mathbf{a} \in \text{convex hull}(\mathcal{A}_n)) = \begin{cases} 1 & \text{if } n > 2^{p/2}/p; \\ 0 & \text{if } n < 2^{p/2}/p. \end{cases}$$

- We are most likely in the extrapolation regime [4]

Wrap up!

- The visualizations can be deceiving to understand the high-dimensional behavior
- Are we really in the interpolation regime in machine learning?

Theorem (Probability of interpolation [5])

Given a p -dimensional dataset $\mathcal{A}_n = \{\mathbf{a}_1, \dots, \mathbf{a}_n\}$ with i.i.d. samples, where $\mathbf{a}_i \sim \mathcal{N}(0, I)$ for all $i = 1, \dots, n$, the probability that a new sample $\mathbf{a} \sim \mathcal{N}(0, I)$ is in the interpolation regime (i.e., within the convex hull of \mathcal{A}_n) has the following limiting behavior

$$\lim_{p \rightarrow \infty} p(\mathbf{a} \in \text{convex hull}(\mathcal{A}_n)) = \begin{cases} 1 & \text{if } n > 2^{p/2}/p; \\ 0 & \text{if } n < 2^{p/2}/p. \end{cases}$$

- We are most likely in the extrapolation regime [4]
- Recitation 4 on Friday!

*From neural networks to random features model [14, 21]

1-hidden-layer neural network with m neurons (fully-connected architecture):

Let $\mathbf{X}_1 \in \mathbb{R}^{m \times p}$, $\mathbf{a} \in \mathbb{R}^p$, $\mathbf{X}_2 \in \mathbb{R}^m$, and $\mu_2 \in \mathbb{R}$

$$h_{\mathbf{x}}(\mathbf{a}) := \begin{bmatrix} \mathbf{X}_2 \\ \sigma \left(\underbrace{\begin{bmatrix} \mathbf{X}_1 \\ \mathbf{a} \end{bmatrix} + \begin{bmatrix} \mu_1 \end{bmatrix}}_{\text{hidden layer = fixed random features}} \right) + \begin{bmatrix} \mu_2 \end{bmatrix} \end{bmatrix}, \quad \mathbf{x} := [\mathbf{X}_1, \mathbf{X}_2, \mu_1, \mu_2]$$

The diagram illustrates the computation of the hidden layer output. The input vector \mathbf{a} (green) is combined with a bias μ_1 (blue) and multiplied by the weight matrix \mathbf{X}_1 (weight). The result is passed through an activation function σ (activation). The output of the hidden layer is then combined with a bias μ_2 (bias) and added to the output of the second layer \mathbf{X}_2 .

- ▶ \mathbf{X}_1 : Gaussian initialization and then fixed
- ▶ \mathbf{X}_2 : to be learned
- ▶ over-parameterized model: #neurons $m >$ #training data n

*Double descent: random features model (I)

o high dimensions: #training data n , #neurons m , feature dimension p are comparably large




Figure: Test MSE, Bias, and Variance of RF regression as a function of the ratio m/n on MNIST data set (digit 3 vs. 7) for $p = 784$ and $n = 600$ across the Gaussian kernel. Source: [16].

- ▶ random features regression solved by SGD: interplay between excess risk and optimization
- ▶ bias variance decomposition for understanding multiple randomness sources
- ▶ monotonic decreasing bias and unimodal variance \Rightarrow double descent

*Double descent: random features model (II)

Algorithm	data assumption	solution type	Result on risk curve
[13]	Gaussian	closed-form	variance ↗ ↘
[17]	i.i.d on sphere	closed-form	variance, bias ↗ ↘
[8]	Gaussian	closed-form	refined decomposition on variance
[1]	Gaussian	closed-form	fully decomposition on variance
[15]	general	closed-form	↗ ↘
[3]	Gaussian	GD	variance ↗ ↘
[16]	sub-exponential	SGD	variance ↗ ↘, bias ↘

Table: Comparison of representative random features on double descent.

- o multiple randomness sources: data sampling, label noise, initialization
- o phase transition due to non-monotonic variance

References I

- [1] Ben Adlam and Jeffrey Pennington.
Understanding double descent requires a fine-grained bias-variance decomposition.
In Advances in neural information processing systems, 2020.
- [2] Navid Azizan and Babak Hassibi.
Stochastic gradient/mirror descent: Minimax optimality and implicit regularization.
- [3] Jimmy Ba, Murat A. Erdogdu, Taiji Suzuki, Denny Wu, and Tianzong Zhang.
Generalization of two-layer neural networks: an asymptotic viewpoint.
In International Conference on Learning Representations, pages 1–8, 2020.
- [4] Randall Balestriero, Jerome Pesenti, and Yann LeCun.
Learning in high dimension always amounts to extrapolation.
arXiv preprint arXiv:2110.09485, 2021.
- [5] Imre Bárány and Zoltán Füredi.
On the shape of the convex hull of random points.
Probability theory and related fields, 77(2):231–240, 1988.
- [6] Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler.
Benign overfitting in linear regression.
Proceedings of the National Academy of Sciences, 117(48):30063–30070, 2020.

References II

- [7] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal.
Reconciling modern machine-learning practice and the classical bias–variance trade-off.
Proceedings of the National Academy of Sciences, 116(32):15849–15854, 2019.
- [8] Stéphane d’Ascoli, Maria Refinetti, Giulio Biroli, and Florent Krzakala.
Double trouble in double descent: Bias and variance (s) in the lazy regime.
pages 2280–2290, 2020.
- [9] Stuart Geman, Elie Bienenstock, and René Doursat.
Neural networks and the bias/variance dilemma.
Neural computation, 4(1):1–58, 1992.
- [10] Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro.
Characterizing implicit bias in terms of optimization geometry.
- [11] Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro.
Implicit regularization in matrix factorization.
In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors,
Advances in Neural Information Processing Systems 30, pages 6151–6159. Curran Associates, Inc.

References III

- [12] Moritz Hardt, Ben Recht, and Yoram Singer.
Train faster, generalize better: Stability of stochastic gradient descent.
In International Conference on Machine Learning, pages 1225–1234. PMLR, 2016.
- [13] Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani.
Surprises in high-dimensional ridgeless least squares interpolation.
arXiv preprint arXiv:1903.08560, 2019.
- [14] Guang-Bin Huang, Qin-Yu Zhu, and Chee-Kheong Siew.
Extreme learning machine: theory and applications.
Neurocomputing, 70(1-3):489–501, 2006.
- [15] Zhenyu Liao, Romain Couillet, and Michael Mahoney.
A random matrix analysis of random fourier features: beyond the gaussian kernel, a precise phase transition, and the corresponding double descent.
In Neural Information Processing Systems, 2020.
- [16] Fanghui Liu, Johan A.K. Suykens, and Volkan Cevher.
On the double descent of random features models trained with sgd.
arXiv preprint arXiv:2110.06910, 2021.

References IV

- [17] Song Mei and Andrea Montanari.
The generalization error of random features regression: Precise asymptotics and double descent curve.
arXiv preprint arXiv:1908.05355, 2019.
- [18] Mor Shpigel Nacson, Nathan Srebro, and Daniel Soudry.
Stochastic gradient descent on separable data: Exact convergence with a fixed learning rate.
- [19] Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever.
Deep double descent: Where bigger models and more data hurt.
- [20] Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro.
The role of over-parametrization in generalization of neural networks.
In International Conference on Learning Representations, 2019.
- [21] Ali Rahimi and Benjamin Recht.
Random features for large-scale kernel machines.
In Advances in Neural Information Processing Systems, pages 1177–1184, 2007.
- [22] Noam Razin and Nadav Cohen.
Implicit regularization in deep learning may not be explainable by norms.
- [23] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro.
The implicit bias of gradient descent on separable data.