

A Mean Field Analysis Of Deep ResNet:

Towards Provable Optimization Via Overparameterization From Depth

Joint work with Chao Ma, Yulong Lu, Jianfeng Lu and Lexing Ying

Presenter:
Yiping Lu

Contact:

yplu@stanford.edu,

<https://web.stanford.edu/~yplu/>

Stanford
University



Global Convergence Proof Of NN

- Neural Tangent Kernel([Jacot et al.2019]):Linearize the model

$$f_{\text{NN}}(\cdot) = f_{\text{NN}}(\text{init}) + \langle r, f_{\text{NN}}(\text{init}) \rangle, \quad \text{init} >$$

Global Convergence Proof Of NN

- Neural Tangent Kernel([Jacot et al.2019]):Linearize the model

$$f_{\text{NN}}(\cdot) = f_{\text{NN}}(\text{init}) + \langle r, f_{\text{NN}}(\text{init}) \rangle, \quad \text{init} >$$

- **Pro:** can provide proof of convergence for any structure of NN. ([Li et al. 2019])
- **Con:** Feature is lazy learned, *i.e.* not data dependent. ([Chizat and Bach 2019.][Ghorbani et al.2019])

Global Convergence Proof Of NN

- Neural Tangent Kernel([Jacot et al.2019]):**Linearize the model**
$$f_{\text{NN}}(\theta) = f_{\text{NN}}(\theta_{\text{init}}) + \langle r, f_{\text{NN}}(\theta_{\text{init}}) \rangle_{\theta_{\text{init}}}$$
 - **Pro:** can provide proof of convergence for any structure of NN. ([Li et al. 2019])
 - **Con:** Feature is lazy learned, *i.e.* not data dependent. ([Chizat and Bach 2019.][Ghorbani et al.2019])

- Mean Field Regime([Bengio et al.2006][Bach et al.2014][Suzuki et al.2015]): **We consider properties of the loss landscape with respect to the distribution of weights** $L(\theta) = kE_{x \sim p} \|g(\theta, x) - f(x)\|_2^2$, the objective is a convex function

Global Convergence Proof Of NN

- Neural Tangent Kernel([Jacot et al.2019]):**Linearize the model**
 $f_{\text{NN}}(\cdot) = f_{\text{NN}}(\text{init}) + \langle r, f_{\text{NN}}(\text{init}) \rangle, \quad \text{init} \gg$
 - **Pro:** can provide proof of convergence for any structure of NN. ([Li et al. 2019])
 - **Con:** Feature is lazy learned, *i.e.* not data dependent. ([Chizat and Bach 2019.][Ghorbani et al.2019])

- Mean Field Regime([Bengio et al.2006][Bach et al.2014][Suzuki et al.2015]): **We consider properties of the loss landscape with respect to the distribution of weights $L(\cdot) = kE \int g(\cdot, x) f(x)k_2^2$, the objective is a convex function**
 - **Pro:** SGD = Wasserstein Gradient Flow ([Mei et al.2018][Chizat et al.2018][Rotskoff et al.2018])
 - **Con:** Hard to generalize beyond two layer

Neural Ordinary Differential Equation

Figure: ResNet can be seen as the Euler discretization of a time evolving ODE

Neural Ordinary Differential Equation

Figure: ResNet can be seen as the Euler discretization of a time evolving ODE

- Limit of depth! $\rightarrow 1$
- This analogy does not directly provide guarantees of global convergence even in the continuum limit.

Mean Field ResNet

Our Aim: Provide a **new** continuous limit for ResNet with **good limiting landscape**.

Idea: We consider properties of the loss landscape with **respect to the distribution of weights**.

Mean Field ResNet

Our Aim: Provide a **new** continuous limit for ResNet with **good limiting landscape**.

Idea: We consider properties of the loss landscape with **respect to the distribution of weights**.

Here:

- Input data is the initial condition $X(x, 0) = \langle w_2, x \rangle$
- X is the feature, t represents the depth.

- Loss function: $E(\mu) = \mathbb{E}_x \int \frac{1}{2} \langle w_1, X(x, 1) \rangle^2 \mu(x)$



Adjoint Equation

To optimize the Mean Field model, we calculate the gradient $\frac{\partial E}{\partial \theta}$ via the *adjoint sensitivity method*.

Model

The loss function can be written as

$$E(x; \theta) := E_x \left[\frac{1}{2} \|w_1 \cdot X(x, 1) - y(x)\|^2 \right] \quad (1)$$

where X satisfies the equation $X(x, t) = \int_0^t f(X(x, s), \theta) ds$,

Adjoint Equation

To optimize the Mean Field model, we calculate the gradient $\frac{\partial E}{\partial \theta}$ via the *adjoint sensitivity method*.

Model

The loss function can be written as

$$E_{\theta} = E(x; \theta) := E_x \left[\frac{1}{2} \|w_1 \cdot X(x, 1) - y(x)\|^2 \right] \quad (1)$$

where X satisfies the equation $\dot{X}(x, t) = f(X(x, t), \theta)(x, t)$.

Adjoint Equation. The gradient can be represented as a second backwards-in-time augmented ODE.

$$\begin{aligned} \dot{p}(x, t) &= -\nabla_x H(p, x, t) \\ &= -p(x, t) \nabla_x f(X(x, t), \theta)(x, t) \end{aligned}$$

Here the Hamiltonian is defined as $H(p, x, t) = p(x, t) \nabla_x f(x, \theta)(x, t)$.



Adjoint Equation

Theorem

For \mathcal{P}^2 let $E(x, t) = E_x f(X(x, t), y) p(x, t)$, where p is the solution to the backward equation $p(x, t) = p(x, t) - r_x f(X(x, t), y) p(x, t) dt$. Then for every \mathcal{P}^2 , we have

$$E(x + \Delta x, t) = E(x, t) + \frac{\partial E}{\partial x} \Delta x + o(\Delta x)$$

for the convex combination $(1 - \alpha)x + \alpha x'$ with $\alpha \in [0, 1]$.



Adjoint equation is equivalent to the back propagation

Li Q, Chen L, Tai C, et al. Maximum principle based algorithms for deep learning. JMLR 2019

Zhang D, Zhang T, Lu Y, et al. You only propagate once: Painless adversarial training using maximal principle Neurips2019

Deep Residual Network Behaves Like an Ensemble Of Shallow Models

$$x^1 = x^0 + \frac{1}{L} \sum_{l=0}^{L-1} (x^0)^{(l)}$$

Deep Residual Network Behaves Like an Ensemble Of Shallow Models

$$x^1 = x^0 + \frac{1}{L} \sum_0^Z (x^0)^{(0)} d^{(0)}$$

$$\begin{aligned}
 x^2 &= x^0 + \frac{1}{L} \sum_0^Z (x^0)^{(0)} d^{(0)} + \frac{1}{L} \sum_1^Z (x^0 + \frac{1}{L} \sum_0^Z (x^0)^{(0)} d^{(0)})^{(1)} d^{(1)} \\
 &= x^0 + \frac{1}{L} \sum_0^Z (x^0)^{(0)} d^{(0)} + \frac{1}{L} \sum_1^Z (x^0)^{(1)} d^{(1)} + \frac{1}{L^2} \sum_{0,1}^Z (x^0)^{(0)} d^{(0)} d^{(1)} \\
 &\quad + \text{h.o.t.}
 \end{aligned}$$

Deep Residual Network Behaves Like an Ensemble Of Shallow Models

$$x^1 = x^0 + \frac{1}{L} \sum_0^Z (x^0)^{(0)} (d)^{(0)}$$

$$\begin{aligned} x^2 &= x^0 + \frac{1}{L} \sum_0^Z (x^0)^{(0)} (d)^{(0)} + \frac{1}{L} \sum_1^Z (x^0 + \frac{1}{L} \sum_0^Z (x^0)^{(0)} (d)^{(0)})^{(1)} (d)^{(1)} \\ &= x^0 + \frac{1}{L} \sum_0^Z (x^0)^{(0)} (d)^{(0)} + \frac{1}{L} \sum_1^Z (x^0)^{(1)} (d)^{(1)} + \frac{1}{L^2} \sum_{0,1}^Z (x^0)^{(0)} (d)^{(0)} (d)^{(1)} \\ &+ h.o.t. \end{aligned}$$

Iterating this expansion gives rise to

$$x^L = x^0 + \frac{1}{L} \sum_{a=0}^{L-1} (x^0)^{(a)} (d)^{(a)} + \frac{1}{L^2} \sum_{b>a}^Z (x^0)^{(b)} (d)^{(b)} (d)^{(a)} + h.o.t.$$

Veit A, Wilber M J, Belongie S. **Residual networks behave like ensembles of relatively shallow networks.** Advances in neural information processing systems. 2016: 550-558.



Deep Residual Network Behaves Like an Ensemble Of Shallow Models

Difference of back propagation process of two-layer net and ResNet.

Two-layer Network

ResNet

$$\frac{E}{\partial \theta}(\cdot, t) = E_x (f(x, \theta))(X \quad y(x)) \quad \frac{E}{\partial \theta}(\cdot, t) = E_x (f(X(x, t), \theta))p(x, t)$$

We aim to show that **the two gradient are similar.**

Deep Residual Network Behaves Like an Ensemble Of Shallow Models

Difference of back propagation process of two-layer net and ResNet.

Two-layer Network

ResNet

$$\frac{d}{dt} E(x, t) = E_x (f(x, t)) (X - y(x)) \quad \frac{d}{dt} E(x, t) = E_x (f(X(x, t), t)) p(x, t)$$

We aim to show that **the two gradient are similar.**

Lemma

The norm of the solution to the adjoint equation can be bounded by the loss

$$\|k p(x, t)\| \leq e^{(C_1 + C_2 t)} E(x), \quad t \in [0, 1]$$

Local = Global

Theorem

If $E(\cdot) > 0$ for distribution $\mu \in \mathcal{P}^2$ that is supported on one of the nested sets Q_r , we can always construct a descend direction $\mu \in \mathcal{P}^2$, i.e.

$$\inf_{\mu \in \mathcal{P}^2} \frac{E}{\mu}(\cdot) < 0$$

Local = Global

Theorem

If $E(\mu) > 0$ for distribution $\mu \in \mathcal{P}^2$ that is supported on one of the nested sets Q_r , we can always construct a descend direction $\mu \in \mathcal{P}^2$, i.e.

$$\inf_{\mu \in \mathcal{P}^2} \frac{E(\mu)}{E(\mu_0)} < 0$$

Corollary

Consider a stationary solution to the *Wasserstein gradient flow* which is full support (informal), then it's a global minimizer.

Numerical Scheme

We may consider using a parametrization of μ_t with n particles as

$$\mu_t = \sum_{i=1}^n \delta_{x_i(t)}$$

The characteristic function $\delta_{x_i(t)}$ can be viewed as a relaxation of the Dirac delta mass $\delta(x)$.

Given: A collection of residual blocks $(f_i)_{i=1}^n$

while training do

Sort (f_i) based on λ_i to be (f_{i^*}) where $\lambda_{i^*} \leq \lambda_{i+1}$.

Define the ResNet as $X^{k+1} = X^k + (f_{i^*})^{\otimes k}(X^k)$ for $0 \leq k < n$.

Use gradient descent to update both λ_i and f_i .

end while

Numerical Results

	Vanilla	mean-field	Dataset
ResNet20	8.75	8.19	CIFAR10
ResNet32	7.51	7.15	CIFAR10
ResNet44	7.17	6.91	CIFAR10
ResNet56	6.97	6.72	CIFAR10
ResNet110	6.37	6.10	CIFAR10
ResNet164	5.46	5.19	CIFAR10
ResNeXt29(864d)	17.92	17.53	CIFAR100
ResNeXt29(1664d)	17.65	16.81	CIFAR100

Table: Comparison of the stochastic gradient descent and mean-field training (Algorithm 1.) of ResNet On CIFAR Dataset. Results indicate that our method performs the Vanilla SGD consistently.

Take Home Message

- We propose a new continuous limit for deep resnet

$$\dot{X}(x, t) = \int f(X(x, t), \cdot) (\cdot, t) d\cdot,$$

with initial $X(x, 0) = hw_2, xi$

- Local minimizer is global in \mathcal{L}_2 space.
- A potential scheme to approximate.

Take Home Message

- We propose a new continuous limit for deep resnet

$$X(x, t) = \int_0^t f(X(x, t), \cdot)(\cdot, t) d\cdot,$$

with initial $X(x, 0) = hw_2, xi$

- Local minimizer is global in \mathbb{R}^2 space.
- A potential scheme to approximate.

TO DO List.

- Analysis of Wasserstein gradient flow. (Global Existence)
- Refined analysis of numerical scheme
- h.o.t in the expansion from ResNet to ensemble of small networks.



Thanks

Reference: Yiping Lu*, Chao Ma, Yulong Lu, Jianfeng Lu, Lexing Ying. "A Mean-field Analysis of Deep ResNet and Beyond: Towards Provable Optimization Via Overparameterization From Depth" arXiv:2003.05508 ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations. **(Contributed Talk)**

Contact: yplu@stanford.edu

