Benefits of Overparameterized Convolutional Residual Networks: Function Approximation under Smoothness Constraint

Hao Liu[⋆], *Minshuo Chen*[†], Siawpeng Er[†], Wenjing Liao[†], Tong Zhang[⋄] and Tuo Zhao[†]

[⋆]Hong Kong Baptist University [†]Georgia Institute of Technology
[⋄]Hong Kong University of Science and Technology

Adversarial Robustness

'Panda'

$+ .007 \times$

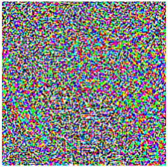

noise

$=$

'Gibbon'

– Credit: Goodfellow et al. ICLR 2015.

Smoothness Ties to Robustness

- A close tie between **smoothness** and adversarial **robustness** [Gu & Rigazio, 2014; Madry et al., 2017; Miyato et al. 2018; Bubeck & Sellke, 2021].

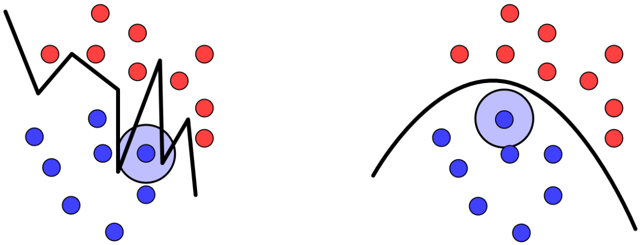

Vulnerable **Robust**

Overparameterization Promotes Smoothness

- Large neural networks favor smoothness and yield good robustness [Madry et al., 2017; Bubeck & Sellke, 2021].

Overparameterization Promotes Smoothness

- Large neural networks favor smoothness and yield good robustness [Madry et al., 2017; Bubeck & Sellke, 2021].

- Theoretical explanation is still missing.

Contributions

- **Approximation theory** for *overparameterized* Convolutional Residual Networks, with smoothness guarantees.
- **Adversarial risk bound** of overparameterized Convolutional Residual Networks.
- Extension to low-dimensional manifold data; **no curse of ambient dimensionality**.

Function Approximation Perspective

- Function value approximation [Yarotsky 2017; Suzuki 2019; Zhou 2020; Peterson & Voigtlaender 2020; Oono & Suzuki 2019; Liu et al. 2021].
- Smoothness of approximation [Hornik et al., 1990; Cardaliaguet & Euvrard, 1992; Gühring et al. 2020; Hon and Yang 2021]

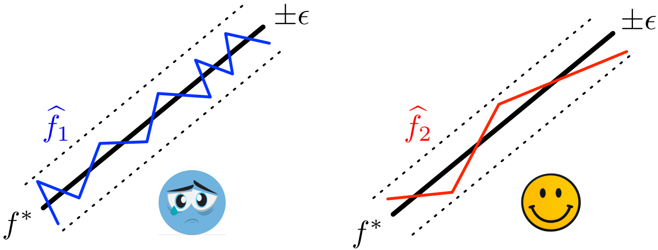

Function Approximation with Smoothness Guarantees

Theorem

Width-$\widetilde{M}$ depth-$\widetilde{J}$ convolutional residual networks can approximate any Sobolev function $f \in W^{\alpha,p}((0,1)^D)$, i.e., there exists $\widetilde{f}$ with

$$\|\widetilde{f} - f\|_\infty \leq C(\widetilde{M}\widetilde{J})^{-\frac{\alpha-1}{D}} \quad \text{and}$$

$$\|\widetilde{f}\|_{\mathrm{Lip}} \leq \|f\|_{\mathrm{Lip}} + C\sqrt{D}(\widetilde{M}\widetilde{J})^{-\frac{\alpha-1}{D}}$$

for some constant C depending on D, α, p.

Function Approximation with Smoothness Guarantees

Theorem

Width-$\widetilde{M}$ depth-$\widetilde{J}$ convolutional residual networks can approximate any Sobolev function $f \in W^{\alpha,p}((0,1)^D)$, i.e., there exists $\widetilde{f}$ with

$$\|\widetilde{f} - f\|_\infty \leq C(\widetilde{M}\widetilde{J})^{-\frac{\alpha-1}{D}} \quad \text{and}$$
$$\|\widetilde{f}\|_{\mathrm{Lip}} \leq \|f\|_{\mathrm{Lip}} + C\sqrt{D}(\widetilde{M}\widetilde{J})^{-\frac{\alpha-1}{D}}$$

for some constant C depending on D, α, p.

- Increasing $\widetilde{M}, \widetilde{J}$ amplifies approximation power.
- To achieve an ϵ L^∞-error, $\widetilde{M}\widetilde{J} = O(\epsilon^{-D/(\alpha-1)})$ (v.s. $O(\epsilon^{-D/\alpha})$).
- **Extension**: d-dimensional manifold, $\|\widetilde{f} - f\|_{W^{1,\infty}} \leq C(\widetilde{M}\widetilde{J})^{-\frac{\alpha-1}{d}}$.

Adversarial Risk Bound

Adversarial risk

$$R(\widetilde{f}, \delta) = \mathbb{E}_{(\mathbf{x}, y)} \left[\sup_{\mathbf{x}' \in B_\delta(\mathbf{x})} \ell\left(\widetilde{f}(\mathbf{x}'), y\right) \right]$$

Corollary

Suppose there exists an optimal classifier $f^* \in W^{\alpha,p}((0,1)^D)$. In the setup of the Theorem above, convolutional residual network gives rise to $\tilde{f}$ with

$$R(\widetilde{f}, \delta) \leq \underbrace{R(\widetilde{f}, 0)}_{\text{Clean Risk}} + \|\ell\|_{\mathrm{Lip}} \underbrace{\left(\|f^*\|_{\mathrm{Lip}} + C\sqrt{D}(\widetilde{M}\widetilde{J})^{\frac{\alpha-1}{D}} \right)}_{\text{Smoothness of } \widetilde{f}} \delta.$$

Thank You