More Deep Learning: Approximation, Generalization

CPSC 532S: Modern Statistical Learning Theory
16 March 2022
cs.ubc.ca/~dsuth/532S/22/

Admin

- A3 is up; due **next Friday the 25th**
- There will be an A4 due at the end of the term
  - Going to try to post that soon (before A3 is due) so you can start early

- This week only, Thursday office hours were moved to this morning

- Project proposals due tonight
  - An informal paragraph on Piazza: just tell me what you want to do
  - Say who your partners are in the post, I'll add them to be able to see it
  - Again, the scope of these is meant to be *small*
    - A lit survey doesn't require fully understanding the proofs or anything
    - An "extension" could be just reading the proofs and talking about when the assumptions hold, etc

Last time

- Deep learning, particularly fully-connected feedforward networks

Last time

- Deep learning, particularly fully-connected feedforward networks

- Universal approximation results:
  - There *exist networks* (of depth two, but potentially extremely wide) that approximate any continuous function in sup-norm
    - as long as the activation function is non-polynomial

Universal approximation via circuit complexity

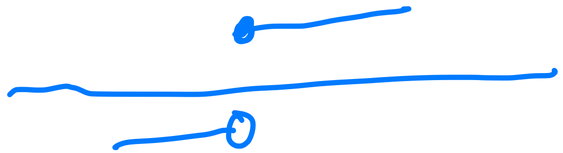

SSBD chapter 20:

- 2 layer nets with sign activations can represent all functions $\{\pm 1\}^d \to \{\pm 1\}$

Universal approximation via circuit complexity

SSBD chapter 20:

- 2 layer nets with sign activations can represent all functions $\{\pm 1\}^d \to \{\pm 1\}$
  - (remember that computers always represent things as $\{0,1\}^d$…)

Universal approximation via circuit complexity

SSBD chapter 20:

- 2 layer nets with sign activations can represent all functions $\{\pm 1\}^d \to \{\pm 1\}$
  - (remember that computers always represent things as $\{0,1\}^d$…)
- …but, it takes exponential width to do that

Universal approximation via circuit complexity

SSBD chapter 20:

- 2 layer nets with sign activations can represent all functions $\{\pm 1\}^d \to \{\pm 1\}$
  - (remember that computers always represent things as $\{0,1\}^d$…)
- …but, it takes exponential width to do that
- …but, there's a network of size $\mathcal{O}(T^2)$ that can implement all boolean functions that can be computed in maximum runtime T

Universal approximation via circuit complexity

SSBD chapter 20:

- 2 layer nets with sign activations can represent all functions $\{\pm 1\}^d \to \{\pm 1\}$
  - (remember that computers always represent things as $\{0,1\}^d$…)
- …but, it takes exponential width to do that
- …but, there's a network of size $\mathcal{O}(T^2)$ that can implement all boolean functions that can be computed in maximum runtime T

Universal approximation via circuit complexity

SSBD chapter 20:

- 2 layer nets with sign activations can represent all functions $\{\pm 1\}^d \to \{\pm 1\}$
  - (remember that computers always represent things as $\{0,1\}^d$…)
- …but, it takes exponential width to do that
- …but, there's a network of size $\mathcal{O}(T^2)$ that can implement all boolean functions that can be computed in maximum runtime T

Circuit Complexity and Neural Networks, Ian Parberry (1994) - UBC access

Proceedings of Machine Learning Research vol 125:1–22, 2020 33rd Annual Conference on Learning Theory

Universal Approximation with Deep Narrow Networks

Patrick Kidger KIDGER@MATHS.OX.AC.UK
Terry Lyons TLYONS@MATHS.OX.AC.UK
Mathematical Institute, University of Oxford

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

The classical Universal Approximation Theorem holds for neural networks of arbitrary width and bounded depth. Here we consider the natural 'dual' scenario for networks of bounded width and arbitrary depth. Precisely, let n be the number of inputs neurons, m be the number of output neurons, and let ρ be any nonaffine continuous function, with a continuous nonzero derivative at some point. Then we show that the class of neural networks of arbitrary depth, width $n + m + 2$, and activation function ρ, is dense in $C(K; \mathbb{R}^m)$ for $K \subseteq \mathbb{R}^n$ with K compact. This covers every activation function possible to use in practice, and also includes polynomial activation functions, which is unlike the classical version of the theorem, and provides a qualitative difference between deep narrow networks and shallow wide networks. We then consider several extensions of this result. In particular we consider nowhere differentiable activation functions, density in noncompact domains with respect to the L^p-norm, and how the width may be reduced to just $n + m + 1$ for 'most' activation functions.

Why deep instead of wide?

- Deep networks much better at learning *compositional structure*

Why deep instead of wide?

- Deep networks much better at learning *compositional structure*
  - “We claim that most functions that can be represented compactly by deep architectures cannot be represented by a compact shallow architecture.”

– Y. Bengio & LeCun (2007)

Why deep instead of wide?

- Deep networks much better at learning *compositional structure*
  - “We claim that most functions that can be represented compactly by deep architectures cannot be represented by a compact shallow architecture.”

    – Y. Bengio & LeCun (2007)
  - Lots of empirical evidence, but theoretical support pretty limited until recently

Why deep instead of wide?

- Deep networks much better at learning *compositional structure*
  - “We claim that most functions that can be represented compactly by deep architectures cannot be represented by a compact shallow architecture.”
    – Y. Bengio & LeCun (2007)
  - Lots of empirical evidence, but theoretical support pretty limited until recently

  - Telgarsky notes section 5 give a particular such function:
    shallow net needs huge width to approximate,
    but narrow not-super-deep net can approximate it efficiently

Why deep instead of wide?

- Deep networks much better at learning *compositional structure*
  - “We claim that most functions that can be represented compactly by deep architectures cannot be represented by a compact shallow architecture.”
    – Y. Bengio & LeCun (2007)
  - Lots of empirical evidence, but theoretical support pretty limited until recently

  - Telgarsky notes section 5 give a particular such function:
    shallow net needs huge width to approximate,
    but narrow not-super-deep net can approximate it efficiently
  - Also proved for a certain class of functions by Mhaskar, Liao, Poggio (2016)

Why deep instead of wide?

- Deep networks much better at learning *compositional structure*
  - "We claim that most functions that can be represented compactly by deep architectures cannot be represented by a compact shallow architecture."
    – Y. Bengio & LeCun (2007)
  - Lots of empirical evidence, but theoretical support pretty limited until recently

  - Telgarsky notes section 5 give a particular such function:
    shallow net needs huge width to approximate,
    but narrow not-super-deep net can approximate it efficiently
  - Also proved for a certain class of functions by Mhaskar, Liao, Poggio (2016)
  - Lu et al. (2017): approximating wide nets with deep nets easier(ish) than vice versa

Why deep instead of wide?

- Deep networks much better at learning *compositional structure*
  - "We claim that most functions that can be represented compactly by deep architectures cannot be represented by a compact shallow architecture."
    – Y. Bengio & LeCun (2007)
  - Lots of empirical evidence, but theoretical support pretty limited until recently

  - Telgarsky notes section 5 give a particular such function:
    shallow net needs huge width to approximate,
    but narrow not-super-deep net can approximate it efficiently
  - Also proved for a certain class of functions by Mhaskar, Liao, Poggio (2016)
  - Lu et al. (2017): approximating wide nets with deep nets easier(ish) than vice versa
  - Liang and Srikant (2017): can approximate piecewise-constant funcs with exponentially smaller deep nets than shallow

- We have some universal approximation results
- A lot of people use this to say “neural networks can do anything! 💪”

- We have some universal approximation results
- A lot of people use this to say “neural networks can do anything! 💪”

but…

- We have some universal approximation results
- A lot of people use this to say “neural networks can do anything! 💪”

but…

- These kind of approximation results don’t tell us:
  - What practically-sized networks can do

- We have some universal approximation results
- A lot of people use this to say “neural networks can do anything! 💪”

but…

- These kind of approximation results don’t tell us:
  - What practically-sized networks can do
    - Gaussian kernels can also do anything (💪)…with ridiculously large norm

- We have some universal approximation results
- A lot of people use this to say “neural networks can do anything! 💪”

but…

- These kind of approximation results don’t tell us:
  - What practically-sized networks can do
    - Gaussian kernels can also do anything (💪)…with ridiculously large norm
    - Neural nets can do anything…if they’re ridiculously large (or large norm)

- We have some universal approximation results
- A lot of people use this to say “neural networks can do anything! 💪”

but…

- These kind of approximation results don’t tell us:
  - What practically-sized networks can do
    - Gaussian kernels can also do anything (💪)…with ridiculously large norm
    - Neural nets can do anything…if they’re ridiculously large (or large norm)

  - Even if our class approximates, do we generalize? (Does ERM, RLM, … work?)

- We have some universal approximation results
- A lot of people use this to say “neural networks can do anything! 💪”

but…

- These kind of approximation results don’t tell us:
  - What practically-sized networks can do
    - Gaussian kernels can also do anything (💪)…with ridiculously large norm
    - Neural nets can do anything…if they’re ridiculously large (or large norm)
  - Even if our class approximates, do we generalize? (Does ERM, RLM, … work?)
  - Does (S)GD find an approximate ERM / RLM / something that generalizes?

- We have some universal approximation results
- A lot of people use this to say “neural networks can do anything! 💪”

but…

- These kind of approximation results don’t tell us:
  - What practically-sized networks can do
    - Gaussian kernels can also do anything (💪)…with ridiculously large norm
    - Neural nets can do anything…if they’re ridiculously large (or large norm)
  - Even if our class approximates, do we generalize? (Does ERM, RLM, … work?)
  - Does (S)GD find an approximate ERM / RLM / something that generalizes?
    - We (pretty much) know it doesn’t *always* find an (approximate) ERM: ERM with deep nets (even for square loss) is NP-hard

so, if you can prove that it *does*, let me know =)

Generalization: VC dimension

- For ReLU (or general piecewise-linear) nets with P params, VCdim = $\mathcal{O}(PL \log P)$
  - and $\Omega\left(PL \log \frac{P}{L}\right)$, so nearly tight – Bartlett/Harvey/Liaw/Mehrabian (2019)

Generalization: VC dimension

- For ReLU (or general piecewise-linear) nets with P params, VCdim = $\mathcal{O}(PL \log P)$
  - and $\Omega\left(PL \log \frac{P}{L}\right)$, so nearly tight – Bartlett/Harvey/Liaw/Mehrabian (2019)
  - $P = \prod_{\ell=1}^{L} d_{\ell-1} d_\ell$ for fully-connected networks

Generalization: VC dimension

- For ReLU (or general piecewise-linear) nets with P params, VCdim = $\mathcal{O}(PL \log P)$
  - and $\Omega\left(PL \log \frac{P}{L}\right)$, so nearly tight – Bartlett/Harvey/Liaw/Mehrabian (2019)
  - $P = \prod_{\ell=1}^{L} d_{\ell-1} d_\ell$ for fully-connected networks
- For piecewise-constant, e.g. threshold functions, VCdim = $\Theta(P \log P)$

Generalization: VC dimension

- For ReLU (or general piecewise-linear) nets with P params, VCdim = $\mathcal{O}(PL \log P)$
  - and $\Omega\left(PL \log \frac{P}{L}\right)$, so nearly tight – Bartlett/Harvey/Liaw/Mehrabian (2019)
  - $P = \prod_{\ell=1}^{L} d_{\ell-1} d_\ell$ for fully-connected networks
- For piecewise-constant, e.g. threshold functions, VCdim = $\Theta(P \log P)$
- For piecewise-polynomial, $\mathcal{O}(PL^2 + PL \log P)$, $\mathcal{O}(PU)$ with U units

Generalization: VC dimension

- For ReLU (or general piecewise-linear) nets with P params, VCdim = $\mathcal{O}(PL \log P)$
  - and $\Omega\left(PL \log \frac{P}{L}\right)$, so nearly tight – Bartlett/Harvey/Liaw/Mehrabian (2019)
  - $P = \prod_{\ell=1}^{L} d_{\ell-1} d_\ell$ for fully-connected networks
- For piecewise-constant, e.g. threshold functions, VCdim = $\Theta(P \log P)$
- For piecewise-polynomial, $\mathcal{O}(PL^2 + PL \log P)$, $\mathcal{O}(PU)$ with U units
- For sigmoids/similar, $\mathcal{O}(P^2U^2)$ and $\Omega(P^2)$

Generalization: VC dimension

- For ReLU (or general piecewise-linear) nets with P params, VCdim = $\mathcal{O}(PL \log P)$
  - and $\Omega\left(PL \log \frac{P}{L}\right)$, so nearly tight – Bartlett/Harvey/Liaw/Mehrabian (2019)
  - $P = \prod_{\ell=1}^{L} d_{\ell-1} d_\ell$ for fully-connected networks
- For piecewise-constant, e.g. threshold functions, VCdim = $\Theta(P \log P)$
- For piecewise-polynomial, $\mathcal{O}(PL^2 + PL \log P)$, $\mathcal{O}(PU)$ with U units
- For sigmoids/similar, $\mathcal{O}(P^2U^2)$ and $\Omega(P^2)$
  - Theorem 8.13/8.14 of Anthony & Bartlett (1999) textbook - UBC access

Problems with parameter counting

- We use networks with a **lot** of parameters
  - ResNet-50 has ~25 million parameters and depth 50: VCdim > 1 billion

Problems with parameter counting

- We use networks with a **lot** of parameters
  - ResNet-50 has ~25 million parameters and depth 50: VCdim > 1 billion

- We can train our networks to get zero error even for **random labels**

Problems with parameter counting

- We use networks with a **lot** of parameters
  - ResNet-50 has ~25 million parameters and depth 50: VCdim > 1 billion

- We can train our networks to get zero error even for **random labels**
  - Even AlexNet can shatter CIFAR-10, *almost* shatter ImageNet
  - Neyshabur et al. (2015), Zhang et al. (2017)

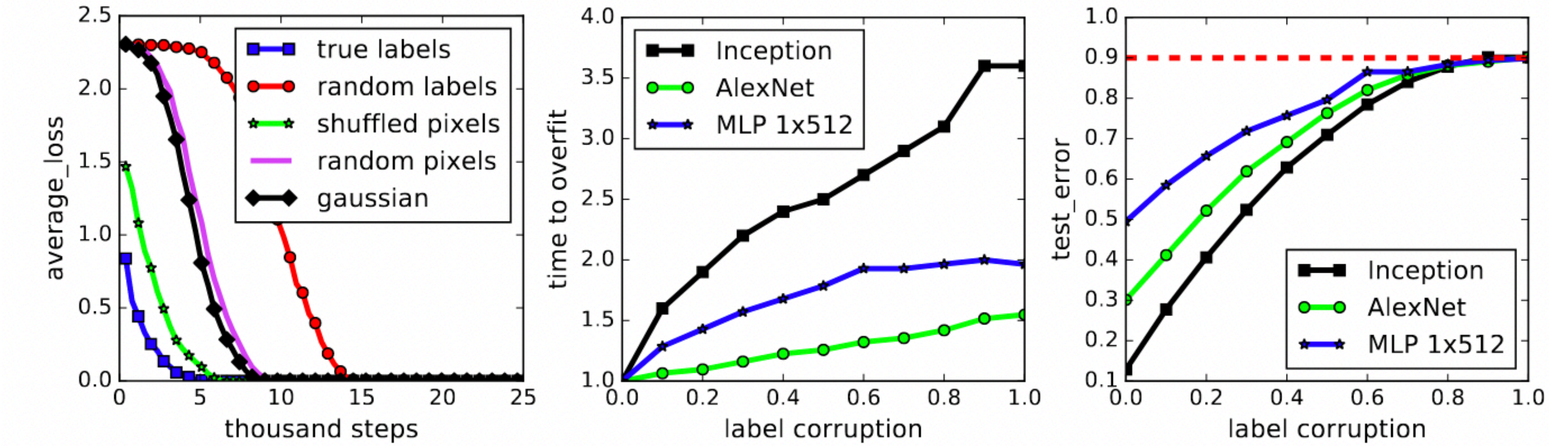

(a) learning curves (b) convergence slowdown (c) generalization error growth

Figure 1: Fitting random labels and random pixels on CIFAR10. (a) shows the training loss of various experiment settings decaying with the training steps. (b) shows the relative convergence time with different label corruption ratio. (c) shows the test error (also the generalization error since training error is 0) under different label corruptions.

Problems with parameter counting

- We use networks with a **lot** of parameters
  - ResNet-50 has ~25 million parameters and depth 50: VCdim > 1 billion

- We can train our networks to get zero error even for **random labels**
  - Even AlexNet can shatter CIFAR-10, *almost* shatter ImageNet
  - Neyshabur et al. (2015), Zhang et al. (2017)
  - But these architectures do generalize well – VC of arch. can't explain that

Problems with parameter counting

- We use networks with a **lot** of parameters
  - ResNet-50 has ~25 million parameters and depth 50: VCdim > 1 billion

- We can train our networks to get zero error even for **random labels**
  - Even AlexNet can shatter CIFAR-10, *almost* shatter ImageNet
  - Neyshabur et al. (2015), Zhang et al. (2017)
  - But these architectures do generalize well – VC of arch. can't explain that
  - Uniform stability can't either, since it's data-independent;
    on-average replace-one stability always can, but hard

Problems with parameter counting

- We use networks with a **lot** of parameters
  - ResNet-50 has ~25 million parameters and depth 50: VCdim > 1 billion

- We can train our networks to get zero error even for **random labels**
  - Even AlexNet can shatter CIFAR-10, *almost* shatter ImageNet
  - Neyshabur et al. (2015), Zhang et al. (2017)
  - But these architectures do generalize well – VC of arch. can't explain that
  - Uniform stability can't either, since it's data-independent;
    on-average replace-one stability always can, but hard

- Making hidden layers wider can often improve generalization,
  but worsens parameter counting-based bounds

Problems with parameter counting

- We use networks with a **lot** of parameters
  - ResNet-50 has ~25 million parameters and depth 50: VCdim > 1 billion

- We can train our networks to get zero error even for **random labels**
  - Even AlexNet can shatter CIFAR-10, *almost* shatter ImageNet
  - Neyshabur et al. (2015), Zhang et al. (2017)
  - But these architectures do generalize well – VC of arch. can't explain that
  - Uniform stability can't either, since it's data-independent;
    on-average replace-one stability always can, but hard

- Making hidden layers wider can often improve generalization,
  but worsens parameter counting-based bounds
- Remember that $\mathscr{H}_k$ has infinite VCdim for universal kernels,
  but we can still learn with *small-norm* predictors

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathcal{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathcal{F}_L) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \left\| \left(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \right) \right\|_c$$

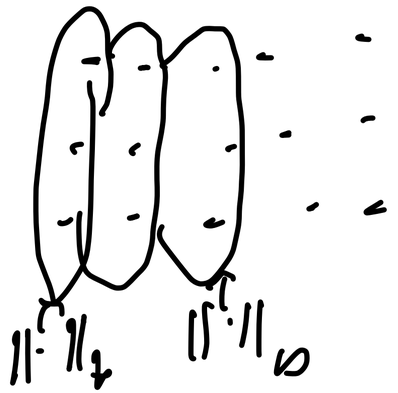

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \left\| \left(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \right) \right\|_c$$

Base case, $L = 0$:

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \left\| \left(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \right) \right\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\})$$

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \Big\|\big(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b\big)\Big\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\}) \le \frac{1}{n}\Big(\max_j \|(x_{1,j}, \ldots, x_{n,j})\|_2\Big)\sqrt{2\log d}$$

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \Big\|\big(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b\big)\Big\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\}) \le \frac{1}{n}\Big(\max_j \|(x_{1,j}, \ldots, x_{n,j})\|_2\Big)\sqrt{2\log d}$$
$$= \frac{1}{n}\|X\|_{2,\infty}\sqrt{2\log d}$$

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \Big\| \big(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \big) \Big\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\}) \le \frac{1}{n}\Big(\max_j \|(x_{1,j}, \ldots, x_{n,j})\|_2 \Big)\sqrt{2\log d}$$
$$= \frac{1}{n}\|X\|_{2,\infty}\sqrt{2\log d} = \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^0\sqrt{2\log d}$$

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathcal{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathcal{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \left\| \left(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \right) \right\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\}) \le \frac{1}{n}\Big(\max_j \|(x_{1,j}, \ldots, x_{n,j})\|_2 \Big)\sqrt{2\log d}$$
$$= \frac{1}{n}\|X\|_{2,\infty}\sqrt{2\log d} = \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^0\sqrt{2\log d}$$

Inductive step:

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathcal{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathcal{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \left\| \left(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \right) \right\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\}) \le \frac{1}{n}\left(\max_j \|(x_{1,j}, \ldots, x_{n,j})\|_2 \right)\sqrt{2\log d}$$
$$= \frac{1}{n}\|X\|_{2,\infty}\sqrt{2\log d} = \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^0\sqrt{2\log d}$$

Inductive step:

$$\hat{\mathfrak{R}}_S(\mathcal{F}_{\ell+1}) = \hat{\mathfrak{R}}_S\left(\left\{ x \mapsto \sigma_{\ell+1}\left(\|W_{\ell+1}^\top\|_{1,\infty}\, g(x) \right) : g \in \mathrm{conv}\left(-\mathcal{F}_\ell \cup \mathcal{F}_\ell \right) \right\} \right)$$

$f \in \mathcal{F}_{\ell+1} \Rightarrow \sigma_\ell\left(\sum_{i=1}^{d_{\ell-1}} w_i^{(\ell)} g_i(x) \right)$, $g \in \mathcal{F}_\ell$

$\sum \alpha_i f_i$ with $\sum \alpha_i = 1$, $\alpha_i \ge 0$

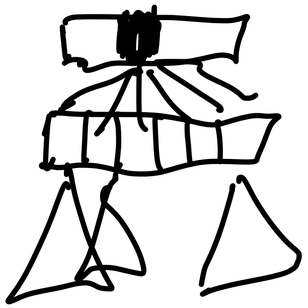

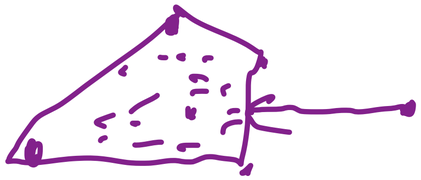

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathcal{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathcal{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \Big\| \big(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \big) \Big\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\}) \le \frac{1}{n}\Big(\max_j \|(x_{1,j}, \ldots, x_{n,j})\|_2 \Big)\sqrt{2\log d}$$
$$= \frac{1}{n}\|X\|_{2,\infty}\sqrt{2\log d} = \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^0\sqrt{2\log d}$$

Inductive step:

$$\hat{\mathfrak{R}}_S(\mathcal{F}_{\ell+1}) = \hat{\mathfrak{R}}_S\left(\left\{ x \mapsto \sigma_{\ell+1}\left(\|W_{\ell+1}^\top\|_{1,\infty}\, g(x) \right) : g \in \mathrm{conv}\big(-\mathcal{F}_\ell \cup \mathcal{F}_\ell \big) \right\} \right)$$
$$\le \rho B\, \hat{\mathfrak{R}}_S\Big(\mathrm{conv}(-\mathcal{F}_\ell \cup \mathcal{F}_\ell) \Big)$$

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \Big\|\big(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b\big)\Big\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\}) \le \frac{1}{n}\Big(\max_j \|(x_{1,j}, \ldots, x_{n,j})\|_2\Big)\sqrt{2\log d}$$
$$= \frac{1}{n}\|X\|_{2,\infty}\sqrt{2\log d} = \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^0\sqrt{2\log d}$$

Inductive step:

$$\hat{\mathfrak{R}}_S(\mathscr{F}_{\ell+1}) = \hat{\mathfrak{R}}_S\left(\left\{x \mapsto \sigma_{\ell+1}\left(\|W_{\ell+1}^\top\|_{1,\infty}\, g(x)\right) : g \in \operatorname{conv}\big(-\mathscr{F}_\ell \cup \mathscr{F}_\ell\big)\right\}\right)$$
$$\le \rho B\, \hat{\mathfrak{R}}_S\left(\operatorname{conv}(-\mathscr{F}_\ell \cup \mathscr{F}_\ell)\right) \quad \hat{\mathfrak{R}}_S(\operatorname{conv}(\mathscr{G})) = \hat{\mathfrak{R}}_S(\mathscr{G})$$

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \Big\| \big(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \big) \Big\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\}) \le \frac{1}{n}\Big(\max_j \|(x_{1,j}, \ldots, x_{n,j})\|_2 \Big)\sqrt{2\log d}$$
$$= \frac{1}{n}\|X\|_{2,\infty}\sqrt{2\log d} = \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^0\sqrt{2\log d}$$

Inductive step:

$$\hat{\mathfrak{R}}_S(\mathscr{F}_{\ell+1}) = \hat{\mathfrak{R}}_S\left(\left\{ x \mapsto \sigma_{\ell+1}\left(\|W_{\ell+1}^\top\|_{1,\infty}\, g(x) \right) : g \in \operatorname{conv}\big(-\mathscr{F}_\ell \cup \mathscr{F}_\ell \big) \right\} \right)$$
$$\le \rho B\, \hat{\mathfrak{R}}_S\left(\operatorname{conv}(-\mathscr{F}_\ell \cup \mathscr{F}_\ell) \right) \quad \hat{\mathfrak{R}}_S(\operatorname{conv}(\mathscr{G})) = \hat{\mathfrak{R}}_S(\mathscr{G})$$
$$\le \rho B\, \hat{\mathfrak{R}}_S\left(-\mathscr{F}_\ell \cup \mathscr{F}_\ell \right)$$

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \Big\| \big(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \big) \Big\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\}) \le \frac{1}{n}\Big(\max_j \|(x_{1,j}, \ldots, x_{n,j})\|_2 \Big)\sqrt{2\log d}$$
$$= \frac{1}{n}\|X\|_{2,\infty}\sqrt{2\log d} = \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^0\sqrt{2\log d}$$

Inductive step:

$$\hat{\mathfrak{R}}_S(\mathscr{F}_{\ell+1}) = \hat{\mathfrak{R}}_S\left(\left\{ x \mapsto \sigma_{\ell+1}\left(\|W_{\ell+1}^\top\|_{1,\infty}\, g(x) \right) : g \in \operatorname{conv}\big(-\mathscr{F}_\ell \cup \mathscr{F}_\ell \big) \right\} \right)$$
$$\le \rho B\, \hat{\mathfrak{R}}_S\left(\operatorname{conv}(-\mathscr{F}_\ell \cup \mathscr{F}_\ell) \right) \quad \hat{\mathfrak{R}}_S(\operatorname{conv}(\mathscr{G})) = \hat{\mathfrak{R}}_S(\mathscr{G})$$
$$\le \rho B\, \hat{\mathfrak{R}}_S\left(-\mathscr{F}_\ell \cup \mathscr{F}_\ell \right) \quad \hat{\mathfrak{R}}_S(A \cup B) \le \hat{\mathfrak{R}}_S(A) + \hat{\mathfrak{R}}_S(B) \text{ if } 0 \in A, 0 \in B$$

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathcal{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathcal{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \Big\| \big(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \big) \Big\|_c$$

Base case, $L = 0$:

$$\hat{\mathfrak{R}}_S(\{x \mapsto x_j : j \in [d]\}) \le \frac{1}{n}\Big(\max_j \|(x_{1,j}, \ldots, x_{n,j})\|_2 \Big)\sqrt{2\log d}$$
$$= \frac{1}{n}\|X\|_{2,\infty}\sqrt{2\log d} = \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^0\sqrt{2\log d}$$

Inductive step:

$$\hat{\mathfrak{R}}_S(\mathcal{F}_{\ell+1}) = \hat{\mathfrak{R}}_S\left(\left\{ x \mapsto \sigma_{\ell+1}\left(\|W_{\ell+1}^\top\|_{1,\infty}\, g(x) \right) : g \in \operatorname{conv}\big(-\mathcal{F}_\ell \cup \mathcal{F}_\ell \big) \right\} \right)$$
$$\le \rho B\, \hat{\mathfrak{R}}_S\left(\operatorname{conv}(-\mathcal{F}_\ell \cup \mathcal{F}_\ell) \right) \quad \hat{\mathfrak{R}}_S(\operatorname{conv}(\mathcal{G})) = \hat{\mathfrak{R}}_S(\mathcal{G})$$
$$\le \rho B\, \hat{\mathfrak{R}}_S\left(-\mathcal{F}_\ell \cup \mathcal{F}_\ell \right) \quad \hat{\mathfrak{R}}_S(A \cup B) \le \hat{\mathfrak{R}}_S(A) + \hat{\mathfrak{R}}_S(B) \text{ if } 0 \in A, 0 \in B$$
$$\le 2\rho B\, \hat{\mathfrak{R}}_S\left(\mathcal{F}_\ell \right)$$

Rademacher of convex hull

$$\hat{\mathfrak{R}}_S(\mathrm{conv}(\mathcal{G})) = \frac{1}{n}\mathbb{E}_\sigma \sup_{k\geq 1} \sup_{\alpha\in\Delta_k} \sup_{g_1,\ldots,g_k\in\mathcal{G}} \left\langle \sigma, \sum_{j=1}^{k} \alpha_j (g_j)_S \right\rangle$$

Rademacher of convex hull

$$\hat{\mathfrak{R}}_S(\mathrm{conv}(\mathcal{G})) = \frac{1}{n}\mathbb{E}_\sigma \sup_{k\geq 1} \sup_{\alpha\in\Delta_k} \sup_{g_1,\ldots,g_k\in\mathcal{G}} \left\langle \sigma, \sum_{j=1}^{k} \alpha_j (g_j)_S \right\rangle$$

$$= \frac{1}{n}\mathbb{E}_\sigma \sup_{k\geq 1} \sup_{\alpha\in\Delta_k} \sum_{j=1}^{k} \alpha_j \sup_{g_j\in\mathcal{G}} \left\langle \varepsilon, (g_j)_S \right\rangle$$

Rademacher of convex hull

$$\hat{\mathfrak{R}}_S(\mathrm{conv}(\mathcal{G})) = \frac{1}{n}\mathbb{E}_\sigma \sup_{k\geq 1} \sup_{\alpha\in\Delta_k} \sup_{g_1,\ldots,g_k\in\mathcal{G}} \left\langle \sigma, \sum_{j=1}^{k} \alpha_j (g_j)_S \right\rangle$$

$$= \frac{1}{n}\mathbb{E}_\sigma \sup_{k\geq 1} \sup_{\alpha\in\Delta_k} \sum_{j=1}^{k} \alpha_j \sup_{g_j\in\mathcal{G}} \left\langle \varepsilon, (g_j)_S \right\rangle$$

$$= \frac{1}{n}\mathbb{E}_\sigma \left(\sup_{k\geq 1} \sup_{\alpha\in\Delta_k} \sum_{j=1}^{k} \alpha_j \right) \sup_{g\in\mathcal{G}} \left\langle \varepsilon, g_S \right\rangle$$

Rademacher of convex hull

$$\hat{\mathfrak{R}}_S(\mathrm{conv}(\mathcal{G})) = \frac{1}{n}\mathbb{E}_\sigma \sup_{k\geq 1} \sup_{\alpha\in\Delta_k} \sup_{g_1,\ldots,g_k\in\mathcal{G}} \left\langle \sigma, \sum_{j=1}^{k} \alpha_j (g_j)_S \right\rangle$$

$$= \frac{1}{n}\mathbb{E}_\sigma \sup_{k\geq 1} \sup_{\alpha\in\Delta_k} \sum_{j=1}^{k} \alpha_j \sup_{g_j\in\mathcal{G}} \left\langle \varepsilon, (g_j)_S \right\rangle$$

$$= \frac{1}{n}\mathbb{E}_\sigma \left(\sup_{k\geq 1} \sup_{\alpha\in\Delta_k} \sum_{j=1}^{k} \alpha_j \right) \sup_{g\in\mathcal{G}} \left\langle \varepsilon, g_S \right\rangle$$

$$= \hat{\mathfrak{R}}_S(\mathcal{G})$$

Rademacher of union

$$\hat{\mathfrak{R}}_S(\mathcal{G} \cup \mathcal{H}) = \frac{1}{n}\mathbb{E}_\sigma \sup_{g \in (\mathcal{G} \cup \mathcal{H})} \langle \sigma, g_S \rangle$$

Rademacher of union

$$\hat{\mathfrak{R}}_S(\mathcal{G} \cup \mathcal{H}) = \frac{1}{n}\mathbb{E}_\sigma \sup_{g \in (\mathcal{G} \cup \mathcal{H})} \langle \sigma, g_S \rangle$$

$$\leq \frac{1}{n}\mathbb{E}_\sigma \Big[\sup_{g \in \mathcal{G}} \langle \sigma, g_S \rangle + \sup_{g \in \mathcal{H}} \langle \sigma, g_S \rangle \Big] \quad \text{if } 0 \in \mathcal{G}, 0 \in \mathcal{H}$$

Rademacher of union

$$\hat{\mathfrak{R}}_S(\mathcal{G} \cup \mathcal{H}) = \frac{1}{n}\mathbb{E}_\sigma \sup_{g \in (\mathcal{G} \cup \mathcal{H})} \langle \sigma, g_S \rangle$$

$$\leq \frac{1}{n}\mathbb{E}_\sigma \Big[\sup_{g \in \mathcal{G}} \langle \sigma, g_S \rangle + \sup_{g \in \mathcal{H}} \langle \sigma, g_S \rangle \Big] \quad \text{if } 0 \in \mathcal{G}, 0 \in \mathcal{H}$$

$$= \hat{\mathfrak{R}}_S(\mathcal{G}) + \hat{\mathfrak{R}}_S(\mathcal{H})$$

Rademacher of union

$$\hat{\mathfrak{R}}_S(\mathcal{G} \cup \mathcal{H}) = \frac{1}{n}\mathbb{E}_\sigma \sup_{g \in (\mathcal{G} \cup \mathcal{H})} \langle \sigma, g_S \rangle$$

$$\leq \frac{1}{n}\mathbb{E}_\sigma \Big[\sup_{g \in \mathcal{G}} \langle \sigma, g_S \rangle + \sup_{g \in \mathcal{H}} \langle \sigma, g_S \rangle \Big] \quad \text{if } 0 \in \mathcal{G}, 0 \in \mathcal{H}$$

$$= \hat{\mathfrak{R}}_S(\mathcal{G}) + \hat{\mathfrak{R}}_S(\mathcal{H})$$

or if both sets are **symmetric**:
for all $g \in \mathcal{G}$, also have $-g \in \mathcal{G}$

Rademacher of union

$$\hat{\mathfrak{R}}_S(\mathcal{G} \cup \mathcal{H}) = \frac{1}{n}\mathbb{E}_\sigma \sup_{g \in (\mathcal{G} \cup \mathcal{H})} \langle \sigma, g_S \rangle$$

$$\leq \frac{1}{n}\mathbb{E}_\sigma \Big[\sup_{g \in \mathcal{G}} \langle \sigma, g_S \rangle + \sup_{g \in \mathcal{H}} \langle \sigma, g_S \rangle \Big] \quad \text{if } 0 \in \mathcal{G}, 0 \in \mathcal{H}$$

$$= \hat{\mathfrak{R}}_S(\mathcal{G}) + \hat{\mathfrak{R}}_S(\mathcal{H})$$

or if both sets are **symmetric**:
for all $g \in \mathcal{G}$, also have $-g \in \mathcal{G}$

…or if we otherwise know that

$$\sup_{g \in \mathcal{G}} \langle \sigma, g_S \rangle \geq 0, \quad \sup_{g \in \mathcal{H}} \langle \sigma, g_S \rangle \geq 0$$

for **any** assignment of σ

Rademacher of deep nets

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathcal{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathcal{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \left\|\left(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b\right)\right\|_c$$

$\sigma_L(W_L\,\sigma_{L-1}(\cdots\ \sigma_1(W_1 x))\|\|_6$

take $\sigma_L = \mathrm{id}$, $W_L = \cdots = W_1 = W$, this is W^L

Rademacher of deep nets

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \left\| \left(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \right) \right\|_c$$

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each 1-Lipschitz, positive homogenous ($\sigma_\ell(ax) = a\sigma_\ell(x)$ for $a > 0$).
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell\|_F \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_F B^L \left(1 + \sqrt{2L\log 2}\right)$.

(More complicated proof: Golowich/Rakhlin/Shamir, COLT 2018 / Telgarsky's 14.2.)

Rademacher of deep nets

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each ρ-Lipschitz with $\sigma_\ell(0) = 0$.
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell^\top\|_{1,\infty} \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_{2,\infty}(2\rho B)^L\sqrt{2\log d}$.

$$\|M\|_{b,c} = \left\| \left(\|M_{\cdot 1}\|_b, \ldots, \|M_{\cdot d}\|_b \right) \right\|_c$$

Theorem: Fix $\sigma_1, \ldots, \sigma_L$ each 1-Lipschitz, positive homogenous ($\sigma_\ell(ax) = a\sigma_\ell(x)$ for $a > 0$).
Let $\mathscr{F}_L$ be the set of L-layer no-intercept nets, $f^{(\ell)} = \sigma_\ell(W_\ell f^{(\ell-1)})$,
with $\|W_\ell\|_F \le B$. Then $\hat{\mathfrak{R}}_n(\mathscr{F}) \le \frac{1}{n}\|X\|_F B^L \left(1 + \sqrt{2L\log 2}\right)$.

(More complicated proof: Golowich/Rakhlin/Shamir, COLT 2018 / Telgarsky's 14.2.)

Can get a slightly better rate via covering numbers: see Telgarsky's section 16.2.

So, does this solve it?

- Experiment by Dziugaite/Roy (2017):
  training a small network on MNIST (0-4 vs 5-9),
  plotting a Rademacher-based margin bound
  using a different (but similarly[?] tight)
  upper bound on the Rademacher complexity

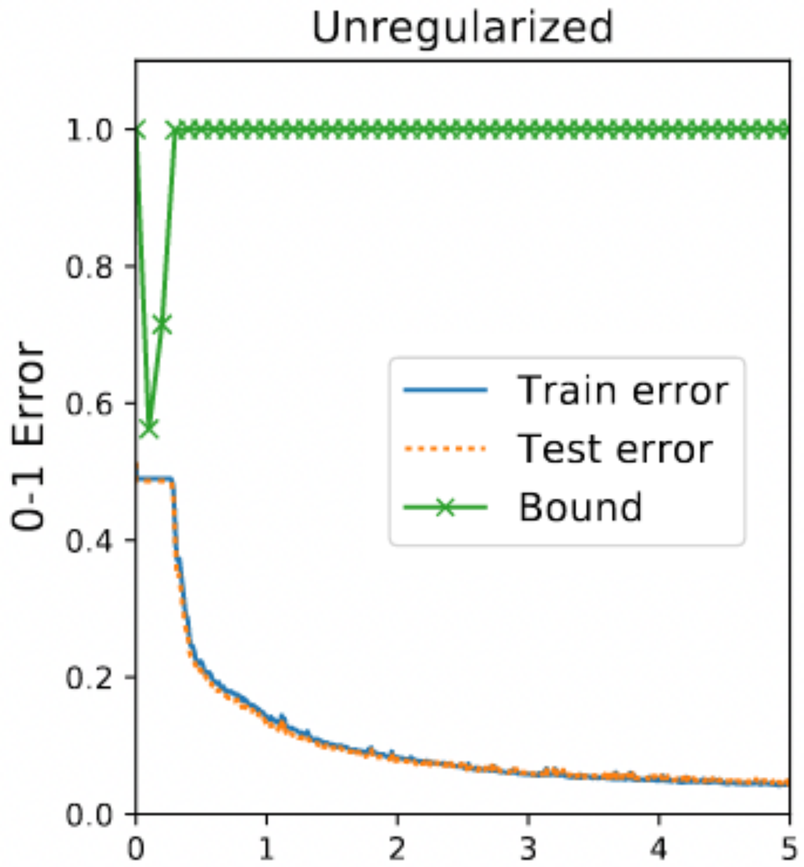

What's left?

- We've shown some bounds on **approximation** and **generalization** (each with significant limitations)

- If we could run ERM, this combination could be enough
- …but ERM is NP-hard even for square loss even with *one* ReLU

- What's the **optimization** error for SGD/similar?

Nonconvex optimization

- Neural nets are not convex

Nonconvex optimization

- Neural nets are not convex
- Even **deep linear networks** are not convex

Nonconvex optimization

- Neural nets are not convex
- Even **deep linear networks** are not convex

- But we do know that SGD converges to a *critical point* under fairly mild conditions

Nonconvex optimization

- Neural nets are not convex
- Even **deep linear networks** are not convex

- But we do know that SGD converges to a *critical point* under fairly mild conditions
  - e.g.: if $f \geq f^{\inf}$ is differentiable and β-smooth, and there are A, B, C s.t. for all x, $\mathbb{E}\left[\|\hat{g}(x)\|^2\right] \leq 2A(f(x) - f^{\inf}) + B\|\nabla f(X)\|^2 + C$, then the *best* iterate has $\mathbb{E}\left[\|\nabla f(x)\|^2\right] \leq \varepsilon^2$ in $\mathcal{O}(\varepsilon^{-4})$ steps (Khaled/Richtárik 2020)

Nonconvex optimization

- Neural nets are not convex
- Even **deep linear networks** are not convex

- But we do know that SGD converges to a *critical point* under fairly mild conditions
  - e.g.: if $f \geq f^{\inf}$ is differentiable and β-smooth, and there are A, B, C s.t. for all x, $\mathbb{E}\left[\|\hat{g}(x)\|^2\right] \leq 2A(f(x) - f^{\inf}) + B\|\nabla f(X)\|^2 + C$, then the *best* iterate has $\mathbb{E}\left[\|\nabla f(x)\|^2\right] \leq \varepsilon^2$ in $\mathcal{O}(\varepsilon^{-4})$ steps (Khaled/Richtárik 2020)

- In deep linear nets, local minima are global minima (Kawaguchi 2016, Laurent/von Brecht 2019)

Nonconvex optimization

- Neural nets are not convex
- Even **deep linear networks** are not convex

- But we do know that SGD converges to a *critical point* under fairly mild conditions
  - e.g.: if $f \geq f^{\inf}$ is differentiable and β-smooth, and there are A, B, C s.t. for all x, $\mathbb{E}\left[\|\hat{g}(x)\|^2\right] \leq 2A(f(x) - f^{\inf}) + B\|\nabla f(X)\|^2 + C$, then the *best* iterate has $\mathbb{E}\left[\|\nabla f(x)\|^2\right] \leq \varepsilon^2$ in $\mathcal{O}(\varepsilon^{-4})$ steps (Khaled/Richtárik 2020)

- In deep linear nets, local minima are global minima (Kawaguchi 2016, Laurent/von Brecht 2019)
  - …but there are saddle points, including "bad" ones where $\lambda_{\min}(\nabla^2 f) = 0$

Nonconvex optimization

- Neural nets are not convex
- Even **deep linear networks** are not convex

- But we do know that SGD converges to a *critical point* under fairly mild conditions
  - e.g.: if $f \geq f^{\inf}$ is differentiable and β-smooth, and there are A, B, C s.t. for all x, $\mathbb{E}\left[\|\hat{g}(x)\|^2\right] \leq 2A(f(x) - f^{\inf}) + B\|\nabla f(X)\|^2 + C$, then the *best* iterate has $\mathbb{E}\left[\|\nabla f(x)\|^2\right] \leq \varepsilon^2$ in $\mathcal{O}(\varepsilon^{-4})$ steps (Khaled/Richtárik 2020)

- In deep linear nets, local minima are global minima (Kawaguchi 2016, Laurent/von Brecht 2019)
  - …but there are saddle points, including "bad" ones where $\lambda_{\min}(\nabla^2 f) = 0$
  - …but gradient descent doesn't get stuck, under some conditions (Arora et al. 2019)

Nonconvex optimization

- Neural nets are not convex
- Even **deep linear networks** are not convex

- But we do know that SGD converges to a *critical point* under fairly mild conditions
  - e.g.: if $f \geq f^{\inf}$ is differentiable and β-smooth, and there are A, B, C s.t. for all x, $\mathbb{E}\left[\|\hat{g}(x)\|^2\right] \leq 2A(f(x) - f^{\inf}) + B\|\nabla f(X)\|^2 + C$, then the *best* iterate has $\mathbb{E}\left[\|\nabla f(x)\|^2\right] \leq \varepsilon^2$ in $\mathcal{O}(\varepsilon^{-4})$ steps (Khaled/Richtárik 2020)

- In deep linear nets, local minima are global minima (Kawaguchi 2016, Laurent/von Brecht 2019)
  - …but there are saddle points, including "bad" ones where $\lambda_{\min}(\nabla^2 f) = 0$
  - …but gradient descent doesn't get stuck, under some conditions (Arora et al. 2019)

- **Next time**: optimization "acts convex" in the **neural tangent kernel** regime