Benign, Tempered, or Catastrophic: A Taxonomy of Overfitting

Neil Mallinar*
UC San Diego

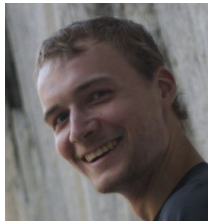

James B. Simon*
UC Berkeley

Amirhesam Abedsoltan
UC San Diego

Parthe Pandit
UC San Diego

Mikhail Belkin
UC San Diego

Preetum Nakkiran
Apple

*Co-first author

paper at tinyurl.com/TemperedOverfitting

NeurIPS 2022

A taxonomy of overfitting

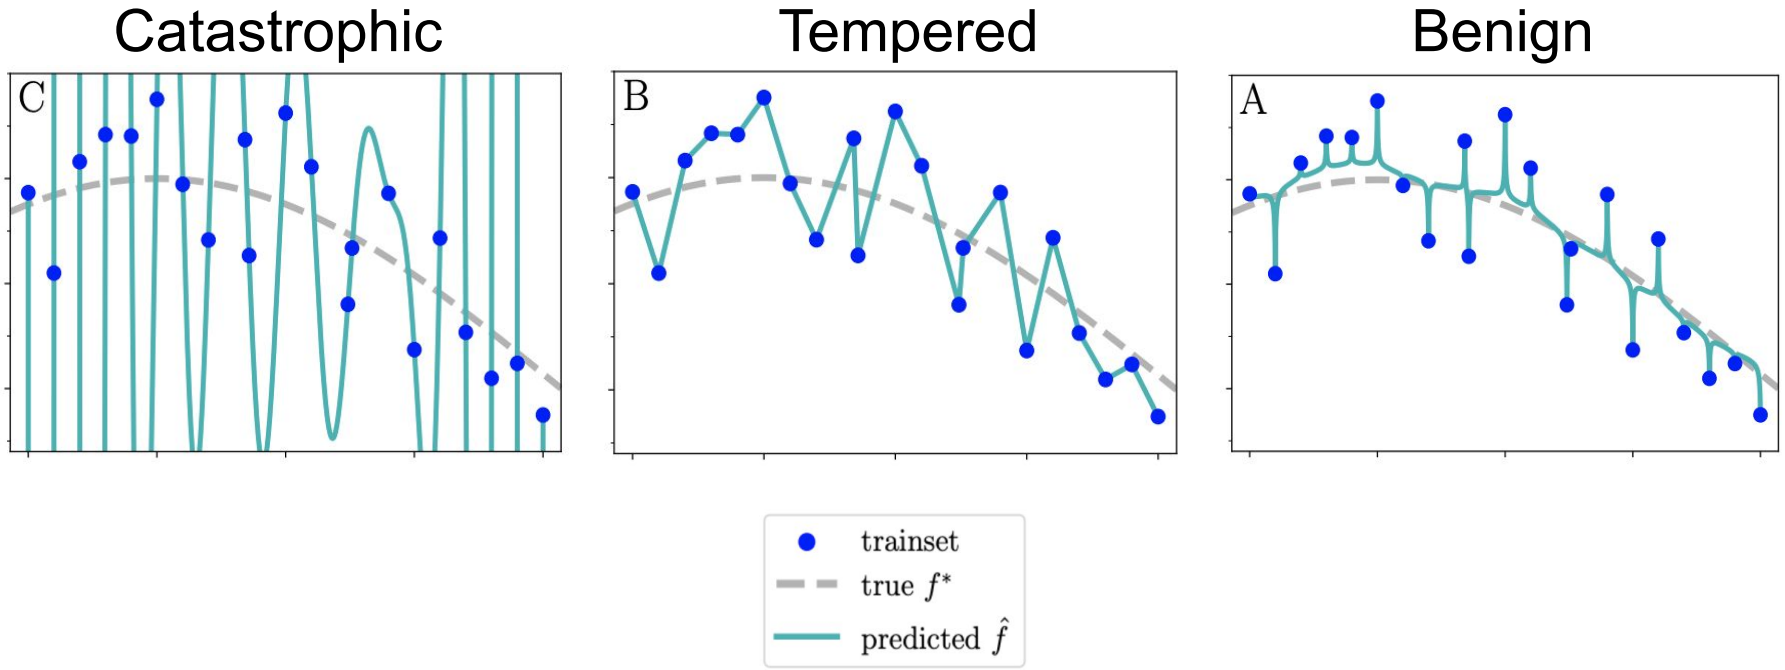

How harmful is overfitting in realistic settings?
Is it actually benign?

Noise profiles

- We propose noise profiles as a tool to measure the harm of overfitting
- Fitting to various label noise helps make the effects of overfitting clear
- This is a binary classification noise profile

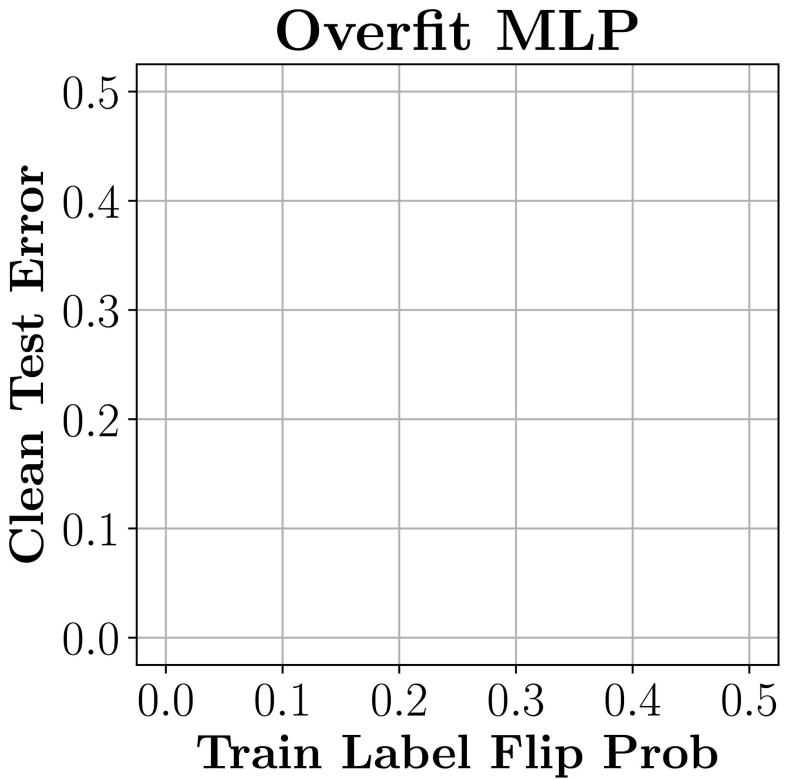

Noise profiles

- A catastrophic method will behave as a random classifier after overfitting noisy labels

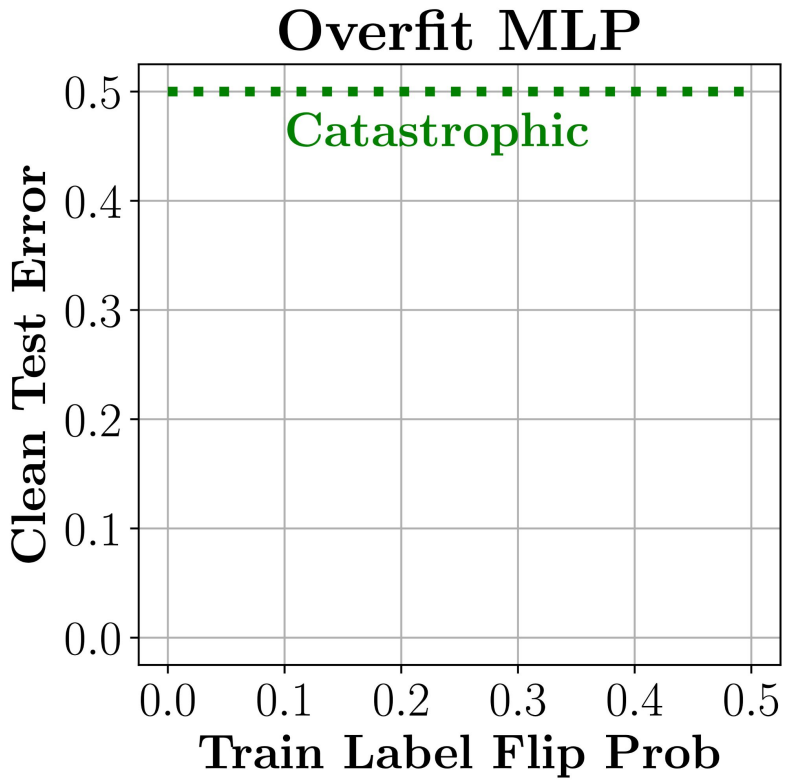

Noise profiles

- A benign method is asymptotically optimal

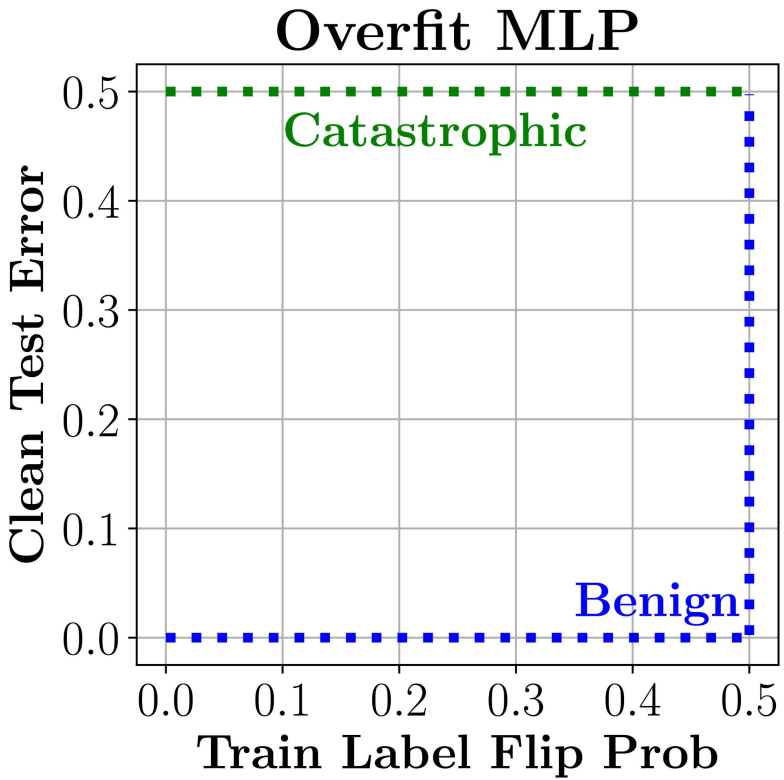

Noise profiles

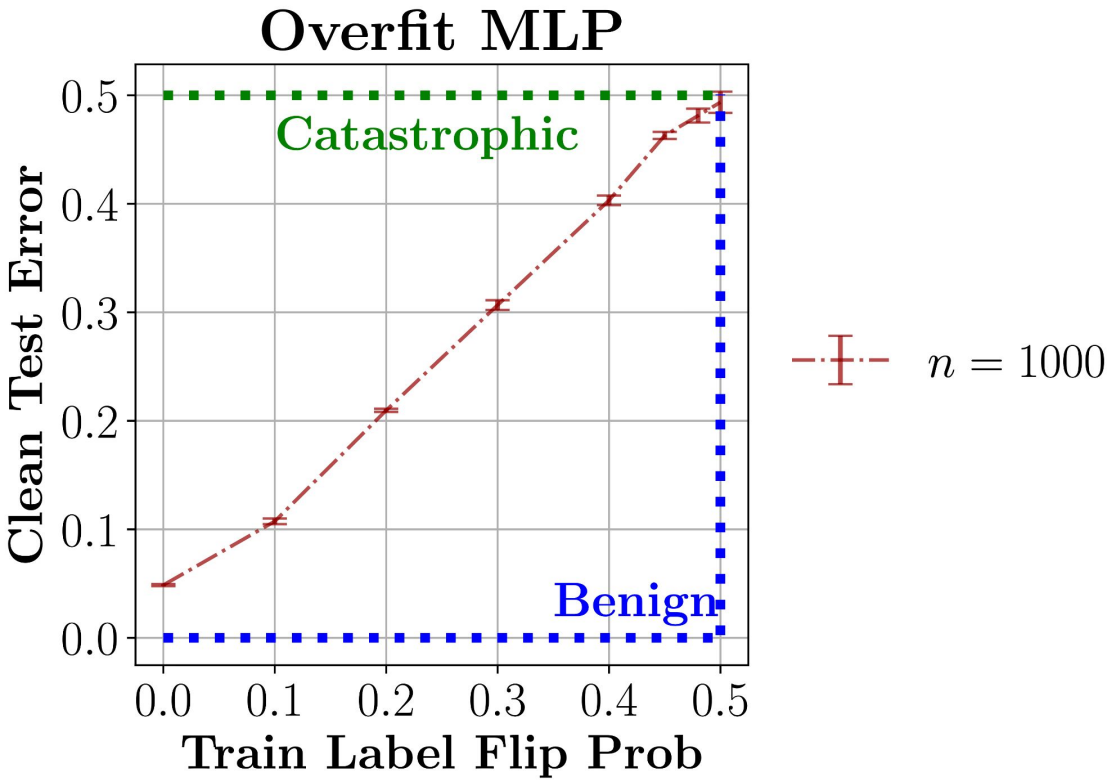

Noise profiles

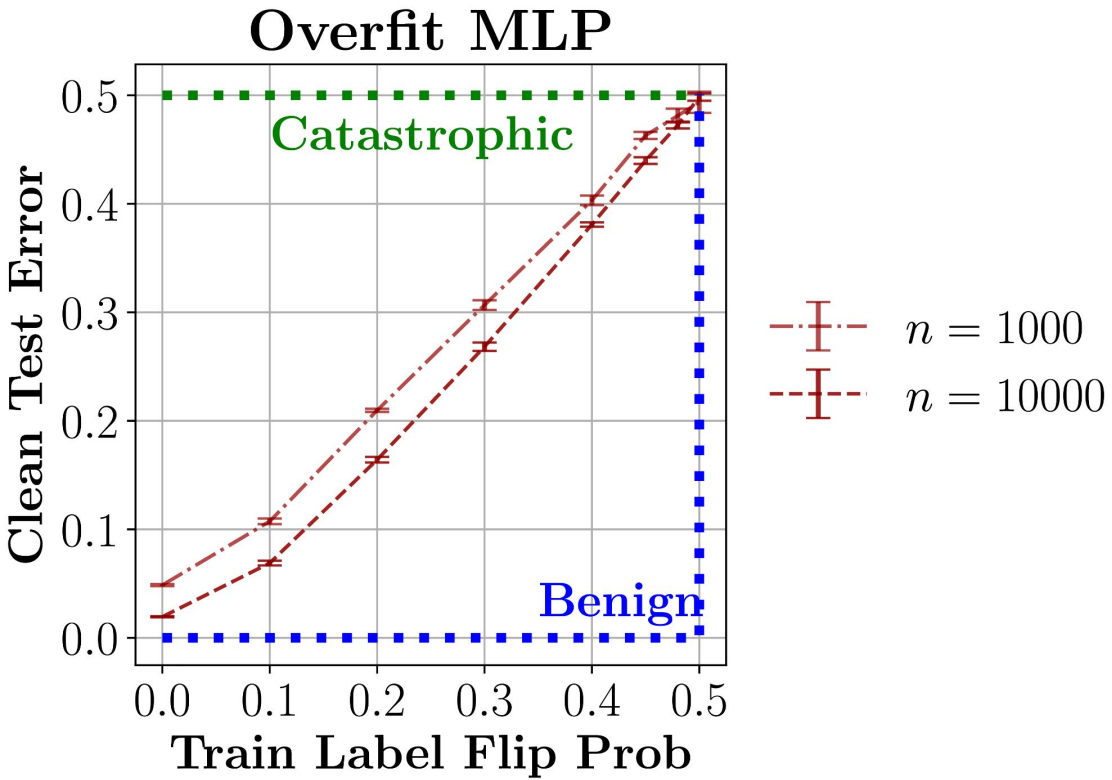

Noise profiles

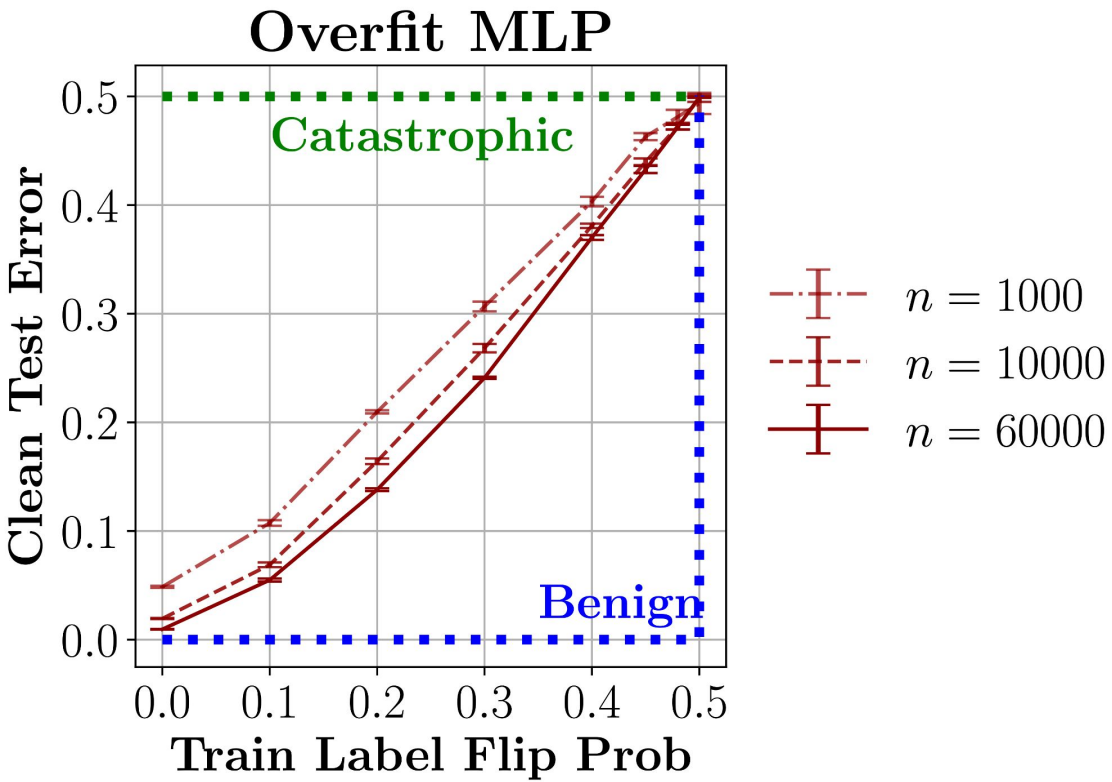

Noise profiles

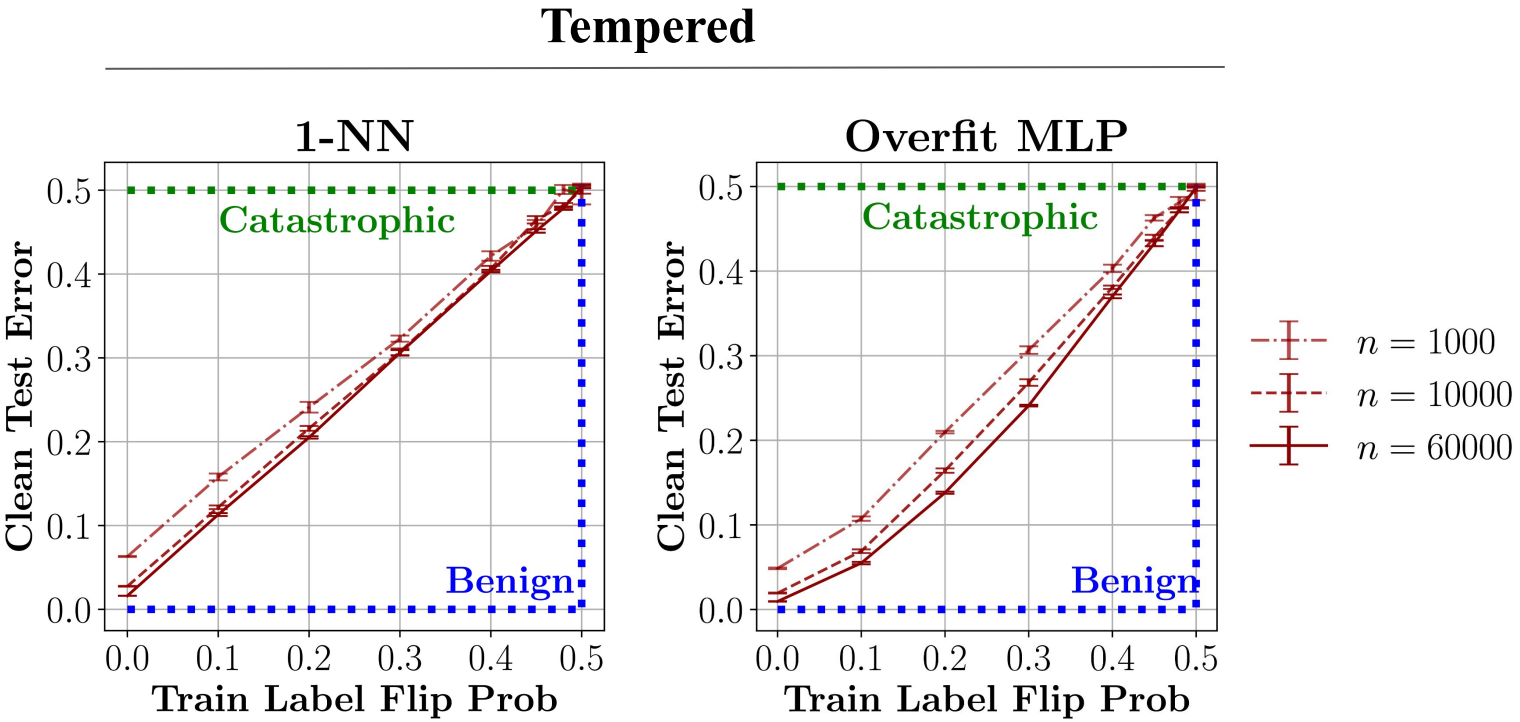

Kernel regression can exhibit all three types of fitting

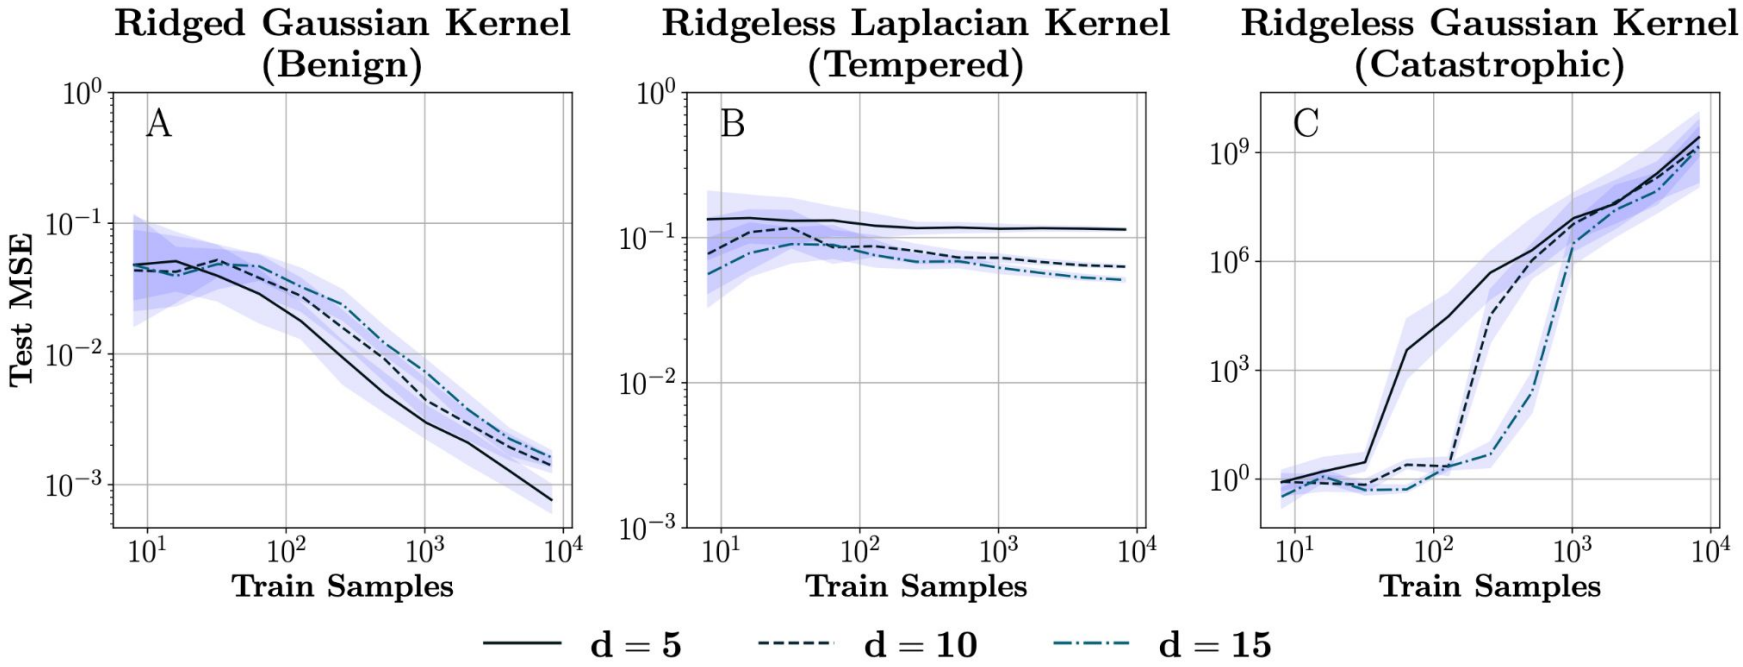

d - data dimension

The Trichotomy Theorem

Spectrum	Limiting risk	
$\delta > 0$ $\delta = 0$ *and* $\lambda_i = i^{-1} \log^{-\alpha} i$ *for some* $\alpha > 1$	$\lim_{n\to\infty} \mathcal{E}_n = \sigma^2$	BENIGN
$\delta = 0$ *and* $\lambda_i = i^{-\alpha}$ *for some* $\alpha > 1$	$\lim_{n\to\infty} \mathcal{E}_n = \alpha\sigma^2$	TEMPERED
$\delta = 0$ *and* $\frac{\lambda_i}{\lambda_{i+1}} \geq \frac{i^{-\log i}}{(i+1)^{-\log(i+1)}}$	$\lim_{n\to\infty} \mathcal{E}_n = \infty$	CATASTROPHIC

λ_i - eigenvalues of data kernel matrix, sorted in descending order

δ - ridge parameter

$\mathcal{E}_n$ - expected risk of model trained on n samples

σ^2 - optimal risk

Summary of contributions

- There are three ways to overfit
- Common interpolating methods fall into the intermediate regime
- Noise profiles are a tool to empirically study overfitting
- For KR, ridge + kernel spectrum determine the regime

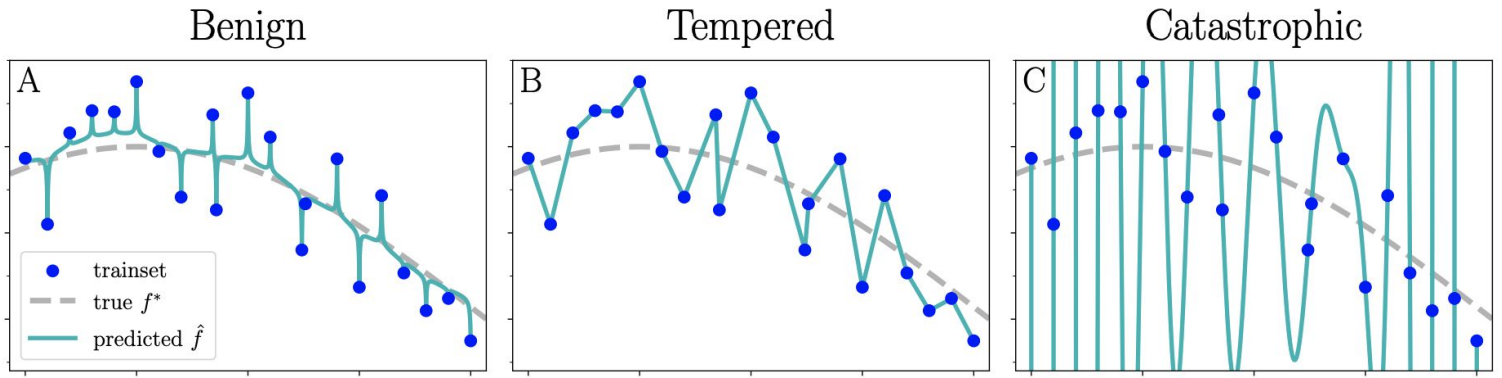

paper at tinyurl.com/TemperedOverfitting