# From Tempered to Benign Overfitting in ReLU Neural Networks 

Guy Kornowski* Gilad Yehudai* Ohad Shamir<br>Weizmann Institute of Science

Spotlight presentation

Overfitting puzzle

## Overfitting puzzle

Main Question: Why do overparameterized neural networks generalize?

## Overfitting puzzle

Main Question: Why do overparameterized neural networks generalize?

- Even when trained to fit noisy samples, even without regularization...


## Overfitting puzzle

Main Question: Why do overparameterized neural networks generalize?

- Even when trained to fit noisy samples, even without regularization...



## Overfitting puzzle

Main Question: Why do overparameterized neural networks generalize?

- Even when trained to fit noisy samples, even without regularization...



## Overfitting puzzle

Main Question: Why do overparameterized neural networks generalize?

- Even when trained to fit noisy samples, even without regularization...



## Overfitting puzzle

Main Question: Why do overparameterized neural networks generalize?

- Even when trained to fit noisy samples, even without regularization...



## Overfitting puzzle

Main Question: Why do overparameterized neural networks generalize?

- Even when trained to fit noisy samples, even without regularization...

- Seems to defy classical learning theory, "Occam's razor"...


## Setting

## Setting

- Noisy "classification" data: $S=\left(x_{i}, y_{i}\right)_{i=1}^{m} \subset \mathbb{R}^{d} \times\{ \pm 1\}$


## Setting

- Noisy "classification" data: $S=\left(x_{i}, y_{i}\right)_{i=1}^{m} \subset \mathbb{R}^{d} \times\{ \pm 1\}$
- Ground truth $f^{*}\left(x_{i}\right) \equiv 1$, each $y_{i}$ flipped w.p. $p \in\left[0, \frac{1}{2}\right)$



## Setting

- Noisy "classification" data: $S=\left(x_{i}, y_{i}\right)_{i=1}^{m} \subset \mathbb{R}^{d} \times\{ \pm 1\}$
- Ground truth $f^{*}\left(x_{i}\right) \equiv 1$, each $y_{i}$ flipped w.p. $p \in\left[0, \frac{1}{2}\right)$
- Two-layer ReLU neural network $N_{\theta}(x):=\sum_{j=1}^{n} a_{j} \cdot\left[w_{j} \cdot x+b_{j}\right]_{+}$



## Setting

- Noisy "classification" data: $S=\left(x_{i}, y_{i}\right)_{i=1}^{m} \subset \mathbb{R}^{d} \times\{ \pm 1\}$
- Ground truth $f^{*}\left(x_{i}\right) \equiv 1$, each $y_{i}$ flipped w.p. $p \in\left[0, \frac{1}{2}\right)$
- Two-layer ReLU neural network $N_{\theta}(x):=\sum_{j=1}^{n} a_{j} \cdot\left[w_{j} \cdot x+b_{j}\right]_{+}$
- Network interpolates dataset: $y_{i} N_{\theta}\left(x_{i}\right)>0, \forall i \in[m]$



## Types of overfitting (following Mallinare etal. 22)



## Types of overfitting (folowing Mallinare etal. 22)

- Analyze the clean test error $L\left(N_{\theta}\right):=\operatorname{Pr}_{x}\left[N_{\theta}(x) \leq 0\right]$



## Types of overfitting (following Mallinare etal. 22)

- Analyze the clean test error $L\left(N_{\theta}\right):=\operatorname{Pr}_{x}\left[N_{\theta}(x) \leq 0\right]$
- The overfitting is called "benign" if $L\left(N_{\theta}\right) \rightarrow 0$ [Bartlett et al. '20]



## Types of overfitting (following Mallinare etal. 22)

- Analyze the clean test error $L\left(N_{\theta}\right):=\operatorname{Pr}_{x}\left[N_{\theta}(x) \leq 0\right]$
- The overfitting is called "benign" if $L\left(N_{\theta}\right) \rightarrow 0$ [Bartlett et al. '20]
- The overfitting is called "tempered" if $L\left(N_{\theta}\right) \in\left(0, \frac{1}{2}\right)$



## Types of overfitting (folowing Mallinare etal. 22)

- Analyze the clean test error $L\left(N_{\theta}\right):=\operatorname{Pr}_{x}\left[N_{\theta}(x) \leq 0\right]$
- The overfitting is called "benign" if $L\left(N_{\theta}\right) \rightarrow 0$ [Bartlett et al. '20]
- The overfitting is called "tempered" if $L\left(N_{\theta}\right) \in\left(0, \frac{1}{2}\right)$
* Special case of ineterest is when $L\left(N_{\theta}\right)$ scales with $p$, e.g. $L\left(N_{\theta}\right) \approx p$



## Types of overfitting (folowing Mallinare etal. 22)

- Analyze the clean test error $L\left(N_{\theta}\right):=\operatorname{Pr}_{x}\left[N_{\theta}(x) \leq 0\right]$
- The overfitting is called "benign" if $L\left(N_{\theta}\right) \rightarrow 0$ [Bartlett et al. '20]
- The overfitting is called "tempered" if $L\left(N_{\theta}\right) \in\left(0, \frac{1}{2}\right)$
* Special case of ineterest is when $L\left(N_{\theta}\right)$ scales with $p$, e.g. $L\left(N_{\theta}\right) \approx p$
- The overfitting is called "catastrophic" if $L\left(N_{\theta}\right) \rightarrow \frac{1}{2}$



## Main technical tool: Implicit bias

Gradient based training with certain losses (e.g. logistic) drives $\theta$ towards a KKT point of the margin maximization problem

$$
\min \|\theta\|^{2} \quad \text { s.t } \quad y_{i} N_{\theta}\left(x_{i}\right) \geq 1 \forall i \in[m]
$$

[Lyu \& Li '20, Ji \& Telgarsky '20]

From tempered to benign overfitting

## From tempered to benign overfitting

Theorem: In dimension $d=1$, with noise level $p$, w.h.p. over the sample any KKT point $\theta$ satisfies $L\left(N_{\theta}\right) \in\left(p^{5}, \sqrt{p}\right)$.

Moreover, any local minimum of max margin $\theta$ satisfies $L\left(N_{\theta}\right) \approx p$.

## From tempered to benign overfitting

Theorem: In dimension $d=1$, with noise level $p$, w.h.p. over the sample any KKT point $\theta$ satisfies $L\left(N_{\theta}\right) \in\left(p^{5}, \sqrt{p}\right)$.

Moreover, any local minimum of max margin $\theta$ satisfies $L\left(N_{\theta}\right) \approx p$.

Theorem: In dimension $d \gtrsim \operatorname{poly}(m) \log (1 / \epsilon)$, under some assumptions*, w.h.p. over the sample, KKT points $\theta$ satisfy $L\left(N_{\theta}\right) \leq \epsilon$.

Equivalently, $L\left(N_{\theta}\right) \lesssim \exp (-d)$.

## From tempered to benign overfitting

Theorem: In dimension $d=1$, with noise level $p$, w.h.p. over the sample any KKT point $\theta$ satisfies $L\left(N_{\theta}\right) \in\left(p^{5}, \sqrt{p}\right)$.

Moreover, any local minimum of max margin $\theta$ satisfies $L\left(N_{\theta}\right) \approx p$.

Theorem: In dimension $d \gtrsim \operatorname{poly}(m) \log (1 / \epsilon)$, under some assumptions*, w.h.p. over the sample, KKT points $\theta$ satisfy $L\left(N_{\theta}\right) \leq \epsilon$. Equivalently, $L\left(N_{\theta}\right) \lesssim \exp (-d)$.

## Empirical study of intermediate dimensions

## Empirical study of intermediate dimensions



## Empirical study of intermediate dimensions




## Thanks!

