

Practical - Analysis of new phenomena in machine/deep learning

Pascal M. Esser

Outline

Machine learning and deep learning research

- Empirical studies, providing benchmark and demonstrating pitfalls.
- Rigorously explain why ML / DL works by analysing theoretical models or algorithms.

Focus:

- Insights for new algorithmic development (example: boosting, methods for regularisation).
- Brings concepts from mathematics to ML (example: Random graphs, Geometry).

Machine learning and deep learning research


This Practical:

- Understand recent advances
- Reproduce existing results
- Extend research (empirically)

Topics



Adversarial ML / Robustness

- Performance of NNs significantly affected if data is slightly perturbed.
- Why? How can we robust ML models / guarantee robustness?



Generalisation in neural networks


- Over-parameterised NNs deviate from bias-variance trade-off - NNs may perform best in zero training loss / interpolating regime.
- Currently, this behaviour has been analytically derived in simpler settings.



Over-parameterised NN (infinite width)


Analyse Over-parametrised NNs asymptotically as width goes to infinity

- Under small learning rate, (S)GD training \equiv Neural Tangent Kernel (NTK), a dot product kernel in gradient space of the NN parameters
- Finite width networks can deviate from the kernel regime.



Unsupervised Learning

- What representations do we learn?
- What are (provable) better embeddings?



Let's look at an example

On Exact Computation with an Infinitely Wide Neural Net*

Sanjeev Arora[†]

Simon S. Du[‡]

Wei Hu[§]

Zhiyuan Li[¶]

Ruslan Salakhutdinov^{||}

Ruosong Wang^{**}

Abstract

How well does a classic deep net architecture like AlexNet or VGG19 classify on a standard dataset such as CIFAR-10 when its “width”—namely, number of channels in convolutional layers, and number of nodes in fully-connected internal layers — is allowed to increase to infinity? Such questions have come to the forefront in the quest to theoretically understand deep learning and its mysteries about optimization and generalization. They also connect deep learning to notions such as *Gaussian processes* and *kernels*. A recent paper [Jacot et al., 2018] introduced the *Neural Tangent Kernel (NTK)* which captures the behavior of fully-connected deep nets in the infinite width limit trained by gradient descent; this object was implicit in some other recent papers. An attraction of such ideas is that a pure kernel-based method is used to capture the power of a fully-trained deep net of infinite width.

The current paper gives the first efficient exact algorithm for computing the extension of NTK to convolutional neural nets, which we call *Convolutional NTK (CNTK)*, as well as an efficient GPU implementation of this algorithm. This results in a significant new benchmark for performance of a pure kernel-based method on CIFAR-10, being 10% higher than the methods reported in [Novak et al., 2019], and only 6% lower than the performance of the corresponding finite deep net architecture (once batch normalization etc. are turned off). Theoretically, we also give the first *non-asymptotic* proof showing that a fully-trained sufficiently wide net is indeed equivalent to the kernel regression predictor using NTK.

General Setup

This Paper shows

Our contributions. We give an exact and efficient dynamic programming algorithm to compute CNTKs for ReLU activation (namely, to compute $\ker(\mathbf{x}, \mathbf{x}')$ given \mathbf{x} and \mathbf{x}'). Using this algorithm — as well as implementation tricks for GPUs — we can settle the question of the performance of fully-trained infinitely wide nets with a variety of architectures. For instance, we find that their performance on CIFAR-10 is within 5% of the performance of the same architectures in the finite case (note that the proper comparison in the finite case involves turning off batch norm, data augmentation, etc., in the optimization). In particular, the CNTK corresponding to a 11-layer convolutional net with global average pooling achieves 77% classification accuracy. This is 10% higher than the best reported performance of a Gaussian process with fixed kernel on CIFAR-10 [Novak et al., 2019].⁸

Furthermore, we give a more rigorous, non-asymptotic proof that the NTK captures the behavior of a fully-trained wide neural net under weaker condition than previous proofs. We also experimentally show that the random feature methods for approximating CNTK in earlier work do not compute good approximations, which is clear from their much worse performance on CIFAR.

What exactly does this paper prove / what do we reproduce

1.1 Notation

We use bold-faced letters for vectors, matrices and tensors. For a vector \mathbf{a} , let $[\mathbf{a}]_i$ be its i -th entry; for a matrix \mathbf{A} , let $[\mathbf{A}]_{i,j}$ be its (i, j) -th entry; for a 4th-order tensor \mathbf{T} , let $[\mathbf{T}]_{ij,i'j'}$ be its (i, j, i', j') -th entry. Let \mathbf{I} be the identity matrix, and $[n] = \{1, 2, \dots, n\}$. Let \mathbf{e}_i be an indicator vector with i -th entry being 1 and other entries being 0, and let $\mathbf{1}$ denote the all-one vector. We use \odot to denote the entry-wise product and \otimes to denote the tensor product. We use $\langle \cdot, \cdot \rangle$ to denote the standard inner product. We use $\text{diag}(\cdot)$ to transform a vector to a diagonal matrix. We use $\sigma(\cdot)$ to denote the activation function, such as the rectified linear unit (ReLU) function: $\sigma(z) = \max\{z, 0\}$, and $\dot{\sigma}(\cdot)$ to denote the derivative of $\sigma(\cdot)$. Denote by $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ the Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$.

2 Related Work

What is new in this paper?

From a Gaussian process (GP) viewpoint, the correspondence between infinite neural networks and kernel machines was first noted by Neal [1996]. Follow-up work extended this corre-

3 Neural Tangent Kernel

In this section we describe fully-connected deep neural net architecture and its infinite width limit, and how training it with respect to the ℓ_2 loss gives rise to a kernel regression problem involving the neural tangent kernel (NTK). We denote by $f(\boldsymbol{\theta}, \mathbf{x}) \in \mathbb{R}$ the output of a neural network where $\boldsymbol{\theta} \in \mathbb{R}^N$ is all the parameters in the network and $\mathbf{x} \in \mathbb{R}^d$ is the input.⁹ Given a training dataset $\{(\mathbf{x}_i, y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \mathbb{R}$, consider training the neural network by minimizing the squared loss over training data: $\ell(\boldsymbol{\theta}) = \frac{1}{2} \sum_{i=1}^n (f(\boldsymbol{\theta}, \mathbf{x}_i) - y_i)^2$. The proof of the following lemma uses simple differentiation and appears in Section C.

General
Setup

convenience. We define an L -hidden-layer fully-connected neural network recursively:

$$\mathbf{f}^{(h)}(\mathbf{x}) = \mathbf{W}^{(h)} \mathbf{g}^{(h-1)}(\mathbf{x}) \in \mathbb{R}^{d_h}, \quad \mathbf{g}^{(h)}(\mathbf{x}) = \sqrt{\frac{c_\sigma}{d_h}} \sigma \left(\mathbf{f}^{(h)}(\mathbf{x}) \right) \in \mathbb{R}^{d_h}, \quad h = 1, 2, \dots, L, \quad (6)$$

where $\mathbf{W}^{(h)} \in \mathbb{R}^{d_h \times d_{h-1}}$ is the weight matrix in the h -th layer ($h \in [L]$), $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ is a coordinate-wise activation function, and $c_\sigma = \left(\mathbb{E}_{z \sim \mathcal{N}(0,1)} [\sigma(z)^2] \right)^{-1}$. The last layer of the neural network is

$$\begin{aligned} f(\boldsymbol{\theta}, \mathbf{x}) &= f^{(L+1)}(\mathbf{x}) = \mathbf{W}^{(L+1)} \cdot \mathbf{g}^{(L)}(\mathbf{x}) \\ &= \mathbf{W}^{(L+1)} \cdot \sqrt{\frac{c_\sigma}{d_L}} \sigma \left(\mathbf{W}^{(L)} \cdot \sqrt{\frac{c_\sigma}{d_{L-1}}} \sigma \left(\mathbf{W}^{(L-1)} \dots \sqrt{\frac{c_\sigma}{d_1}} \sigma \left(\mathbf{W}^{(1)} \mathbf{x} \right) \right) \right), \end{aligned}$$

where $\mathbf{W}^{(L+1)} \in \mathbb{R}^{1 \times d_L}$ is the weights in the final layer, and $\boldsymbol{\theta} = (\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(L+1)})$ represents all the parameters in the network.

We initialize all the weights to be i.i.d. $\mathcal{N}(0, 1)$ random variables, and consider the limit of large hidden widths: $d_1, d_2, \dots, d_L \rightarrow \infty$. The scaling factor $\sqrt{c_\sigma/d_h}$ in Equation (6) ensures that the norm of $\mathbf{g}^{(h)}(\mathbf{x})$ for each $h \in [L]$ is approximately preserved at initialization (see [Du et al., 2018b]). In particular, for ReLU activation, we have $\mathbb{E} [\|\mathbf{g}^{(h)}(\mathbf{x})\|^2] = \|\mathbf{x}\|^2$ ($\forall h \in [L]$).

Recall from [Lee et al., 2018] that in the infinite width limit, the pre-activations $\mathbf{f}^{(h)}(\mathbf{x})$ at every hidden layer $h \in [L]$ has all its coordinates tending to i.i.d. centered Gaussian processes of covariance $\Sigma^{(h-1)}: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ defined recursively as: for $h \in [L]$,

$$\begin{aligned} \Sigma^{(0)}(\mathbf{x}, \mathbf{x}') &= \mathbf{x}^\top \mathbf{x}', \\ \boldsymbol{\Lambda}^{(h)}(\mathbf{x}, \mathbf{x}') &= \begin{pmatrix} \Sigma^{(h-1)}(\mathbf{x}, \mathbf{x}) & \Sigma^{(h-1)}(\mathbf{x}, \mathbf{x}') \\ \Sigma^{(h-1)}(\mathbf{x}', \mathbf{x}) & \Sigma^{(h-1)}(\mathbf{x}', \mathbf{x}') \end{pmatrix} \in \mathbb{R}^{2 \times 2}, \\ \Sigma^{(h)}(\mathbf{x}, \mathbf{x}') &= c_\sigma \mathbb{E}_{(u,v) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Lambda}^{(h)})} [\sigma(u) \sigma(v)]. \end{aligned} \quad (7)$$

To give the formula of NTK, we also need to define a derivative covariance:

$$\dot{\Sigma}^{(h)}(\mathbf{x}, \mathbf{x}') = c_\sigma \mathbb{E}_{(u,v) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Lambda}^{(h)})} [\dot{\sigma}(u) \dot{\sigma}(v)]. \quad (8)$$

The final NTK expression for the fully-connected neural network is

$$\Theta^{(L)}(\mathbf{x}, \mathbf{x}') = \sum_{h=1}^{L+1} \left(\Sigma^{(h-1)}(\mathbf{x}, \mathbf{x}') \cdot \prod_{h'=h}^{L+1} \dot{\Sigma}^{(h')}(\mathbf{x}, \mathbf{x}') \right), \quad (9)$$

Exact definition of the model.

Important for reproducing results

Corresponding NTK

this formula. Rigorously, for ReLU activation, we have the following theorem that gives a concrete bound on the hidden widths that is sufficient for convergence to the NTK at initialization:

Theorem 3.1 (Convergence to the NTK at initialization). Fix $\epsilon > 0$ and $\delta \in (0, 1)$. Suppose $\sigma(z) = \max(0, z)$ and $\min_{h \in [L]} d_h \geq \Omega(\frac{L^6}{\epsilon^4} \log(L/\delta))$. Then for any inputs $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^{d_0}$ such that $\|\mathbf{x}\| \leq 1, \|\mathbf{x}'\| \leq 1$, with probability at least $1 - \delta$ we have:

$$\left| \left\langle \frac{\partial f(\boldsymbol{\theta}, \mathbf{x})}{\partial \boldsymbol{\theta}}, \frac{\partial f(\boldsymbol{\theta}, \mathbf{x}')}{\partial \boldsymbol{\theta}} \right\rangle - \Theta^{(L)}(\mathbf{x}, \mathbf{x}') \right| \leq (L + 1)\epsilon.$$

Theoretical
Result 1

Equivalence between wide neural net and kernel regression with NTK. Built on Theorem 3.1, we can further incorporate the training process and show the equivalence between a fully-trained sufficiently wide neural net and the kernel regression solution using the NTK, as described in Lemma 3.1 and the discussion after it.

Recall that the training data are $\{(\mathbf{x}_i, y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \mathbb{R}$, and $\mathbf{H}^* \in \mathbb{R}^{n \times n}$ is the NTK evaluated on these training data, i.e., $[\mathbf{H}^*]_{i,j} = \Theta^{(L)}(\mathbf{x}_i, \mathbf{x}_j)$. Denote $\lambda_0 = \lambda_{\min}(\mathbf{H}^*)$. For a testing point $\mathbf{x}_{te} \in \mathbb{R}^d$, we let $\ker_{ntk}(\mathbf{x}_{te}, \mathbf{X}) \in \mathbb{R}^n$ be the kernel evaluated between the testing point and n training points, i.e., $[\ker_{ntk}(\mathbf{x}_{te}, \mathbf{X})]_i = \Theta^{(L)}(\mathbf{x}_{te}, \mathbf{x}_i)$. The prediction of kernel regression using NTK on this testing point is $f_{ntk}(\mathbf{x}_{te}) = (\ker_{ntk}(\mathbf{x}_{te}, \mathbf{X}))^\top (\mathbf{H}^*)^{-1} \mathbf{y}$.

Since the above solution corresponds to the linear dynamics in Equation (4) with zero initialization, in order to establish equivalence between neural network and kernel regression, we would like the initial output of the neural network to be small. Therefore, we apply a small multiplier $\kappa > 0$, and let the final output of the neural network be $f_{nn}(\boldsymbol{\theta}, \mathbf{x}) = \kappa f(\boldsymbol{\theta}, \mathbf{x})$. We let $f_{nn}(\mathbf{x}_{te}) = \lim_{t \rightarrow \infty} f_{nn}(\boldsymbol{\theta}(t), \mathbf{x}_{te})$ be the prediction of the neural network at the end of training.

The following theorem establishes the equivalence between the fully-trained wide neural network f_{nn} and the kernel regression predictor f_{ntk} using the NTK.

Theorem 3.2 (Equivalence between trained net and kernel regression). *Suppose $\sigma(z) = \max(0, z)$, $1/\kappa = \text{poly}(1/\epsilon, \log(n/\delta))$ and $d_1 = d_2 = \dots = d_L = m$ with $m \geq \text{poly}(1/\kappa, L, 1/\lambda_0, n, \log(1/\delta))$. Then for any $\mathbf{x}_{te} \in \mathbb{R}^d$ with $\|\mathbf{x}_{te}\| = 1$, with probability at least $1 - \delta$ over the random initialization, we have*

$$|f_{nn}(\mathbf{x}_{te}) - f_{ntk}(\mathbf{x}_{te})| \leq \epsilon.$$

Theoretical
Result 2

Note that $\Sigma(\mathbf{x}, \mathbf{x}')$ and $\dot{\Sigma}(\mathbf{x}, \mathbf{x}')$ share similar structures as their NTK counterparts in Equations (7) and (8). The only difference is that we have one more step, taking the trace over patches. This step represents the convolution operation in the corresponding CNN. Next, we can use a recursion to compute the CNTK:

1. First, we define $\Theta^{(0)}(\mathbf{x}, \mathbf{x}') = \Sigma^{(0)}(\mathbf{x}, \mathbf{x}')$.
2. For $h = 1, \dots, L - 1$ and $(i, j, i', j') \in [P] \times [Q] \times [P] \times [Q]$, we define

$$\left[\Theta^{(h)}(\mathbf{x}, \mathbf{x}') \right]_{ij, i'j'} = \text{tr} \left(\left[\dot{\mathbf{K}}^{(h)}(\mathbf{x}, \mathbf{x}') \odot \Theta^{(h-1)}(\mathbf{x}, \mathbf{x}') + \mathbf{K}^{(h)}(\mathbf{x}, \mathbf{x}') \right]_{D_{ij, i'j'}} \right).$$

3. For $h = L$, we define $\Theta^{(L)}(\mathbf{x}, \mathbf{x}') = \dot{\mathbf{K}}^{(L)}(\mathbf{x}, \mathbf{x}') \odot \Theta^{(L-1)}(\mathbf{x}, \mathbf{x}') + \mathbf{K}^{(L)}(\mathbf{x}, \mathbf{x}')$.
4. The final CNTK value is defined as $\text{tr}(\Theta^{(L)}(\mathbf{x}, \mathbf{x}'))$.

Main Algorithm to compute CNTK

Summery of Empirical Results

Depth	CNN-V	CNTK-V	CNTK-V-2K	CNN-GAP	CNTK-GAP	CNTK-GAP-2K
3	59.97%	64.47%	40.94%	63.81%	70.47%	49.71%
4	60.20%	65.52%	42.54%	80.93%	75.93%	51.06%
6	64.11%	66.03%	43.43%	83.75%	76.73%	51.73%
11	69.48%	65.90%	43.42%	82.92%	77.43%	51.92%
21	75.57%	64.09%	42.53%	83.30%	77.08%	52.22%

Table 1: Classification accuracies of CNNs and CNTKs on the CIFAR-10 dataset. CNN-V represents vanilla CNN and CNTK-V represents the kernel corresponding to CNN-V. CNN-GAP represents CNN with GAP and CNTK-GAP represents the kernel corresponding to CNN-GAP. CNTK-V-2K and CNTK-GAP-2K represent training CNTKs with only 2,000 training data.

5 Experiments

We evaluate the performances of CNNs and their corresponding CNTKs on the CIFAR-10 dataset. The implementation details are in Section A. We also compare the performances between CNTKs and their corresponding random feat Due to space limit, we defer these results on random features to Section B.

Results. We test two types of architectures, vanilla CNN and CNN with global average pooling (GAP), as described in Sections 4 and H. We also test CNTKs with only 2,000 training data to see whether their performances are consistent with CNTKs and CNNs using the full training set. The results are summarized in Table 1. Notice that in Table 1, depth is the total number of layers (including both convolution layers and fully-connected layers).

Several comments are in sequel. First, CNTKs are very powerful kernels. The best kernel, 11-layer CNTK with GAP, achieves 77.43% classification accuracy on CIFAR-10. This results in a significant new benchmark for performance of a pure kernel-based method on CIFAR-10, being 10% higher than methods reported in [Novak et al., 2019].

Second, we find that for both CNN and CNTK, depth can affect the classification accuracy. This observation demonstrates that depth not only matters in deep neural networks but can also affect the performance of CNTKs.

Third, the global average pooling operation can significantly increase the classification accuracy by 8% - 10% for both CNN and CNTK. Based on this finding, we expect that many techniques that

Experiment
Setup

Analysis of
Results

Possible Extensions / Further Experiments

- How does data pre-processing influence the performance of CNN and CNTK?
- Analyze the influence of depth for CNN and CNTK
- CNTK with vector output and global average pooling: What helps performance improve with depth?

Practical Structure

Structure

Groups of 2 students - 2 research papers per group

1. understand the main theoretical ideas of the paper and reproduce the empirical findings.
2. extend on the empirical observations with further experiments.

Timeline

- **First or second week of semester:** Introduction Meeting
- **Mid June:** Reproducibility report submission
- **End of semester:** Empirical extensions + functional code submission
- **End of semester:** Final presentations

Weekly:

- ~10 min update presentation from every student
- Office Hours for further questions

Grading

- Report on reproducibility (40%)
- Report on extensions (20%)
- Final group presentation (40%).

Prerequisite

- Machine learning (IN2064)
- Introduction to deep learning (IN2346)
- Statistical foundations of learning (IN2378) - optional

Survey: <https://forms.gle/so6b5GL9PCBofiXb7>

(Link also on website of
Theoretical Foundations of Artificial Intelligence)

THIS DOES NOT REPLACE THE MATCHING SYSTEM

