

Learning with norm-based neural networks: model capacity, function spaces, and computational-statistical gaps

Fanghui LIU

fanghui.liu@warwick.ac.uk


Department of Computer Science, University of Warwick, UK
Centre for Discrete Mathematics and its Applications (DIMAP), Warwick
[joint work with Leello Dadi, Zhenyu Zhu, Volkan Cevher (EPFL)]

at INRIA, Paris, 2024




The
Alan Turing
Institute

In the era of deep learning



```
... This code is not working like I expect - how do I fix it?  
resnet152 = resnet152(pretrained=True)  
for name, param in resnet152.named_parameters():  
    if param.requires_grad: param.data *= 0.01  
optimizer = optim.Adam(resnet152.parameters())  
trainer = GradientDescentTrainer(resnet152, optimizer)  
trainer.train()  
return resnet152
```

... It's difficult to say what's wrong with the code without more context. Can you provide more information about what the code is supposed to do and what isn't working as expected? Also, is this the entire code or just a part of it?



Scaling law: under compute budget

scaling law [13]

test loss = A **Model Size** a + B Data Size b + C

under limited compute budget

- data-parameter trade-o
- time-space trade-o






Figure 1: Scaling law under compute-optimal configuration [21].

Model size is a “right” complexity?

- double descent [4] (Belkin, Hsu, Ma, Mandal, 2019)



(a) Results on ResNet18 [18]




(b) Optimal early stopping [18].


- Empirically: neural network pruning [16], lottery ticket hypothesis [11], fine-tuning with large dropout [28]
- Theoretically: how much over-parameterization is sufficient? [7, 26]

Model size is a “right” complexity?

- double descent [4] (Belkin, Hsu, Ma, Mandal, 2019)



(a) Results on ResNet18 [18]



(b) Optimal early stopping [18].

- Empirically: neural network pruning [16], lottery ticket hypothesis [11], fine-tuning with large dropout [28]
- Theoretically: how much over-parameterization is sufficient? [7, 26]

What is the “right” model complexity?

Complexity of a prediction rule, e.g.,

- number of parameters
- norm of parameters

[2] (Bartlett, 1998)

The size of the weights is more important than the size of the network!

Norm-based capacity: [19, 24, 20, 8]

name	definition	rank correlation
Parameter Frobenius norm	$\sum_{i=1}^L \sum_k W_{ik}^2$	0.073
Frobenius distance to initialization [17]	$\sum_{i=1}^L \sum_k W_{ik}^2 - W_{ik}^0$	0.263
Spectral complexity [3]	$\sum_{i=1}^L \sum_k W_{ik}^2$ (with $2=3$)	0.537
Fisher-Rao [14]	$\frac{(L+1)^2}{n} \sum_{i=1}^n \sum_{j=1}^L W_{ij}^2 (h_W(x_i); y_i)$	0.078
Path-norm [19]	$\sum_{j=1}^L \sum_{i=1}^n W_{ij}^2$	0.373

Table 1: Complexity measures compared in the empirical study [12], and their correlation with generalization

What is the “right” model complexity?

Complexity of a prediction rule, e.g.,

- number of parameters
- norm of parameters

[2] (Bartlett, 1998)

The size of the weights is more important than the size of the network!

Norm-based capacity: [19, 24, 20, 8]

name	definition	rank correlation
Parameter Frobenius norm	$\sum_{i=1}^L \sum_k W_{ik}^2$	0.073
Frobenius distance to initialization [17]	$\sum_{i=1}^L \sum_k W_{ik}^2 - W_{ik}^0$	0.263
Spectral complexity [3]	$\sum_{i=1}^L \sum_k W_{ik}^2$ (with $2=3$)	0.537
Fisher-Rao [14]	$\frac{(L+1)^2}{n} \sum_{i=1}^n \sum_{j=1}^L W_{ij}^2$	0.078
Path-norm [19]	$\sum_{(i_1, \dots, i_L)} \sum_{j=1}^L W_{i_j, j}$	0.373

Table 1: Complexity measures compared in the empirical study [12], and their correlation with generalization

What is the “right” model complexity?

Complexity of a prediction rule, e.g.,

- number of parameters
- norm of parameters

[2] (Bartlett, 1998)

The size of the weights is more important than the size of the network!

Norm-based capacity: [19, 24, 20, 8]

name	definition	rank correlation
Parameter Frobenius norm	$\sum_{i=1}^L \sum_k W_{ik}^2$	0.073
Frobenius distance to initialization [17]	$\sum_{i=1}^L \sum_k W_{ik}^0 - W_{ik}^2$	0.263
Spectral complexity [3]	$\sum_{i=1}^L \sum_k W_{ik}^2$ (with $2=3$)	0.537
Fisher-Rao [14]	$\frac{(L+1)^2}{n} \sum_{i=1}^n \sum_{j=1}^L W_{ij}^2$	0.078
Path-norm [19]	$\sum_{j=1}^L \sum_{i=1}^n W_{ij}^2$	0.373

Table 1: Complexity measures compared in the empirical study [12], and their correlation with generalization

What is the “right” model complexity?

Complexity of a prediction rule, e.g.,

- number of parameters
- norm of parameters

[2] (Bartlett, 1998)


The size of the weights is more important than the size of the network!

Norm-based capacity: [19, 24, 20, 8]

name	definition	rank correlation
Parameter Frobenius norm	$\sum_{i=1}^L \sum_{k=1}^K W_{ik}^2$	0.073
Frobenius distance to initialization [17]	$\sum_{i=1}^L \sum_{k=1}^K W_{ik}^2 - \sum_{i=1}^L \sum_{k=1}^K W_{ik}^0$	0.263
Spectral complexity [3]	$\sum_{i=1}^L \sum_{k=1}^K W_{ik}^2 - \frac{(\sum_{i=1}^L \sum_{k=1}^K W_{ik}^2)^2}{\sum_{i=1}^L \sum_{k=1}^K W_{ik}^4}$	0.537
Fisher-Rao [14]	$\frac{(L+1)^2}{\mathcal{P}} \sum_{i=1}^n h_{W_i}^T r_{W_i} \cdot (h_{W_i}(x_i); y_i)$	0.078
Path-norm [19]	$\sum_{(i_0, \dots, i_L)} \sum_{j=1}^L W_{i_j; i_{j-1}}^2$	0.373

Table 1: Complexity measures compared in the empirical study [12], and their correlation with generalization

Two-layer neural networks, path norm




$$P_m = f(\cdot) := \frac{1}{m} \sum_{k=1}^m a_k \cdot h(w_k; i)$$

ℓ_1 -path norm

$$\|k\|_{k_P} := \frac{1}{m} \sum_{k=1}^m |a_k| \cdot |w_k|$$

- semi-norm
- representation cost
- relations to Barron spaces $B [1, 10]$
- $\|k\|_{k_B} \leq \|k\|_{k_P} \leq 2\|k\|_{k_B}$

Two-layer neural networks, path norm




$$P_m = f(\cdot) := \frac{1}{m} \sum_{k=1}^m a_k \quad \|w_k\|_1$$

ℓ_1 -path norm

$$\|k\|_{k_P} := \frac{1}{m} \sum_{k=1}^m |a_k| \|w_k\|_1$$

- semi-norm
- representation cost
- relations to Barron spaces B [1, 10]
- $\|k\|_{k_P} \leq 2\|k\|_B$

Two-layer neural networks, path norm




$$P_m = f(\cdot) := \frac{1}{m} \sum_{k=1}^m a_k \cdot \|w_k\|_1$$

ℓ_1 -path norm

$$\|k\|_{k_P} := \frac{1}{m} \sum_{k=1}^m \|a_k\|_1 \|w_k\|_1$$

- semi-norm
- representation cost
- relations to Barron spaces B [1, 10]
- $\|k\|_{k_P} \leq 2\|k\|_B$

Two-layer neural networks, path norm



$$P_m = f(\cdot) := \frac{1}{m} \sum_{k=1}^m a_k \quad \|w_k\|_1$$

ℓ_1 -path norm

$$\|k\|_{k_P} := \frac{1}{m} \sum_{k=1}^m |a_k| \|w_k\|_1$$

- semi-norm
- representation cost
- relations to Barron spaces B [1, 10]
- $\|k\|_{k_P} \leq 2\|k\|_B$

Path norm, Barron spaces, RKHS

Consider a random features model [22, 15]

- first layer: $w \stackrel{iid}{\sim} P(W)$; only train the second layer

infinite many features $f_a(x) = \int_W a(w) \phi(x; w) d(w)$

Definition (RKHS and Barron space [1, 1])

$$F_{p; \gamma} := \{f_a : \|a\|_{L^p(\gamma)} < 1\}; \quad \|f\|_{F_{p; \gamma}} := \inf_{f=f_a} \|a\|_{L^p(\gamma)}$$

For any $1 \leq p < q$, we have

$$B = \left[\int P(W) F_{p; \gamma}; \quad \|f\|_B = \inf_{P(W)} \|f\|_{F_{p; \gamma}} \right]$$

- RFMs $\hat{=}$ kernel methods by taking $p = 2$ using Representer theorem [23]
- RFMs $\hat{=}$ kernel methods if $p < 2$
- function space: $F_1 \subset F_p \subset F_q \subset F_1$; if $p < q$

Path norm, Barron spaces, RKHS

Consider a random features model [22, 15]

- first layer: $w \stackrel{iid}{\sim} P(W)$; only train the second layer

infinite many features $f_a(x) = \int_W a(w) \phi(x; w) d(w)$

Definition (RKHS and Barron space [1, 1])

$$F_{p; P} := \{f_a : \|a\|_{L^p(W)} < 1\}; \quad \|f\|_{F_{p; P}} := \inf_{f=f_a} \|a\|_{L^p(W)}$$

For any $1 \leq p < q$, we have

$$B = \int_{2P(W)} F_{p; P}; \quad \|f\|_B = \inf_{2P(W)} \|f\|_{F_{p; P}}$$

- RFMs \approx kernel methods by taking $p = 2$ using Representer theorem [23]
- RFMs $\not\approx$ kernel methods if $p < 2$
- function space: $F_1 \subset F_p \subset F_q \subset F_1$; if $p < q$

Path norm, Barron spaces, RKHS

Consider a random features model [22, 15]

- first layer: $w \stackrel{iid}{\sim} P(W)$; only train the second layer

infinite many features $f_a(x) = \int_W a(w) \phi(x; w) d(w)$

Definition (RKHS and Barron space [9, 5])

$$F_{p; \gamma} := \{f_a : \|a\|_{L^p(\gamma)} < \gamma\}; \quad \|f\|_{F_{p; \gamma}} := \inf_{f=f_a} \|a\|_{L^p(\gamma)}$$

For any $1 < p < \infty$, we have

$$B = \left[\int P(W) F_{p; \gamma}; \quad \|f\|_B = \inf_{\int P(W)} \|f\|_{F_{p; \gamma}} \right]$$

- RFMs $\hat{=}$ kernel methods by taking $p = 2$ using Representer theorem [23]
- RFMs $\hat{=}$ kernel methods if $p < 2$
- function space: $F_1 \subset F_p \subset F_q \subset F_\infty$ if $p < q$

Path norm, Barron spaces, RKHS

Consider a random features model [22, 15]

- first layer: $w \stackrel{iid}{\sim} P(W)$; only train the second layer

infinite many features $f_a(x) = \int_{\mathcal{W}} a(w) \phi(x; w) d(w)$

Definition (RKHS and Barron space [9, 5])

$$F_{p; P} := \{f_a : \|a\|_{L^p(\mathcal{W})} < 1\}; \quad \|f\|_{F_{p; P}} := \inf_{f=f_a} \|a\|_{L^p(\mathcal{W})}$$

For any $1 \leq p \leq 2$, we have

$$B = \left[\int_{\mathcal{W}} P(dw) F_{p; P} \right]; \quad \|f\|_B = \inf_{f \in F_{p; P}} \|f\|_{F_{p; P}}$$

- RFMs $\hat{=}$ kernel methods by taking $p = 2$ using Representer theorem [23]
- RFMs $\hat{=}$ kernel methods if $p < 2$
- function space: $F_1 \subset F_p \subset F_q \subset F_1$ if $p < q$

Path norm, Barron spaces, RKHS

Consider a random features model [22, 15]

- first layer: $w \stackrel{iid}{\sim} P(W)$; only train the second layer

infinite many features $f_a(x) = \int_{\mathcal{W}} a(w) \phi(x; w) d(w)$

Definition (RKHS and Barron space [9, 5])

$$F_p := \{f_a : \|a\|_{L^p(\cdot)} < 1\}; \quad \|f\|_{F_p} := \inf_{f=f_a} \|a\|_{L^p(\cdot)}$$

For any $1 < p < \infty$, we have

$$B = \left[\int_{\mathcal{W}} F_p; \quad \|f\|_B = \inf_{\int_{\mathcal{W}} F_p} \|f\|_{F_p} \right]$$

- RFMs $\hat{=}$ kernel methods by taking $p = 2$ using Representer theorem [23]
- RFMs $\hat{\neq}$ kernel methods if $p < 2$
- function space: $F_1 \subset F_p \subset F_q \subset F_1$; if $p < q$

Our results: statistical guarantees

For the class of two-layer neural networks $G_R = \{f : \mathbb{R}^d \rightarrow \mathbb{R} \mid f(x) = \sum_{i=1}^m \max(0, w_i \cdot x) + b\}$

$$\hat{f} := \operatorname{argmin}_{f \in G_R} \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i))^2 :$$

Theorem (Liu, Dadi, Cevher, JMLR 2024)

Under standard assumptions (bounded data, $f^* \in \mathcal{L}_x^2$), for two-layer over-parameterized neural networks, we have

$$\|\hat{f} - f^*\|_{L^2_x}^2 \leq \frac{R^2}{m} + R^2 d^{\frac{1}{3}} n^{-\frac{d+2}{2d+2}} \quad \text{w.h.p.}$$

$n^{-\frac{d+2}{2d+2}}$ is always faster than $n^{-\frac{1}{2}}$: No curse of dimensionality!

Our results: statistical guarantees

For the class of two-layer neural networks $G_R = \{f : \mathbb{R}^d \rightarrow \mathbb{R} \mid f(x) = \sum_{i=1}^m \max(0, w_i \cdot x) + b\}$

$$\hat{f} := \operatorname{argmin}_{f \in G_R} \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i))^2 :$$

Theorem (Liu, Dadi, Cevher, JMLR 2024)

Under standard assumptions (bounded data, $f^* \in B$), for two-layer *over-parameterized* neural networks, we have

$$\|\hat{f} - f^*\|_{L^2_X}^2 \leq \frac{R^2}{m} + R^2 d^{\frac{1}{3}} n^{-\frac{d+2}{2d+2}} \quad \text{w.h.p.}$$

$n^{-\frac{d+2}{2d+2}}$ is always faster than $n^{-\frac{1}{2}}$: No curse of dimensionality!

Our results: statistical guarantees

For the class of two-layer neural networks $G_R = \{f : \mathbb{R}^d \rightarrow \mathbb{R} \mid f(x) = \sum_{i=1}^m \max(0, w_i \cdot x) + b\}$

$$\hat{f} := \operatorname{argmin}_{f \in G_R} \frac{1}{n} \sum_{i=1}^n (y_i - f(x_i))^2 :$$

Theorem (Liu, Dadi, Cevher, JMLR 2024)

Under standard assumptions (bounded data, $f^* \in \mathcal{L}_2(B)$), for two-layer *over-parameterized* neural networks, we have

$$\|\hat{f} - f^*\|_{L^2_X}^2 \leq \frac{R^2}{m} + R^2 d^{\frac{1}{3}} n^{-\frac{d+2}{2d+2}} \quad \text{w.h.p.}$$

$n^{-\frac{d+2}{2d+2}}$ is always *faster* than $n^{-\frac{1}{2}}$: No curse of dimensionality!

Sample complexity

Proposition (metric entropy)

For bounded data $\|x\|_1 \leq 1$, denote $G_R = \{f \in P_m : \|f\|_P \leq R\}$, the metric entropy of G_1 can be bounded by

$$\log N_2(G_1; \cdot) \leq C d^{\frac{2d}{d+2}}; \quad \delta > 0 \text{ and } d \geq 5;$$

with some universal constant C independent of d .


Sample complexity

Proposition (metric entropy)

For bounded data $\|x\|_1 \leq 1$, denote $G_R = \{f \in P_m : \|f\|_p \leq R\}$, the metric entropy of G_1 can be bounded by

$$\log N_2(G_1; \cdot) \leq Cd^{\frac{2d}{d+2}}; \quad \delta > 0 \text{ and } d \geq 5;$$

with some universal constant C independent of d .




Sample complexity

Proposition (metric entropy)

For bounded data $\|x\|_1 \leq 1$, denote $G_R = \{f \in P_m : \|f\|_P \leq R\}$, the metric entropy of G_1 can be bounded by

$$\log N_2(G_1; \cdot) \leq Cd^{\frac{2d}{d+2}}; \quad \delta > 0 \text{ and } d \geq 5;$$

with some universal constant C independent of d .



The "best" trade-off between $\log d$ and d .

Optimization in Barron spaces is NP hard: curse of dimensionality!


Computational-to-statistical gaps

Optimization in Barron spaces is NP hard: curse of dimensionality!




Computational-to-statistical gaps

Optimization in Barron spaces is NP hard: curse of dimensionality!



Computational-to-statistical gaps

Optimization in Barron spaces is NP hard: curse of dimensionality!



Do some Barron functions can be learned by two-layer NNs, both statistically and computationally efficient?

Learning with multiple ReLU neurons

Can we learn **multiple ReLU neurons** by two-layer NNs, both statistically and computationally efficient?

$$f^*(x) = \sum_{j=1}^k a_j \max(0, v_j \cdot x); k = O(1)$$

$k^2 \cdot f^*_{K^2(d)}$ from $f^*(x_i); f^*(x_i) g_{i=1}^n$ with $x_i \sim N(0; I_d)$

Theorem ([1]) PAC learning f^* under Gaussian measure)

There exists an *algorithm* that requires time/samples at $(d =)^{O(k^2)}$

- correlational statistical query (CSQ): $j^q \mathbb{E}_{x,y}[\max(0, v_j \cdot x)y]$

Learning with multiple ReLU neurons

Can we learn **multiple ReLU neurons** by two-layer NNs, both statistically and computationally efficient?

$$f^*(x) = \sum_{j=1}^k a_j (h_{v_j}; x); k = O(1)$$

$k^{\frac{1}{2}} \|f^*\|_{L^2(d)}$ from $f(x_i); f^*(x_i) g_{i=1}^n$ with $x_i \sim N(0; I_d)$

Theorem ([1]) PAC learning f^* under Gaussian measure)

There exists an *algorithm* that requires time/samples at $(d =)^{O(k^2)}$

- correlational statistical query (CSQ): $j \epsilon \mathbb{E}_{x,y} [(x) y]$

Learning with multiple ReLU neurons

Can we learn **multiple ReLU neurons** by two-layer NNs, both statistically and computationally efficient?

$$f^*(x) = \sum_{j=1}^k a_j \max(v_j, x_i); k = O(1)$$

$k f^* f^*_{kL^2(d)}$ from $f^* x_i; f^*(x_i) g_{i=1}^n$ with $x_i \sim N(0; I_d)$

Theorem ([1]) PAC learning f^* under Gaussian measure)

There exists an algorithm that requires time/samples at $(d =)^{O(k^2)}$

- correlational statistical query (CSQ): $j q \sim E_{x,y} [\max(x,y)]$

Learning with multiple ReLU neurons

Can we learn **multiple ReLU neurons** by two-layer NNs, both statistically and computationally efficient?

$$f^*(x) = \sum_{j=1}^k a_j \max(v_j, x_i); k = O(1)$$

$\|f^* - \hat{f}\|_{L^2(d)}$ from $\{x_i, f^*(x_i)\}_{i=1}^n$ with $x_i \sim N(0, I_d)$

Theorem ([6] PAC learning f^* under Gaussian measure)

There exists an *algorithm* that requires time/samples at $(d =)^{O(k^2)}$

- correlational statistical query (CSQ): $\int \phi(x) y \, dP_{x,y}$

Learning with multiple ReLU neurons

Can we learn **multiple ReLU neurons** by two-layer NNs, both statistically and computationally efficient?

$$f^*(x) = \sum_{j=1}^k a_j \max(v_j, x_i); k = O(1)$$

$k f^* f^*_{kL^2(d)}$ from $f x_i; f^*(x_i) g_{i=1}^n$ with $x_i \sim N(0; I_d)$

Theorem ([6] PAC learning f^* under Gaussian measure)

There exists an *algorithm* that requires time/samples at $(d =)^{O(k^2)}$

- correlational statistical query (CSQ): $j q \quad E_{x,y}[\max(x,y)]$

How does student(s) become teacher(s) under GD training?

Learning multi ReLU neurons by two-layer NN via online SGD

$$L(W) = \frac{1}{2} \mathbb{E}_{x \sim N(0; I_d)} \sum_{i=1}^n (hw_i; x_i - f^*(x))^2$$

- Gaussian initialization $w_i \sim N(0; \sigma^2 I_d)$
- angle: $\theta_{ij} = \angle(w_i; v_j)$

Assumption

- *diverse teacher neurons: $\{v_j\}_{j=1}^k$ are orthogonal and $\|v_j\|_2 = \text{const}$*
- *warm start: the smallest angle not close to orthogonal*
weak recovery: $\|hw_i; v_i\| \gg \|hw_i; v_j\|$

How does student(s) become teacher(s) under GD training?

Learning multi ReLU neurons by two-layer NN via online SGD

$$L(W) = \frac{1}{2} \mathbb{E}_x \sum_{i=1}^n (h w_i; x_i - f^*(x))^2$$

- Gaussian initialization $w_i \sim N(0; I_d)$
- angle: $\angle_{ij} = \angle(w_i; v_j)$

Assumption

- *diverse teacher neurons: $\{v_j\}_{j=1}^k$ are orthogonal and $\|v_j\|_2 = \text{const}$*
- *warm start: the smallest angle not close to orthogonal*
- *weak recovery: $\langle h w_i; v_i \rangle > \langle h w_i; v_j \rangle$*

How does student(s) become teacher(s) under GD training?

Learning multi ReLU neurons by two-layer NN via online SGD

$$L(W) = \frac{1}{2} \mathbb{E}_x \sum_{i=1}^n \max(0, (w_i; x_i)) - f^*(x)^2$$

- Gaussian initialization $w_i \sim N(0; I_d)$
- angle: $\angle_{ij} = \angle(w_i; v_j)$

Assumption

- *diverse teacher neurons*: $\{v_j\}_{j=1}^k$ are *orthogonal* and $\|v_j\|_2 = \text{const}$
- *warm start*: the *smallest* angle not close to orthogonal
- *weak recovery*: $\langle w_i; v_i \rangle > \langle w_i; v_j \rangle$

How does student(s) become teacher(s) under GD training?

Learning multi ReLU neurons by two-layer NN via online SGD

$$L(W) = \frac{1}{2} E_{x \sim N(0; I_d)} \sum_{i=1}^n (h w_i; x_i - f(x))^2$$

- Gaussian initialization $w_i \sim N(0; I_d)$
- angle: $\angle_{ij} = \angle(w_i; v_j)$

Assumption



- *diverse teacher neurons*: $\{v_j\}_{j=1}^k$ are *orthogonal* and $\|v_j\|_2 = \text{const}$
- *warm start*: the *smallest* angle not close to orthogonal
- *weak recovery*: $\langle h w_i; v_i \rangle > \langle h w_i; v_j \rangle$

How does student(s) become teacher(s) under GD training?

- align $i \neq 0$

norm converge

then fit



Theorem (Zhu, Liu, Cevher, 2024)

For sufficiently small initialization and step-size $\eta = o(m^{-k^2})$, then there exists a time $T_2 = \frac{1}{\eta}$ such that $8T \geq 2N$ and $i \geq 2[m]$,



$$L(W(T + T_2)) = O\left(\frac{1}{T^3}\right); \|kw_i(T + T_2)k_2\| = \frac{kkvk_2}{m} \quad w.h.p.$$

How does student(s) become teacher(s) under GD training?

- align $i \neq 0$

norm converge

then fit



Theorem (Zhu, Liu, Cevher, 2024)

For sufficiently small initialization and step-size ; $= o(m^{-k^2})$, then there exists a time $T_2 = \frac{1}{\epsilon}$ such that $8T \geq 2N$ and $i \geq [m]$,

$$L(W(T + T_2)) \leq \frac{1}{T^3} ; \|w_i(T + T_2)\|_2 = \frac{k\|v\|k_2}{m} \quad w.h.p :$$

Take-away messages

- model size -> size of weights -> path norm -> Barron spaces
- statistical guarantees with improved sample complexity
- computational-statistical gap -> learning with multiple ReLU neurons

We're organizing one workshop at NeurIPS 2024!

Fine-Tuning in Modern Machine Learning: Principles and Scalability

<https://sites.google.com/view/neurips2024-ftw/home>

Take-away messages

- ^ model size -> size of weights -> path norm -> Barron spaces
- ^ statistical guarantees with improved sample complexity
- ^ computational-statistical gap -> learning with multiple ReLU neurons

Fine-Tuning in Modern Machine Learning: Principles and Scalability
<https://sites.google.com/view/neurips2024-ftw/home>

Take-away messages

- ^ model size -> size of weights -> path norm -> Barron spaces
- ^ statistical guarantees with improved sample complexity
- ^ computational-statistical gap -> learning with multiple ReLU neurons

Fine-Tuning in Modern Machine Learning: Principles and Scalability
<https://sites.google.com/view/neurips2024-ftw/home>

Take-away messages

- ^ model size -> size of weights -> path norm -> Barron spaces
- ^ statistical guarantees with improved sample complexity
- ^ computational-statistical gap -> learning with multiple ReLU neurons

Wu, et al. (2024) "Fine-Tuning in Modern Machine Learning: Principles and Scalability"
<https://sites.google.com/view/neurips2024-ftw/home>

Andrew R Barron.

Universal approximation bounds for superpositions of a sigmoidal function.

IEEE Transactions on Information Theory, 39(3):930-945, 1993.

Peter Bartlett.

The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network.

IEEE Transactions on Information Theory, 44(2):525-536, 1998.

Peter Bartlett, Dylan Foster, and Matus Telgarsky.

Spectrally-normalized margin bounds for neural networks.

In Advances in Neural Information Processing Systems, pages 6241-6250, 2017.

Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal.
Reconciling modern machine-learning practice and the classical
bias variance trade-off.

the National Academy of Sciences, 16(32):15849–15854, 2019.

Hongrui Chen, Jihao Long, and Lei Wu.

A duality framework for generalization analysis of random feature
models and two-layer neural networks.

arXiv preprint arXiv:2305.05642, 2023.

Sitan Chen and Shyam Narayanan.

A faster and simpler algorithm for learning shallow networks.

arXiv preprint arXiv:2307.12496, 2023.

- Zixiang Chen, Yuan Cao, Difan Zou, and Quanquan Gu.
How much over-parameterization is sufficient to learn deep relu networks?
In International Conference on Learning Representations, 2020.
- Carles Domingo-Enrich and Youssef Mroueh.
Tighter sparse approximation bounds for relu neural networks.
In International Conference on Learning Representations, 2022.
- Weinan E, Chao Ma, and Lei Wu.
A priori estimates of the population risk for two-layer neural networks.
Communications in Mathematical Sciences, 17(5):1407–1425, 2019.

Weinan E, Chao Ma, and Lei Wu.

The barron space and the low-induced function spaces for neural network models.

Constructive Approximation pages 1-38, 2021.

Jonathan Frankle and Michael Carbin.

The lottery ticket hypothesis: Finding sparse, trainable neural networks.

In International Conference on Learning Representations, 2019.

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, 2020.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei.

Scaling laws for neural language models.

arXiv preprint arXiv:2001.08361, 2020.

Tengyuan Liang, Tomaso Poggio, Alexander Rakhlin, and James Stokes. Fisher-rao metric, geometry, and complexity of neural networks.

In International conference on Artificial Intelligence and Statistics, pages 888–896, 2019.

Fanghui Liu, Xiaolin Huang, Yudong Chen, and Johan AK Suykens. Random features for kernel approximation: A survey on algorithms, theory, and beyond.

IEEE Transactions on Pattern Analysis and Machine Intelligence 44(10):7128-7148, 2021.

Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz.

Pruning convolutional neural networks for resource efficient inference.

In International Conference on Learning Representations, 2017.

Vaishnavh Nagarajan and J Zico Kolter.

Generalization in deep networks: The role of distance from initialization.

arXiv preprint arXiv:1901.01672, 2019.

Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever.

Deep double descent: Where bigger models and more data hurt. In International Conference on Learning Representations, 2019.

Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro.

Norm-based capacity control in neural networks.

In Conference on Learning Theory, pages 1376–1401. PMLR, 2015.

Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro.
A function space view of bounded norm in finite width relu nets:
The multivariate case.

In International Conference on Learning Representations, 2020.

Elliot Paquette, Courtney Paquette, Lechao Xiao, and Jeffrey
Pennington.

4+3 phases of compute-optimal neural scaling laws.

arXiv preprint arXiv:2405.15074, 2024.

Ali Rahimi and Benjamin Recht.

Random features for large-scale kernel machines.

In Advances in Neural Information Processing Systems, pages 1177–1184,
2007.

Ali Rahimi and Benjamin Recht.

Uniform approximation of functions with random bases.

In Annual Allerton Conference on Communication, Control, and Computing pages 555–561. IEEE, 2008.

Pedro Savarese, Itay Evron, Daniel Soudry, and Nathan Srebro.

How do infinite width bounded norm networks look in function space?

In Conference on Learning Theory pages 2667–2690. PMLR, 2019.

Jonathan W Siegel and Jinchao Xu.

Sharp bounds on the approximation rates, metric entropy, and n -widths of shallow neural networks.

arXiv preprint arXiv:2101.12365, 2021.

- Taiji Suzuki, Denny Wu, Kazusato Oko, and Atsushi Nitanda.
Feature learning via mean-eld langevin dynamics: classifying
sparse parities and beyond.
In Advances in Neural Information Processing Systems 2023.
- Lei Wu and Jihao Long.
A spectral-based analysis of the separation between two-layer
neural networks and linear methods.
Journal of Machine Learning Research 19:1 34, 2022.
- Jianyu Zhang and Léon Bottou.
Fine-tuning with very large dropout.
arXiv preprint arXiv:2403.00946 2024.