Fast Finite Width Neural Tangent Kernel

Roman Novak [1] **Jascha Sohl-Dickstein** [1] **Samuel S. Schoenholz** [1]

Abstract

The Neural Tangent Kernel (NTK), defined as $\Theta_\theta^f(x_1, x_2) = \left[\partial f(\theta, x_1)/\partial\theta\right] \left[\partial f(\theta, x_2)/\partial\theta\right]^T$ where $\left[\partial f(\theta, \cdot)/\partial\theta\right]$ is a neural network (NN) Jacobian, has emerged as a central object of study in deep learning. In the infinite width limit, the NTK can sometimes be computed analytically and is useful for understanding training and generalization of NN architectures. At finite widths, the NTK is also used to better initialize NNs, compare the conditioning across models, perform architecture search, and do meta-learning. Unfortunately, the finite width NTK is notoriously expensive to compute, which severely limits its practical utility. We perform the first in-depth analysis of the compute and memory requirements for NTK computation in finite width networks. Leveraging the structure of neural networks, we further propose two novel algorithms that change the *exponent* of the compute and memory requirements of the finite width NTK, dramatically improving efficiency. Our algorithms can be applied in a black box fashion to *any* differentiable function, including those implementing neural networks. We open-source our implementations within the Neural Tangents package (Novak et al., 2020) at github.com/google/neural-tangents.

1. Introduction

The past few years have seen significant progress towards a theoretical foundation for deep learning. Much of this work has focused on understanding the properties of random functions in high dimensions. One significant line of work (Neal, 1994; Lee et al., 2018; Matthews et al., 2018; Borovykh, 2018; Garriga-Alonso et al., 2019; Novak et al., 2019; Yang, 2019; Hron et al., 2020b;a; Hu et al., 2020) established that in the limit of infinite width, randomly initialized Neural Networks (NNs) are Gaussian Processes (called NNGPs). Building on this development, Jacot et al. (2018) showed that in function space the dynamics under gradient descent could be computed analytically using the so-called Neural Tangent Kernel (NTK) and Lee et al. (2019) showed that wide neural networks reduce to their linearizations in weight space throughout training. A related set of results (Belkin et al., 2019; Spigler et al., 2019) showed that the ubiquitous bias-variance decomposition breaks down as high-dimensional models enter the so-called interpolating regime. Together these results describe learning in the infinite width limit and help explain the impressive generalization capabilities of NNs.

Insights from the wide network limit have had significant practical impact. The conditioning of the NTK has been shown to significantly impact trainability and generalization in NNs (Schoenholz et al., 2017; Xiao et al., 2018; 2020). This notion inspired initialization schemes like Fixup (Zhang et al., 2019), MetaInit (Dauphin & Schoenholz, 2019), and Normalizer Free networks (Brock et al., 2021a;b), and has enabled efficient neural architecture search (Park et al., 2020; Chen et al., 2021b). The NTK has additionally given insight into a wide range of phenomena such as: behavior of Generative Adversarial Networks (Franceschi et al., 2021), neural scaling laws (Bahri et al., 2021), and neural irradiance fields (Tancik et al., 2020). Kernel regression using the NTK has further enabled strong performance on small datasets (Arora et al., 2020), and applications such as approximate inference (Khan et al., 2019), dataset distillation (Nguyen et al., 2020; 2021), and uncertainty prediction (He et al., 2020; Adlam et al., 2020).

Despite the significant promise of theory based on the NTK, computing the NTK in practice is challenging. In the infinite width limit, the NTK can sometimes be computed analytically. However, the infinite-width kernel remains intractable for many architectures and finite width corrections are often important to describe actual NNs used in practice (see §I for detailed discussion). The NTK matrix can be computed for finite width networks as the outer product of Jacobians using forward or reverse mode automatic differentiation (AD),

$$\underbrace{\Theta_\theta^f(x_1, x_2)}_{\mathbf{O}\times\mathbf{O}} := \underbrace{\left[\partial f(\theta, x_1)/\partial\theta\right]}_{\mathbf{O}\times\mathbf{P}} \underbrace{\left[\partial f(\theta, x_2)/\partial\theta\right]^T}_{\mathbf{P}\times\mathbf{O}}, \quad (1)$$

[1]Google Brain, Mountain View, California, United States. Correspondence to: Roman Novak <romann@google.com>.

Proceedings of the 39^{th} International Conference on Machine Learning, Baltimore, Maryland, USA, PMLR 162, 2022. Copyright 2022 by the author(s).

where f is the forward pass NN function producing outputs in $\mathbb{R}^{\mathbf{O}}$, $\theta \in \mathbb{R}^{\mathbf{P}}$ are all trainable parameters, and x_1 and x_2 are two inputs to the network. If inputs are batches of sizes $\mathbf{N}_1$ and $\mathbf{N}_2$, the NTK is an $\mathbf{N}_1\mathbf{O} \times \mathbf{N}_2\mathbf{O}$ matrix.

Unfortunately, evaluating Eq. (1) is often infeasible due to time and memory requirements. For modern machine learning tasks $\mathbf{O}$ is often greater (sometimes much greater) than 1000 (e.g. for ImageNet (Deng et al., 2009)), while even modestly sized models feature tens of millions of parameters, or $\mathbf{P} \sim 10^7$. This makes both storing ($[\mathbf{N}_1 + \mathbf{N}_2]\,\mathbf{OP}$ memory) and contracting ($[\mathbf{N}_1\mathbf{N}_2]\,\mathbf{O}^2\mathbf{P}$ time) the Jacobians in Eq. (1) very costly. The theoretical importance of the NTK together with its prohibitive computational costs implies that performance improvements will unlock impactful novel research.

We perform the first in-depth analysis of the compute and memory requirements for the NTK as in Eq. (1). Noting that forward and reverse mode AD are two extremes of a wide range of AD strategies (Naumann, 2004; 2008), we explore other methods for computing the NTK leveraging the structure of NNs used in practice. We propose two novel methods for computing the NTK that exploit different orderings of the computation. We describe the compute and memory requirements of our techniques in fully-connected (FCN) and convolutional (CNN) settings, and show that one is asymptotically more efficient in both settings. We compute the NTK over a wide range of NN architectures and demonstrate that these improvements are robust in practice. We open-source our implementations as general-purpose JAX[1] (Bradbury et al., 2018) function transformations.

2. Related Work

The finite width NTK (denoted simply NTK throughout this work[2]) has been used extensively in many recent works, but to our knowledge implementation details and compute costs were rarely made public. Below we draw comparison to some of these works, but we stress that it only serves as a sanity check to make sure our contribution is valuable relative to the scale of problems that have been attempted. None of these works had efficient NTK computation as their central goal.

In order to compare performance of models based on the NTK and the infinite width NTK, Arora et al. (2019a, Table 2) compute the NTK of up to 20-layer, 128-channel CNN in a binary CIFAR-2 classification setting. In an equivalent setting with the same hardware (NVIDIA V100), we are able to compute the NTK of a 2048-channel CNN, i.e. a network with at least 256 times more parameters.

To demonstrate the stability of the NTK during training for wide networks, Lee et al. (2019, Figure S6) compute the NTK of up to 3-layer 2^{12}-wide or 1-layer 2^{14}-wide FCNs. In the same setting with the same hardware (NVIDIA V100), we can reach widths of at least 2^{14} and 2^{18} respectively, i.e. handle networks with 4 to 16 times more parameters.

To investigate convergence of a WideResNet WRN-28-k (Zagoruyko & Komodakis, 2016) to its infinite width limit, Novak et al. (2020, Figure 2) evaluate the NTK of this model with widening factor k up to 32. In matching setting and hardware, we are able to reach a widening factor of at least 64, i.e. work with models at least 4 times larger.

To meta-learn NN parameters for transfer learning in a MAML-like (Finn et al., 2017) setting, Zhou et al. (2021, Table 7) replace the inner training loop with NTK-based inference. They use up to 5-layer, 200-channel CNNs on MiniImageNet (Oreshkin et al., 2018) with scalar outputs and batch size 25. In same setting we achieve at least 512 channels, i.e. support models at least 6 times larger.

Park et al. (2020, §4.1) use the NTK to predict the generalization performance of architectures in the context of Neural Architecture Search (Zoph & Le, 2016, NAS); however, the authors comment on its high computational burden and ultimately use a different proxy objective. In another NAS setting, Chen et al. (2021a, §3.1.1) use the condition number of NTK to predict a model's trainability. Chen et al. (2021b, Table 1) use the NTK to evaluate the trainability of several ImageNet models such as ResNet 50/152 (He et al., 2016), Vision Transformer (Dosovitskiy et al., 2021) and MLP-Mixer (Tolstikhin et al., 2021). However, due to the prohibitive computational cost, in all of these cases the authors only evaluate a pseudo-NTK, i.e. the NTK of a scalar-valued function,[3] which impacts the quality of the respective trainability/generalization proxy.

By contrast, in this work we can compute the full 1000×1000 (1000 classes) NTK for the same models, i.e. perform a task 1000 times more costly.

Finally, we remark that in all of the above settings, scaling up by increasing width or by working with the true NTK (vs the pseudo-NTK) should lead to improved downstream task performance due to a better infinite width/linearization approximation or a higher-quality trainability/generalization proxy respectively, which makes our work especially relevant to modern research.

[1] Our algorithms are framework-agnostic, but implementation in JAX is easier, as described in §M. We also provide instructions for implementation in other frameworks like Tensorflow (Abadi et al., 2016) and PyTorch (Paszke et al., 2019) in §M.

[2] See §I for a comparison between the finite and infinite width settings.

[3] Precisely, computing the Jacobian only for a single logit or the sum of all 1000 class logits. The result is not the full NTK, but rather a single diagonal block or the sum of its 1000 diagonal blocks (the finite width NTK is a dense matrix, not block-diagonal).

3. Algorithms for Efficient NTK Computation

We now describe our algorithms for fast NTK computation.

In §3.1 we cover preliminaries. We begin by introducing notation used throughout the paper (§3.1.1). We then (§3.1.2) describe primitive building blocks of AD including the Jacobian-vector products (JVP) and vector-Jacobian products (VJP) that correspond to forward and reverse mode AD respectively before discussing the Jacobian (§3.1.3).

In §3.2 we apply the above tools to describe the computational complexity of the baseline approach to computing the NTK that is used in most (likely all) prior works.

In §3.3 and §3.4 we present our two algorithms that each enable accelerating the computation by orders of magnitude in different ways.

3.1. Preliminaries

3.1.1. Notation

Consider a NN $f(\theta, x) \in \mathbb{R}^{\mathbf{O}}$ with $\mathbf{O}$ outputs (e.g. class logits) per input x and a total number $\mathbf{P}$ of trainable parameters $\theta = \operatorname{vec}\left[\theta^0, \dots, \theta^{\mathbf{L}}\right]$, with each θ^l of size $\mathbf{P}^l$, $\mathbf{P} = \sum_{l=0}^{\mathbf{L}} \mathbf{P}^l$. Also assume the network has $\mathbf{K}$ intermediate primitive outputs y^k of size $\mathbf{Y}^k$ each (for example, activations or pre-activations), and let $\mathbf{Y} = \sum_{k=1}^{\mathbf{K}} \mathbf{Y}^k$ be the total size of the outputs (see Fig. 5 and Fig. 6). The NTK is

$$\underbrace{\Theta_\theta^f}_{\mathbf{O}\times\mathbf{O}} := \underbrace{\left[\partial f\left(\theta, x_1\right)/\partial\theta\right]}_{\mathbf{O}\times\mathbf{P}} \underbrace{\left[\partial f\left(\theta, x_2\right)/\partial\theta\right]^T}_{\mathbf{P}\times\mathbf{O}} = \tag{2}$$

$$\sum_{l=0}^{\mathbf{L}} \underbrace{\left[\partial f\left(\theta, x_1\right)/\partial\theta^l\right]}_{\mathbf{O}\times\mathbf{P}^l} \underbrace{\left[\partial f\left(\theta, x_2\right)/\partial\theta^l\right]^T}_{\mathbf{P}^l\times\mathbf{O}}. \tag{3}$$

We denote **FP** to be the (time or memory, depending on context) cost of a single forward pass $f(\theta, x)$. For memory, we exclude the cost of storing all $\mathbf{P}$ weights, but rather define it to be the cost of evaluating f one primitive y^k at a time. This gives a memory cost of at most $\mathcal{O}\left(\max_k \mathbf{Y}^k + \max_l \mathbf{P}^l\right)$, which we denote as simply $\mathbf{Y}^k + \mathbf{P}^l$.[4] Finally, we will consider x_1 and x_2 to be batches of $\mathbf{N}$ inputs each, in which case the NTK will be an $\mathbf{NO}\times\mathbf{NO}$ matrix, obtained by computing Eq. (2) for each pair of inputs. See §B for glossary.

3.1.2. Jacobian-vector Products (JVP) and Vector-Jacobian Products (VJP)

Following Maclaurin et al. (2015) we define

$$\mathrm{JVP}_{(\theta,x)}^f : \theta_t \in \mathbb{R}^{\mathbf{P}} \mapsto \left[\partial f\left(\theta, x\right)/\partial\theta\right] \theta_t \in \mathbb{R}^{\mathbf{O}}; \tag{4}$$

$$\mathrm{VJP}_{(\theta,x)}^f : f_c \in \mathbb{R}^{\mathbf{O}} \mapsto \left[\partial f\left(\theta, x\right)/\partial\theta\right]^T f_c \in \mathbb{R}^{\mathbf{P}}. \tag{5}$$

The JVP can be understood as pushing forward a tangent vector θ_t in weight space to a tangent vector in the space of outputs; by contrast the VJP pulls back a cotangent vector f_c in the space of outputs to a cotangent vector in weight space. These elementary operations enable forward and reverse mode AD respectively and serve as a basis for typical AD computations such as gradients, Jacobians, Hessians, etc.

The time cost of both is comparable to **FP** (see §D and Griewank & Walther (2008)). The memory cost of a JVP is **FP** as well (i.e. $\mathbf{Y}^k + \mathbf{P}^l$), while the memory cost of a VJP is generally $\mathbf{Y} + \mathbf{P}$, since it requires storing all $\mathbf{K}$ intermediate primitive outputs for efficient backprop and all $\mathbf{L}$ output cotangents. However, for the purpose of computing the NTK, we never need to store the whole Jacobian $\partial f/\partial\theta$, but only individual cotangents like $\partial f/\partial\theta^l$ to compute the sum in Eq. (2) layer-by-layer. Hence we consider VJP to cost $\mathbf{Y} + \mathbf{P}^l$ memory. To summarize, for a batch of $\mathbf{N}$ inputs,

- JVP costs $\mathbf{N}$ [**FP**] time; $\mathbf{N}\left[\mathbf{Y}^k\right] + \mathbf{P}$ memory.
- VJP costs $\mathbf{N}$ [**FP**] time; $\mathbf{N}\left[\mathbf{Y} + \mathbf{P}^l\right] + \mathbf{P}$ memory.

3.1.3. Jacobian

For NNs, the (reverse mode) Jacobian $\partial f/\partial\theta$ is most often computed via $\mathbf{O}$ VJP calls on rows of the identity matrix $I_{\mathbf{O}} \in \mathbb{R}^{\mathbf{O}\times\mathbf{O}}$, i.e.

$$\left[\partial f\left(\theta, x\right)/\partial\theta\right]^T = \left[\partial f\left(\theta, x\right)/\partial\theta\right]^T I_{\mathbf{O}} \in \mathbb{R}^{\mathbf{P}\times\mathbf{O}}, \tag{6}$$

and therefore costs $\mathbf{O}$ [VJP] time and memory apart from parameters and primitive outputs that can be reused across VJPs. Therefore, for a batch of $\mathbf{N}$ inputs,

> Jacobian costs $\mathbf{NO}$ [**FP**] time; $\mathbf{NO}\left[\mathbf{Y}^k + \mathbf{P}^l\right] + \mathbf{NY} + \mathbf{P}$ memory.

3.2. Jacobian contraction – the Baseline

This baseline method of evaluating the NTK consists in computing the Jacobians $\partial f/\partial\theta$ and contracting them as in Eq. (2). The contraction costs $\mathbf{N}^2\mathbf{O}^2\mathbf{P}$ time and $\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOP}^l$ memory, to store the result Θ_θ^f and individual layer-by-layer cotangents $\partial f/\partial\theta^l$. Including the cost of computing the cotangents via the batch Jacobian $\partial f/\partial\theta = \left[\partial f/\partial\theta^0, \dots, \partial f/\partial\theta^{\mathbf{L}}\right]$ from §3.1.3 we arrive at

> Jacobian contraction costs $\mathbf{NO}$ [**FP**] $+ \mathbf{N}^2\mathbf{O}^2\mathbf{P}$ time; $\mathbf{N}^2\mathbf{O}^2 + \mathbf{NO}\left[\mathbf{Y}^k + \mathbf{P}^l\right] + \mathbf{NY} + \mathbf{P}$ memory.

In summary, Jacobian contraction performs $\mathbf{NO}$ forward passes followed by an expensive $\mathbf{N}^2\mathbf{O}^2\mathbf{P}$ contraction. Next we demonstrate how to reduce the contraction cost.

[4] To declutter notation throughout this work, in time and memory complexity expressions, we (1) omit the $\mathcal{O}$ symbol, and (2) imply taking the maximum over any free index.

3.3. NTK-vector products

Consider the NTK-vector product function (for $\mathbf{N} = 1$):

$$\Theta_\theta^f \mathrm{VP}\colon v \in \mathbb{R}^{\mathbf{O}} \mapsto \Theta_\theta^f v \in \mathbb{R}^{\mathbf{O}}.$$

Taking the NTK-vector product with $\mathbf{O}$ columns of the identity matrix $I_{\mathbf{O}}$ yields the full NTK, i.e. $\Theta_\theta^f I_{\mathbf{O}} = \Theta_\theta^f$. Expanding $\Theta_\theta^f \mathrm{VP}(v)$ as

$$\Theta_\theta^f v = \left[\partial f\left(\theta, x_1\right)/\partial\theta\right]\left[\partial f\left(\theta, x_2\right)/\partial\theta\right]^T v = \tag{7}$$

$$= \left[\partial f\left(\theta, x_1\right)/\partial\theta\right] \mathrm{VJP}^f_{(\theta, x_2)}\left(v\right) = \tag{8}$$

$$= \mathrm{JVP}^f_{(\theta, x_1)}\left[\mathrm{VJP}^f_{(\theta, x_2)}\left(v\right)\right], \tag{9}$$

where we have observed that the NTK-vector product can be expressed as a composition of a JVP and a VJP. The cost of computing Θ_θ^f is then asymptotically the cost of the Jacobian, since it consists of $\mathbf{O}$ VJPs followed by $\mathbf{O}$ (cheaper) JVPs, therefore $\mathbf{O}\,[\mathbf{FP}]$ time and $\mathbf{O}\left[\mathbf{Y}^k + \mathbf{P}^l\right] + \mathbf{Y} + \mathbf{P}$ memory. In the batched setting Eq. (7) is repeated for each pair of inputs, and therefore time increases by a factor of $\mathbf{N}^2$ to become $\mathbf{N}^2\mathbf{O}\,[\mathbf{FP}]$. However, the memory cost grows only linearly in $\mathbf{N}$ (except for the cost of storing the NTK of size $\mathbf{N}^2\mathbf{O}^2$), since intermediate primitive outputs and tangents/cotangents can be computed for each batch x_1 and x_2 separately and then reused for every pairwise combination. Therefore the memory cost is asymptotically the cost to store the NTK and compute the Jacobian. Altogether,

> NTK-vector products cost $\mathbf{N}^2\mathbf{O}\,[\mathbf{FP}]$ time; $\mathbf{N}^2\mathbf{O}^2 + \mathbf{N}\mathbf{O}\left[\mathbf{Y}^k + \mathbf{P}^l\right] + \mathbf{N}\mathbf{Y} + \mathbf{P}$ memory.

In summary, NTK-vector products eliminate the costly $\mathbf{N}^2\mathbf{O}^2\mathbf{P}$ contraction of Jacobian contraction, but perform $\mathbf{N}^2\mathbf{O}$ forward passes (as opposed to $\mathbf{N}\mathbf{O}$), and the memory requirement is identical. As a result, this method is beneficial for small $\mathbf{N}$, and for networks with a cheap forward pass $\mathbf{FP}$ relative to $\mathbf{OP}$, which is always the case for FCNs (§4.1), but not necessarily for CNNs (§4.2).

3.4. Structured derivatives

Rewriting Θ_θ^f from Eq. (2) using the chain rule in terms of the primitive outputs y^k, we find:

$$\Theta_\theta^f = \sum_{l, k_1, k_2}\left(\frac{\partial f_1}{\partial y_1^{k_1}}\frac{\partial y_1^{k_1}}{\partial\theta^l}\right)\left(\frac{\partial f_2}{\partial y_2^{k_2}}\frac{\partial y_2^{k_2}}{\partial\theta^l}\right)^T \tag{10}$$

$$= \sum_{l, k_1, k_2}\left(\frac{\partial f_1}{\partial y_1^{k_1}}\frac{\partial y_1^{k_1}}{\partial\theta^l}\frac{\partial y_2^{k_2}}{\partial\theta^l}^T\frac{\partial f_2}{\partial y_2^{k_2}}^T\right) \tag{11}$$

$$=: \sum_{l, k_1, k_2}\Theta_\theta^f\left[l, k_1, k_2\right], \tag{12}$$

where we define $f_i := f(\theta, x_i)$, and only consider $\partial y_i^{k_i}/\partial\theta^l$ to be non-zero if θ^l is a direct input to $y_i^{k_i}$. We have also defined $\Theta_\theta^f\left[l, k_1, k_2\right]$ to be individual summands.

Both Jacobian contraction and NTK-vector products perform this sum of contractions via VJPs and JVPs, without explicit instantiation of primitive Jacobians $\partial y_i^{k_i}/\partial\theta^l$. However, while VJPs and JVPs themselves are guaranteed to be computationally optimal (§D), higher order computations like their composition (NTK-vector products) or contraction (Jacobian contraction) are not. Specifically, each $\Theta_\theta^f\left[l, k_1, k_2\right]$ from Eq. (10) is a matrix-Jacobian-Jacobian-matrix product (**MJJMP**), which, as we will show shortly, can't always be evaluated optimally with VJPs and JVPs.

The idea of Structured derivatives is to design rules for efficient computation of MJJMPs, similarly to AD rules for JVPs and VJPs.

From Eq. (10), in the general case this requires hand-made rules for all pairwise combinations of primitives y_1 and y_2, of which there are $136^2 > 10,000$ in JAX, and even more in Tensorflow (Abadi et al., 2016) and PyTorch (Paszke et al., 2019) (see §M). We dramatically reduce this number by:

1. Linearization. It follows from Eq. (4), that $\Theta_\theta^f = \Theta_\theta^{\mathrm{JVP}^f_{(\theta,\cdot)}}$, i.e. the NTK of f evaluated at parameters θ is equal to the NTK of the JVP of f given primal θ. JVP is a linear function of input tangents θ, and therefore we only need to implement efficient MJJMPs for linear primitives, of which JAX has only 56.[5]

2. MJJMPs through structured derivatives. We further reduce the necessary MJJMP rule count from 56^2 down to only 56 by decomposing an MJJMP rule into two parts:

1. **Structured derivative rule**. Given a single primitive y, this rule identifies the smallest subarray of $\partial y/\partial\theta^l$ sufficient to reconstruct the entire primitive Jacobian $\partial y/\partial\theta^l$, and the (constant and negligible in memory size) metadata necessary for the reconstruction. For example, if $x \in \mathbb{R}^{\mathbf{W}}$, $\theta^l \in \mathbb{R}^{\mathbf{W}\times\mathbf{W}}$, and $y\left(\theta^l\right) = \theta^l x \in \mathbb{R}^{\mathbf{W}}$ (matrix-vector multiplication), then $\partial y/\partial\theta^l = I_{\mathbf{W}} \otimes x^T \in \mathbb{R}^{\mathbf{W}\times\mathbf{W}^2}$, and the rule will indicate that (1) only the subarray $\left[\partial y/\partial\theta^l\right]_{1,:\mathbf{W}} \in \mathbb{R}^{1\times\mathbf{W}}$,[6] needs to be computed (which is equal to x^T in this case), and (2) that the entire primitive Jacobian can be reconstructed as $\partial y/\partial\theta^l = I \otimes \left[\partial y/\partial\theta^l\right]_{1,:\mathbf{W}}$. In other words, this rule annotates linear primitives y with the structure of their Jacobians, such as block diagonal, constant-block diagonal, or tiling along certain dimensions.

[5] A linear function can contain nonlinear primitives. However, linearizing any function in JAX is guaranteed to produce only linear primitives (Frostig et al., 2021; Radul et al., 2022).

[6] We define $[A]_{i,:j} := [A_{i,1}, \ldots, A_{i,j}] \in \mathbb{R}^{1\times j}$.

2. **MJJMPs with structured Jacobians.** Given input tensors A, B, C, D, where B and C are provided in the structured form as described above (i.e. only small subarrays along with their metadata) this rule efficiently computes the 4-way contraction $ABCD$ (i.e. the NTK summand $\Theta_\theta^f[l, k_1, k_2]$). This amounts to using `np.einsum` with the optimal contraction order and adjusting its instructions based on provided metadata. For example, if $B = I_\mathbf{W} \otimes b^T \in \mathbb{R}^{\mathbf{W}\times\mathbf{W}^2}$ and $C = I_\mathbf{W} \otimes c^T \in \mathbb{R}^{\mathbf{W}\times\mathbf{W}^2}$ (for $b, c \in \mathbb{R}^\mathbf{W}$), then

$$ABCD = A\left(I \otimes b^T\right)\left(I \otimes c^T\right)^T D = \quad (13)$$
$$= A\left(I \otimes b^T c\right) D = \left(b^T c\right) AD, \quad (14)$$

where were able to pull out $b^T c$ since it is a scalar. As we will see in §4 and §E, this and other similar contraction rules can enable significant speedups.

Therefore we avoid implementing 56^2 MJJMP rules by instead having (1) a single routine to perform 4-way tensor contractions with structured tensors, and (2) 56 rules annotating the structure in the 56 linear primitive Jacobians. We list all these structures and associated MJJMP costs in §E. Our approach does not guarantee optimality for the NTK of an arbitrary function, however, as we show in §4, it is asymptotically better than Jacobian contraction for FCNs and CNNs, and can provide orders of magnitude speedups in much more complex contemporary ImageNet models (§5).

3. Focusing on MJJMPs for typical operations. Many of the 56 linear JAX primitives are trivial to implement or rarely arise in NNs. At the time of writing we have only annotated 21 linear primitives (Table 4), which was sufficient for the empirical speedups observed in §4 and §5.

Summary. Structured derivatives amount to evaluating the sum of MJJMPs in Eq. (10), where (1) only small subarrays of primitive Jacobians $\partial y_i^{k_i}/\partial\theta^l$ are instantiated, and (2) MJJMPs leverage the structure of these primitive Jacobians for efficient contractions. Together, this incurs

1. The cost of computing primitive output cotangents $\partial f_i/\partial y_i^{k_i}$ for Eq. (10), which is equivalent to the cost of the reverse mode Jacobian (§3.1.3), *less* the cost of computing (**NOP**) and storing (**NP**^{l}) weight-space cotangents $\partial f_i/\partial\theta^l$, since they aren't used in Eq. (10), i.e. **NO** [**FP**] time[7] and $\mathbf{NOY}^k + \mathbf{NY} + \mathbf{P}$ memory.

2. The cost of computing primitive Jacobian $\partial y_i^{k_i}/\partial\theta^l$ subarrays, denoted as $\mathbf{J}_l^{k_i}$ with $\mathbf{J} := \sum_{l,k_1}\mathbf{J}_l^{k_1} + \sum_{l,k_2}\mathbf{J}_l^{k_2}$. This amounts to **NJ** time and $\mathbf{NJ}_l^k$ memory.

3. The cost of evaluating $\Theta_\theta^f[l, k_1, k_2]$ via the efficient MJJMP, which we denote as simply **MJJMP** and substitute specific values based on the primitive. Required memory to store the result is $\mathbf{N}^2\mathbf{O}^2$.

[7]Note that **NOP** time saving is not reflected in the asymptotic cost since it is dominated by **NO** [**FP**] required to compute primitive output cotangents. However, as we will see in Fig. 1, it often provides a substantial practical benefit, up to allowing to compute the NTK faster than computing the two Jacobians themselves.

We add up these costs below and in Table 1, and show in §4 and §5 how they are beneficial in most practical settings.

> Structured derivatives cost **NO** [**FP**] + **MJJMP** + **N** [**J** − **OP**] time; $\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOY}^k + \mathbf{NJ}_l^k + \mathbf{NY} + \mathbf{P}$ memory.

4. Examples

The three algorithms in §3 (and our implementations) apply to any differentiable function f, but the resulting complexities depend on variables such as **FP** and **P**, which depend on f (Table 1). Below we compute and compare all three complexities for a deep FCN (§4.1) and CNN (§4.2), summarizing the results in Table 2 and Table 3 respectively.

4.1. FCN NTK Complexity

We apply our algorithms from §3 to FCNs with **L** hidden layers of width **W**. For simplicity we assume no biases, and $x \in \mathbb{R}^\mathbf{W}$, i.e. inputs of same size as the width. We define $y^l := \theta^l x^l$, and $x^l := \phi\left(y^{l-1}\right)$ for $l > 0$, with $x^0 := x$. Output is $f(x, \theta) := y^\mathbf{L}$. See Fig. 5 (top) for $\mathbf{L} = 2$.

In this case $\mathbf{K} = 2\mathbf{L}+1$ ($\mathbf{L}+1$ matrix-vector multiplications and **L** nonlinearities), $\mathbf{P}^l = \mathbf{W}^2$ for $l < \mathbf{L}$ and **OW** for the top layer $l = \mathbf{L}$, $\mathbf{P} = \mathbf{LW}^2 + \mathbf{OW}$, $\mathbf{Y}^k = \mathbf{W}$ for $k < \mathbf{K}$ and **O** for $k = \mathbf{K}$, and $\mathbf{Y} \sim \mathbf{LW} + \mathbf{O}$. Finally, a single forward pass $\mathbf{FP} \sim \mathbf{LW}^2 + \mathbf{OW}$ time and $\mathbf{W}^2 + \mathbf{OW}$ memory.

Plugging the above into the cost of the baseline algorithm Jacobian contraction in §3.2, we obtain

> FCN Jacobian contraction costs $\mathbf{N}^2\mathbf{O}^2\mathbf{LW}^2 + \mathbf{N}^2\mathbf{O}^3\mathbf{W}$ time; $\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOW}^2 + \mathbf{NO}^2\mathbf{W} + \mathbf{NLW} + \mathbf{LW}^2$ memory.

Similarly, the cost of NTK-vector products from §3.3 is

> FCN NTK-vector products cost $\mathbf{N}^2\mathbf{OLW}^2 + \mathbf{N}^2\mathbf{O}^2\mathbf{W}$ time; $\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOW}^2 + \mathbf{NO}^2\mathbf{W} + \mathbf{NLW} + \mathbf{LW}^2$ memory.

For Structured derivatives (§3.4), we additionally need to derive values of **J** and **MJJMP**. For an arbitrary primitive, **J** and **MJJMP** costs are derived by (1) looking up the type of structure in the primitive Jacobian in Table 4, followed by (2) extracting the costs for a given structure from Table 5

Method	Time	Memory	Use when
Jacobian contraction	$\mathbf{N}\ \mathbf{O}\,[\mathbf{FP}] + \mathbf{N}^2\mathbf{O}^2\mathbf{P}$	$\mathbf{N}^2\mathbf{O}^2 + \mathbf{NO}\left[\mathbf{Y}^k + \mathbf{P}^l\right] + \mathbf{NY} + \mathbf{P}$	$\mathbf{P} < \mathbf{Y}$, small $\mathbf{O}$
NTK-vector products	$\mathbf{N}^2\mathbf{O}\,[\mathbf{FP}]$	$\mathbf{N}^2\mathbf{O}^2 + \mathbf{NO}\left[\mathbf{Y}^k + \mathbf{P}^l\right] + \mathbf{NY} + \mathbf{P}$	$\mathbf{FP} < \mathbf{OP}$, large $\mathbf{O}$, small $\mathbf{N}$
Structured derivatives	$\mathbf{N}\ \mathbf{O}\,[\mathbf{FP}] + \mathbf{MJJMP} + \mathbf{NJ}$	$\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOY}^k + \mathbf{NJ}^k_l + \mathbf{NY} + \mathbf{P}$	$\mathbf{FP} > \mathbf{OP}$, large $\mathbf{O}$, large $\mathbf{N}$

Table 1. **Generic NTK computation costs**. NTK-vector products trade-off contractions for more **FP**. Structured derivatives usually save both time and memory. See §3.1.1 and §3.4 for notation, and §B for a glossary of symbols.

Method	Time	Memory	Use when
Jacobian contraction	$\mathbf{N}^2\mathbf{O}^2\mathbf{LW}^2 + \mathbf{N}^2\mathbf{O}^3\mathbf{W}$	$\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOW}^2 + \mathbf{NO}^2\mathbf{W} + \mathbf{NLW} + \mathbf{LW}^2$	Don't
NTK-vector products	$\mathbf{N}^2\mathbf{OLW}^2 + \mathbf{N}^2\mathbf{O}^2\ \ \mathbf{W}$	$\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOW}^2 + \mathbf{NO}^2\mathbf{W} + \mathbf{NLW} + \mathbf{LW}^2$	$\mathbf{O} > \mathbf{W}$ or $\mathbf{N} = 1$
Structured derivatives	$\mathbf{N}\ \mathbf{OLW}^2 + \mathbf{N}^2\mathbf{O}^2\mathbf{LW}\ + \mathbf{N}^2\mathbf{O}^3$	$\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOW}\ + \mathbf{NLW} + \mathbf{LW}^2$	$\mathbf{O} < \mathbf{W}$ or $\mathbf{L} = 1$

Table 2. **FCN NTK computation cost.** The costs are obtained by substituting into Table 1 specific values for **FP**, **P**, **Y**, **J**, and **MJJMP** that correspond to an FCN as described in §4.1. NTK-vector products allow a reduction of the time complexity, while Structured derivatives reduce both time and memory complexity. See Fig. 1 for empirical confirmation with FLOPs and wall-clock time. See §4.1 for discussion and notation (§B for the full glossary of symbols).

Method	Time	Memory	Use when
Jacobian contraction	$\mathbf{N}\ \mathbf{O}\left[\mathbf{LDFW}^2 + \mathbf{OW}\right] + \mathbf{N}^2\mathbf{O}^2\left[\mathbf{LFW}^2 + \mathbf{OW}\right]$	$\mathbf{N}^2\mathbf{O}^2 + \mathbf{NO}\left[\mathbf{DW} + \mathbf{FW}^2 + \mathbf{OW}\right] + \mathbf{N}\left[\mathbf{LDW}\right] + \left[\mathbf{LFW}^2 + \mathbf{OW}^2\right]$	$\mathbf{D} > \mathbf{OW}$
NTK-vector products	$\mathbf{N}^2\mathbf{O}\left[\mathbf{LDFW}^2 + \mathbf{OW}\right]$	$\mathbf{N}^2\mathbf{O}^2 + \mathbf{NO}\left[\mathbf{DW} + \mathbf{FW}^2 + \mathbf{OW}\right] + \mathbf{N}\left[\mathbf{LDW}\right] + \left[\mathbf{LFW}^2 + \mathbf{OW}^2\right]$	$\mathbf{N} = 1$
Structured derivatives	$\mathbf{N}\ \mathbf{O}\left[\mathbf{LDFW}^2 + \mathbf{OW}\right] + \mathbf{N}^2\mathbf{O}^2\left[\mathbf{L}\min\left(\mathbf{FW}^2, \mathbf{DW} + \frac{\mathbf{DFW}^2}{\mathbf{O}}, \mathbf{DW} + \frac{\mathbf{D}^2\mathbf{W}}{\mathbf{O}} + \frac{\mathbf{D}^2\mathbf{FW}}{\mathbf{O}^2}\right) + \mathbf{O}\right]$	$\mathbf{N}^2\mathbf{O}^2 + \mathbf{NO}\left[\mathbf{DW}\right]\ +\ \mathbf{NDFW}\ + \mathbf{N}\left[\mathbf{LDW}\right] + \left[\mathbf{LFW}^2 + \mathbf{OW}^2\right]$	$\mathbf{D} < \mathbf{OW}$

Table 3. **CNN NTK computation cost** for a CNN with $\mathbf{D}$ pixels and filter size $\mathbf{F}$. Structured derivatives reduce time complexity, and have lower memory cost if $\mathbf{D} < \mathbf{OW}$, which is a common setting. See Fig. 2 for experiments with ResNets, §4.2 for discussion, Table 2 for FCN, Table 1 for generic cost analysis, and §B for a glossary of symbols.

(see §E). We apply this formal approach in §E.9, but for demonstration purposes below present the derivation that does not require referencing the Appendix.

We first note that we only need to consider $k_1 = k_2 = 2l+1$ indices in Eq. (10), since all other summands are zero due to absence of weight sharing between layers. For matrix-vector multiplication $y_i^l = \theta^l x_i^l$ our rules indicate (per example given in §3.4) that $\partial y_i^l / \partial \theta^l = I_{\mathbf{W}} \otimes \left[\partial y_i^l / \partial \theta^l\right]_{1,:\mathbf{W}} \in \mathbb{R}^{\mathbf{W} \times \mathbf{W}^2}$, and command to only compute $\left[\partial y_i^l / \partial \theta^l\right]_{1,:\mathbf{W}} \in \mathbb{R}^{1 \times \mathbf{W}}$ (which is ${x_i^l}^T$ in this case). Therefore $\mathbf{J}_l^{2l+1} = \mathbf{W}$, and $\mathbf{J} = 2\sum_{l=1}^{\mathbf{L}} \mathbf{J}_l^{2l+1} \sim \mathbf{LW}$.

Finally, the efficient MJJMP for this structured $\partial y_i^l / \partial \theta^l$ can be computed, analogously to Eq. (13), as follows for $l < \mathbf{L}$:

$$\underbrace{\Theta_\theta^f\left[l, 2l+1, 2l+1\right]}_{\mathbf{O} \times \mathbf{O}} = \frac{\partial f_1}{\partial y_1^l} \frac{\partial y_1^l}{\partial \theta^l} \frac{\partial y_2^l}{\partial \theta^l}^T \frac{\partial f_2}{\partial y_2^l}^T = \tag{15}$$

$$= \underbrace{\frac{\partial f_1}{\partial y_1^l}}_{\mathbf{O} \times \mathbf{W}} \left(I_{\mathbf{W}} \otimes \underbrace{{x_1^l}^T}_{1 \times \mathbf{W}}\right) \left(I_{\mathbf{W}} \otimes \underbrace{{x_2^l}^T}_{1 \times \mathbf{W}}\right)^T \underbrace{\frac{\partial f_2}{\partial y_2^l}^T}_{\mathbf{W} \times \mathbf{O}} = \tag{16}$$

$$= \left(\underbrace{{x_1^l}^T}_{1 \times \mathbf{W}} \underbrace{x_2^l}_{\mathbf{W} \times 1}\right) \underbrace{\frac{\partial f_1}{\partial y_1^l}}_{\mathbf{O} \times \mathbf{W}} \underbrace{\frac{\partial f_2}{\partial y_2^l}^T}_{\mathbf{W} \times \mathbf{O}}, \tag{17}$$

which can be contracted in only $\mathbf{O}^2\mathbf{W}$ time. An analogous derivation applied to $l = \mathbf{L}$ yields $\mathbf{O}^3 + \mathbf{W}$ time. Therefore the total contraction cost is $\mathbf{MJJMP} \sim \mathbf{N}^2\mathbf{LO}^2\mathbf{W} + \mathbf{N}^2\mathbf{O}^3$, when accounting for depth $\mathbf{L}$ and batch size $\mathbf{N}$. Altogether,

> FCN Structured derivatives cost $\mathbf{NOLW}^2 + \mathbf{N}^2\mathbf{O}^2\mathbf{LW} + \mathbf{N}^2\mathbf{O}^3$ time; $\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOW} + \mathbf{NLW} + \mathbf{LW}^2$ memory.

Summary. We summarize all FCN costs in Table 2. We conclude that Structured derivatives and NTK-vector products allow a reduction in the time cost of NTK computation in different ways, while Structured derivatives also reduce memory requirements. Structured derivatives are beneficial for wide networks, with large $\mathbf{W}$, and NTK-vector products are beneficial for networks with large outputs $\mathbf{O}$.

We confirm our predictions with FLOPs measurements in Fig. 1. We further confirm our methods provide orders of magnitude speed-ups and memory savings on all major hardware platforms in Fig. 1 (right) and Fig. 3. However, we notice that time measurements often deviate from predictions due to unaccounted constant overheads of various methods, hardware specifics, padding, and the behavior of the XLA compiler. We find Structured derivatives to almost always outperform NTK-vector products.

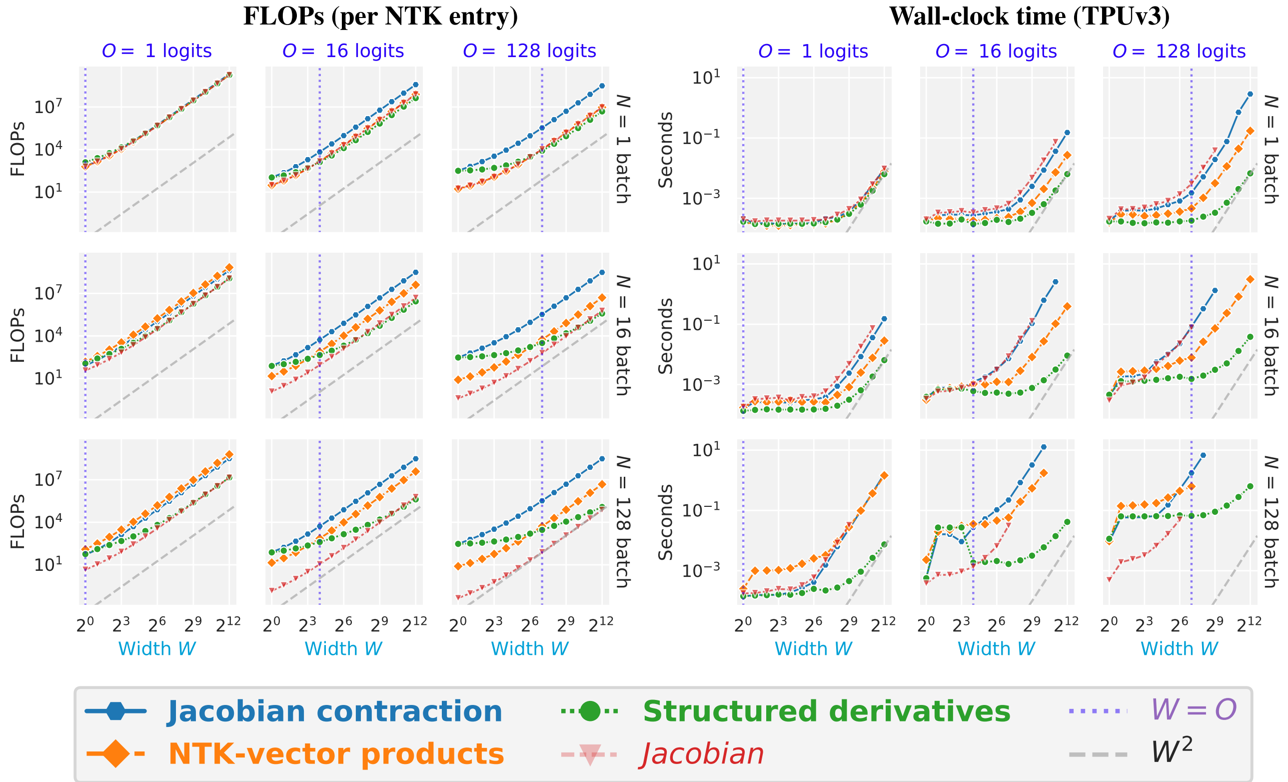

Figure 1. **FLOPs (left) and wall-clock time (right) of computing the NTK for a 10-layer ReLU FCN.** As predicted by Table 2, our methods almost always outperform Jacobian contraction, allowing orders of magnitude speed-ups and memory improvements for realistic problem sizes.

Left: FLOPs per NTK entry. We confirm several specific theoretical predictions from §4.1:

1. NTK-vector products are the best performing method for **N** = 1, and have cost equivalent to Jacobian for any width **W** or output size **O** (top row);
2. NTK-vector products offer an **O**-fold improvement over Jacobian contraction (left to right columns);
3. NTK-vector products are equivalent to Jacobian contraction for **O** = 1 (leftmost column);
4. Structured derivatives outperform NTK-vector products iff **O** < **W** (**O** = **W** are plotted as pale vertical lines, which is where Structured derivatives and NTK-vector products intersect);
5. Structured derivatives approach the cost of Jacobian in the limit of large width **W** (left to right);
6. All methods, as expected, scale quadratically with width **W** (pale grey dashed line depicts $\mathbf{W}^2$ scaling).

More: see Fig. 3 for other hardware platforms, and §N for details.

Right: Wall-clock runtime. In real applications, given the XLA compiler and hardware specifics, we observe that:

1. NTK-vector products improve upon Jacobian contraction for **O** > 1, but the effect is not perfectly robust (see bottom row for small **W** and Fig. 3, notably GPU platforms);
2. Structured derivatives robustly outperform all other methods, including simply computing the Jacobian, as discussed in §3.4;
3. Structured derivatives have lower memory footprint, and reach up to 8x larger widths (bottom right; missing points indicate out-of-memory), i.e. can process models up to 64x larger than other methods, as discussed in §3.4;
4. All methods have a smaller memory footprint than Jacobian (see §3.1.3).

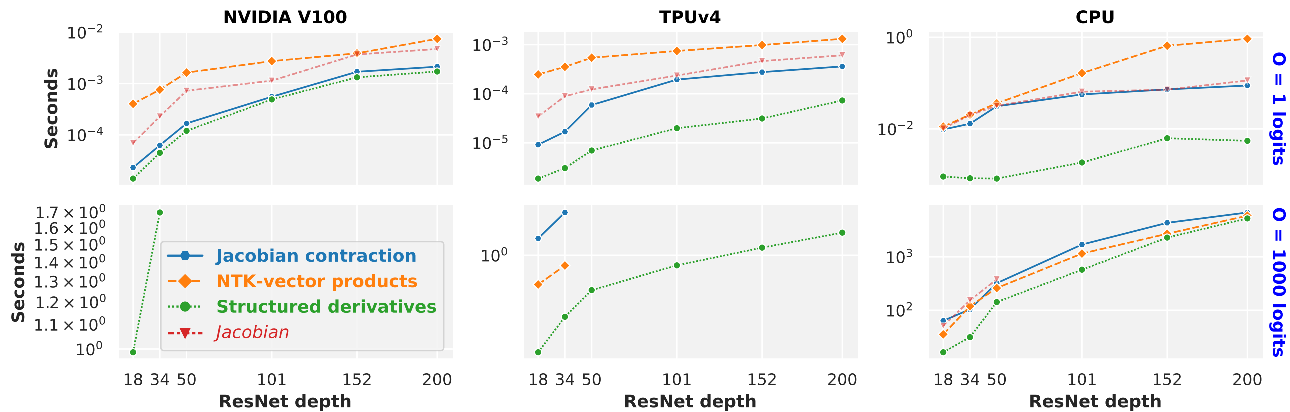

Figure 2. **Wall-clock time of computing an NTK for several ResNet sizes on a pair of ImageNet images**. Structured derivatives allow the NTK to be computed faster and for larger models (see bottom row – missing points indicate out-of-memory). NTK-vector products, as predicted in §3.3 and Table 1, are advantageous for large **O** (bottom row), but are suboptimal when the cost of the forward pass **FP** is large relative to the output size and the number of parameters **OP**, e.g. when there is a lot of weight sharing (see Table 3 and Table 1), which is the case for convolutions, notably for **O** $= 1$ (top). See Fig. 4 for more ImageNet models, §4.2 for CNN NTK complexity analysis, and §N for experimental details.

4.2. CNN NTK Complexity

We perform analogous derivations for a CNN with **W** channels, **D** pixels, and filter size **F** in §F. We arrive at Table 3, and make two observations specific to CNNs.

First, the speeds of NTK-vector products and Jacobian contraction are now much more similar, due to the higher cost of the forward pass **FP** relative to **P** (i.e. weight sharing), and how they perform will depend on the specific values of parameters. We confirm this in our experiments on ImageNet models in §5, where NTK-vector products typically underperform for **O** $= 1$, but outperform for **O** $= 1000$.

Secondly, Structured derivatives continue to perform faster than Jacobian contraction, but the relative memory costs depend on other hyperparameters, requiring **D** $<$ **OW**. This is a common case for ImageNet with **O** $= 1000$, and is confirmed in our experiments in Fig. 2 and Fig. 4 (bottom).

5. ImageNet Experiments

In §4 we have derived asymptotic time and memory benefits of NTK-vector products and Structured derivatives over the baseline Jacobian contraction for FCNs and CNNs. However, contemporary architectures rarely resemble vanilla feedforward networks, but instead result in much more complex computational graphs comprised of many different primitives, making complexity analysis impractical.

We therefore evaluate our methods in the wild, and confirm computational benefits on full ImageNet models in Fig. 2 (ResNets, He et al. (2016)) and Fig. 4 (WideResNets, Zagoruyko & Komodakis (2016); Vision Transformers and Transformer-ResNet hybrids Dosovitskiy et al. (2021); Steiner et al. (2021); and MLP-Mixers Tolstikhin et al. (2021)). Computing the full **O** $\times$ **O** $= 1000 \times 1000$ NTK is often only possible with Structured derivatives.

6. Implementation

All algorithms are implemented in JAX[8] (Bradbury et al., 2018) and integrated into Neural Tangents (Novak et al., 2020). Jacobian contraction and NTK-vector products are built with core operations such as `vjp`, `jvp`, and `vmap`. Structured derivatives are implemented as a Jaxpr interpreter, built on top of the JAX reverse mode AD interpreter.

Owing to the nuanced trade-offs between different methods in the general case, we release all implementations within a single function that allows the user to manually select implementation. We also include an automated setting which will perform FLOPs analysis for each method at compilation time and automatically choose the most efficient one.

7. Conclusion

We have performed the first extensive analysis of the computational complexity of the NTK, and have shown how it can be improved dramatically with mixed-order AD (NTK-vector products), or with a custom interpreter for more efficient higher-order AD operations (Structured derivatives).

The NTK computation is similar to many other objects of interest in machine learning, such as the Gauss-Newton or the Fisher Information matrix, and we look forward to extensions of our algorithms to more settings in future work.

[8] See §M for discussion about other frameworks.

Acknowledgements

We thank Lechao Xiao for useful discussion, review and comments on the initial version of this manuscript, and Jaehoon Lee for useful discussion and code review.

We also thank Shaobo Hou for his work on and help with TF2Jax, and the JAX team for their help and advice on JAX and Jax2TF.

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, 2016. *Cited on page 2, 4, 26.*

Adlam, B., Lee, J., Xiao, L., Pennington, J., and Snoek, J. Exploring the uncertainty properties of neural networks' implicit priors in the infinite-width limit. In *International Conference on Learning Representations*, 2020. *Cited on page 1, 23.*

Arfian, W. Ukraine vectors by vecteezy. https://www.vecteezy.com/vector-art/7506324-stand-with-ukraine-text-with-ukraine-flag-ribbon-and-ukraine-map-vector-design-on-a-dark-blue-background, 2022. *Cited on page 27.*

Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R. R., and Wang, R. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*, pp. 8141–8150. Curran Associates, Inc., 2019a. *Cited on page 2.*

Arora, S., Du, S. S., Hu, W., Li, Z., and Wang, R. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019b. *Cited on page 24.*

Arora, S., Du, S. S., Li, Z., Salakhutdinov, R., Wang, R., and Yu, D. Harnessing the power of infinitely wide deep nets on small-data tasks. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=rkl8sJBYvH. *Cited on page 1.*

Babuschkin, I., Baumli, K., Bell, A., Bhupatiraju, S., Bruce, J., Buchlovsky, P., Budden, D., Cai, T., Clark, A., Danihelka, I., Fantacci, C., Godwin, J., Jones, C., Hennigan, T., Hessel, M., Kapturowski, S., Keck, T., Kemaev, I., King, M., Martens, L., Mikulik, V., Norman, T., Quan, J., Papamakarios, G., Ring, R., Ruiz, F., Sanchez, A., Schneider, R., Sezener, E., Spencer, S., Srinivasan, S., Stokowiec, W., and Viola, F. The DeepMind JAX Ecosystem, 2020. URL http://github.com/deepmind. *Cited on page 27.*

Bahri, Y., Dyer, E., Kaplan, J., Lee, J., and Sharma, U. Explaining neural scaling laws. *arXiv preprint arXiv:2102.06701*, 2021. *Cited on page 1.*

Bai, J., Lu, F., Zhang, K., et al. Onnx: Open neural network exchange. https://github.com/onnx/onnx, 2019. *Cited on page 27.*

Belkin, M., Hsu, D., Ma, S., and Mandal, S. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019. *Cited on page 1.*

Borovykh, A. A gaussian process perspective on convolutional neural networks. *arXiv preprint arXiv:1810.10798*, 2018. *Cited on page 1.*

Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., and Wanderman-Milne, S. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github.com/google/jax. *Cited on page 2, 8, 26, 27.*

Brock, A., De, S., and Smith, S. L. Characterizing signal propagation to close the performance gap in unnormalized resnets. *arXiv preprint arXiv:2101.08692*, 2021a. *Cited on page 1.*

Brock, A., De, S., Smith, S. L., and Simonyan, K. High-performance large-scale image recognition without normalization. *arXiv preprint arXiv:2102.06171*, 2021b. *Cited on page 1.*

Chen, W., Gong, X., and Wang, Z. Neural architecture search on imagenet in four gpu hours: A theoretically inspired perspective. In *International Conference on Learning Representations*, 2021a. *Cited on page 2, 26.*

Chen, X., Hsieh, C.-J., and Gong, B. When vision transformers outperform resnets without pretraining or strong data augmentations, 2021b. *Cited on page 1, 2, 26.*

Dauphin, Y. N. and Schoenholz, S. Metainit: Initializing learning by learning to initialize. *Advances in Neural Information Processing Systems*, 32, 2019. *Cited on page 1, 26.*

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009. *Cited on page 2.*

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=YicbFdNTTy. *Cited on page 2, 8, 14.*

Du, S. S., Hou, K., Salakhutdinov, R. R., Poczos, B., Wang, R., and Xu, K. Graph neural tangent kernel: Fusing graph neural networks with graph kernels. *Advances in neural information processing systems*, 32, 2019. *Cited on page 24.*

Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1126–1135. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/finn17a.html. *Cited on page 2.*

Franceschi, J.-Y., de Bézenac, E., Ayed, I., Chen, M., Lamprier, S., and Gallinari, P. A neural tangent kernel perspective of gans. *arXiv preprint arXiv:2106.05566*, 2021. *Cited on page 1.*

Frostig, R., Johnson, M. J., Maclaurin, D., Paszke, A., and Radul, A. Decomposing reverse-mode automatic differentiation. *arXiv preprint arXiv:2105.09469*, 2021. *Cited on page 4, 26.*

Garriga-Alonso, A., Aitchison, L., and Rasmussen, C. E. Deep convolutional networks as shallow gaussian processes. In *International Conference on Learning Representations*, 2019. *Cited on page 1.*

Griewank, A. and Walther, A. *Evaluating Derivatives*. Society for Industrial and Applied Mathematics, second edition, 2008. doi: 10.1137/1.9780898717761. URL https://epubs.siam.org/doi/abs/10.1137/1.9780898717761. *Cited on page 3, 17.*

Grosse, R. Neural net training dynamics, January 2021. URL https://www.cs.toronto.edu/~rgrosse/courses/csc2541_2021/readings/L02_Taylor_approximations.pdf. *Cited on page 25.*

Hanin, B. and Nica, M. Finite depth and width corrections to the neural tangent kernel. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=SJgndT4KwB. *Cited on page 24.*

He, B., Lakshminarayanan, B., and Teh, Y. W. Bayesian deep ensembles via the neural tangent kernel. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/0b1ec366924b26fc98fa7b71a9c249cf-Abstract.html. *Cited on page 1.*

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016. *Cited on page 2, 8.*

Heek, J., Levskaya, A., Oliver, A., Ritter, M., Rondepierre, B., Steiner, A., and van Zee, M. Flax: A neural network library and ecosystem for JAX, 2020. URL http://github.com/google/flax. *Cited on page 24, 27.*

Hennigan, T., Cai, T., Norman, T., and Babuschkin, I. Haiku: Sonnet for JAX, 2020. URL http://github.com/deepmind/dm-haiku. *Cited on page 24.*

Horace He, R. Z. functorch: Jax-like composable function transforms for pytorch. https://github.com/pytorch/functorch, 2021. *Cited on page 27.*

Hron, J., Bahri, Y., Novak, R., Pennington, J., and Sohl-Dickstein, J. Exact posterior distributions of wide bayesian neural networks, 2020a. *Cited on page 1.*

Hron, J., Bahri, Y., Sohl-Dickstein, J., and Novak, R. Infinite attention: NNGP and NTK for deep attention networks. In *International Conference on Machine Learning*, 2020b. *Cited on page 1, 24.*

Hu, J., Shen, J., Yang, B., and Shao, L. Infinitely wide graph convolutional networks: semi-supervised learning via gaussian processes. *arXiv preprint arXiv:2002.12168*, 2020. *Cited on page 1.*

Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, 2018. *Cited on page 1, 23, 26.*

Khan, M. E. E., Immer, A., Abedi, E., and Korzepa, M. Approximate inference turns deep networks into gaussian processes. In *Advances in neural information processing systems*, 2019. *Cited on page 1.*

Lee, J., Bahri, Y., Novak, R., Schoenholz, S., Pennington, J., and Sohl-dickstein, J. Deep neural networks as gaussian processes. In *International Conference on Learning Representations*, 2018. *Cited on page 1.*

Lee, J., Xiao, L., Schoenholz, S. S., Bahri, Y., Novak, R., Sohl-Dickstein, J., and Pennington, J. Wide neural networks of any depth evolve as linear models under gradient descent. In *Advances in Neural Information Processing Systems*, 2019. *Cited on page 1, 2, 23, 24, 26.*

Lee, J., Schoenholz, S., Pennington, J., Adlam, B., Xiao, L., Novak, R., and Sohl-Dickstein, J. Finite versus infinite neural networks: an empirical study. *Advances in Neural Information Processing Systems*, 33:15156–15172, 2020. *Cited on page 24.*

Maclaurin, D., Duvenaud, D., and Adams, R. P. Autograd: Effortless gradients in numpy. In *ICML 2015 AutoML Workshop*, 2015. URL https://github.com/HIPS/autograd. *Cited on page 3.*

Martens, J. and Grosse, R. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pp. 2408–2417. PMLR, 2015. *Cited on page 25.*

Matthews, A., Hron, J., Rowland, M., Turner, R. E., and Ghahramani, Z. Gaussian process behaviour in wide deep neural networks. In *International Conference on Learning Representations*, 2018. *Cited on page 1.*

Müntz, H. Solution directe de l'équation séculaire et de quelques problèmes analogues transcendants. *C. R. Acad. Sci. Paris*, 156:43–46, 1913. *Cited on page 26.*

Naumann, U. Optimal accumulation of jacobian matrices by elimination methods on the dual computational graph. *Mathematical Programming*, 99(3):399–421, 2004. *Cited on page 2.*

Naumann, U. Optimal jacobian accumulation is np-complete. *Mathematical Programming*, 112(2):427–441, 2008. *Cited on page 2.*

Neal, R. M. Priors for infinite networks (tech. rep. no. crg-tr-94-1). *University of Toronto*, 1994. *Cited on page 1.*

Nguyen, T., Chen, Z., and Lee, J. Dataset meta-learning from kernel ridge-regression. *arXiv preprint arXiv:2011.00050*, 2020. *Cited on page 1.*

Nguyen, T., Novak, R., Xiao, L., and Lee, J. Dataset distillation with infinitely wide convolutional networks. *arXiv preprint arXiv:2107.13034*, 2021. *Cited on page 1.*

Novak, R., Xiao, L., Lee, J., Bahri, Y., Yang, G., Hron, J., Abolafia, D. A., Pennington, J., and Sohl-Dickstein, J. Bayesian deep convolutional networks with many channels are gaussian processes. In *International Conference on Learning Representations*, 2019. *Cited on page 1, 24.*

Novak, R., Xiao, L., Hron, J., Lee, J., Alemi, A. A., Sohl-Dickstein, J., and Schoenholz, S. S. Neural tangents: Fast and easy infinite neural networks in python. In *International Conference on Learning Representations*, 2020. URL https://github.com/google/neural-tangents. *Cited on page 1, 2, 8, 24, 27.*

Oreshkin, B. N., López, P. R., and Lacoste, A. Tadam: Task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, 2018. *Cited on page 2.*

Park, D. S., Lee, J., Peng, D., Cao, Y., and Sohl-Dickstein, J. Towards nngp-guided neural architecture search. *arXiv preprint arXiv:2011.06006*, 2020. *Cited on page 1, 2.*

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019. URL http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf. *Cited on page 2, 4, 26.*

Pennington, J. and Bahri, Y. Geometry of neural network loss surfaces via random matrix theory. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 2798–2806. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/pennington17a.html. *Cited on page 25.*

Radul, A., Paszke, A., Frostig, R., Johnson, M., and Maclaurin, D. You only linearize once: Tangents transpose to gradients. *arXiv preprint arXiv:2204.10923*, 2022. *Cited on page 4, 17.*

Schoenholz, S. S., Gilmer, J., Ganguli, S., and Sohl-Dickstein, J. Deep information propagation. *International Conference on Learning Representations*, 2017. *Cited on page 1.*

Spigler, S., Geiger, M., d'Ascoli, S., Sagun, L., Biroli, G., and Wyart, M. A jamming transition from under-to over-parametrization affects generalization in deep learning. *Journal of Physics A: Mathematical and Theoretical*, 52 (47):474001, 2019. *Cited on page 1.*

Steiner, A., Kolesnikov, A., Zhai, X., Wightman, R., Uszkoreit, J., and Beyer, L. How to train your vit? data, augmentation, and regularization in vision transformers. *arXiv preprint arXiv:2106.10270*, 2021. *Cited on page 8, 14.*

Tancik, M., Srinivasan, P. P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J. T., and Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020. *Cited on page 1.*

Tolstikhin, I., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., Lucic, M., and Dosovitskiy, A. Mlp-mixer: An all-mlp architecture for vision, 2021. *Cited on page 2, 8, 14.*

Xiao, L., Bahri, Y., Sohl-Dickstein, J., Schoenholz, S., and Pennington, J. Dynamical isometry and a mean field theory of CNNs: How to train 10,000-layer vanilla convolutional neural networks. In *International Conference on Machine Learning*, 2018. *Cited on page 1.*

Xiao, L., Pennington, J., and Schoenholz, S. S. Disentangling trainability and generalization in deep learning. In *International Conference on Machine Learning*, 2020. *Cited on page 1, 23.*

Yaida, S. Non-Gaussian processes and neural networks at finite widths. In *Mathematical and Scientific Machine Learning Conference*, 2020. *Cited on page 24.*

Yang, G. Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation. *arXiv preprint arXiv:1902.04760*, 2019. *Cited on page 1, 24.*

Yang, G. Tensor programs ii: Neural tangent kernel for any architecture. *arXiv preprint arXiv:2006.14548*, 2020. *Cited on page 24.*

Yang, G., Pennington, J., Rao, V., Sohl-Dickstein, J., and Schoenholz, S. S. A mean field theory of batch normalization. In *International Conference on Learning Representations*, 2019. *Cited on page 24.*

Zagoruyko, S. and Komodakis, N. Wide residual networks. In *British Machine Vision Conference*, 2016. *Cited on page 2, 8, 14.*

Zhang, H., Dauphin, Y. N., and Ma, T. Fixup initialization: Residual learning without normalization. *arXiv preprint arXiv:1901.09321*, 2019. *Cited on page 1.*

Zhou, Y., Wang, Z., Xian, J., Chen, C., and Xu, J. Meta-learning with neural tangent kernels. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=Ti87Pv5Oc8. *Cited on page 2, 26.*

Zoph, B. and Le, Q. V. Neural architecture search with reinforcement learning, 2016. URL http://arxiv.org/abs/1611.01578. *Cited on page 2.*

Appendix

A. Additional Figures

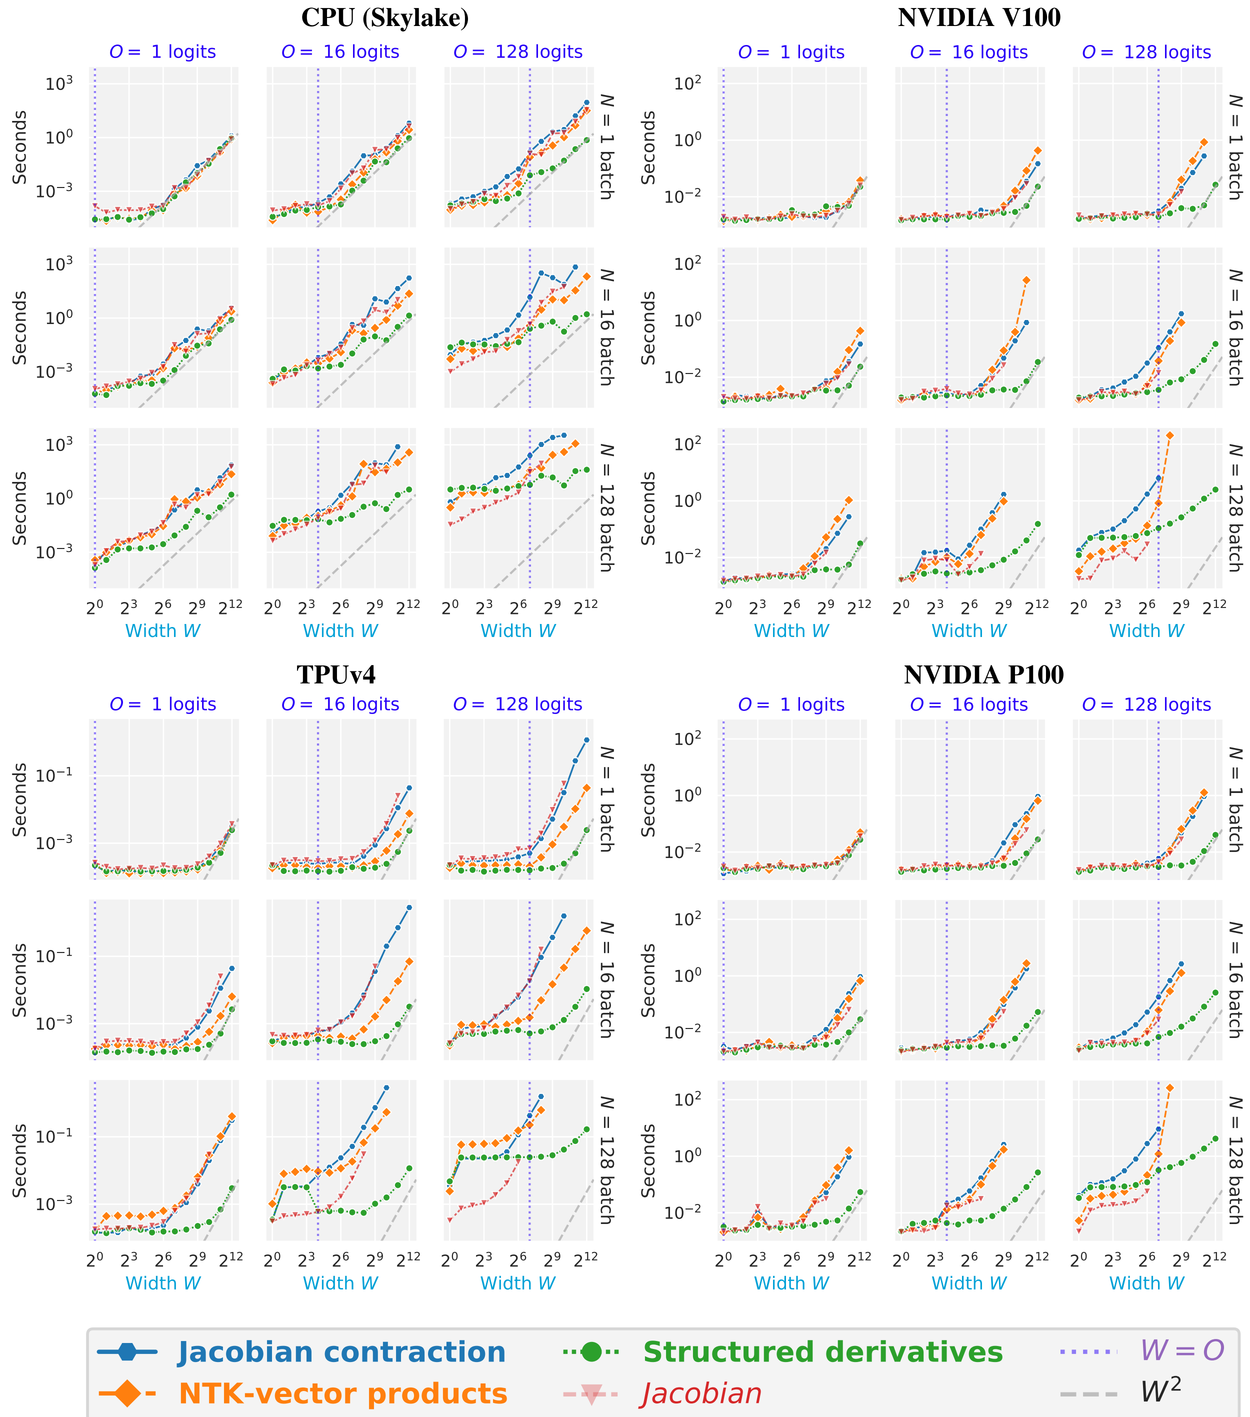

Figure 3. **Wall-clock time of computing the NTK of a 10-layer ReLU FCN on different platforms.** In all settings, Structured derivatives allow orders of magnitude improvement in wall-clock time and memory (missing points indicate out-of-memory error). However, we remark that on GPU platforms (right), NTK-vector products deliver a robust improvement only for large **O** (rightmost column), while for **O** = 16 the cost is comparable or even larger than Jacobian contraction. See Fig. 1 for FLOPs, **TPUv3** platform, and more discussion. See §N for details.

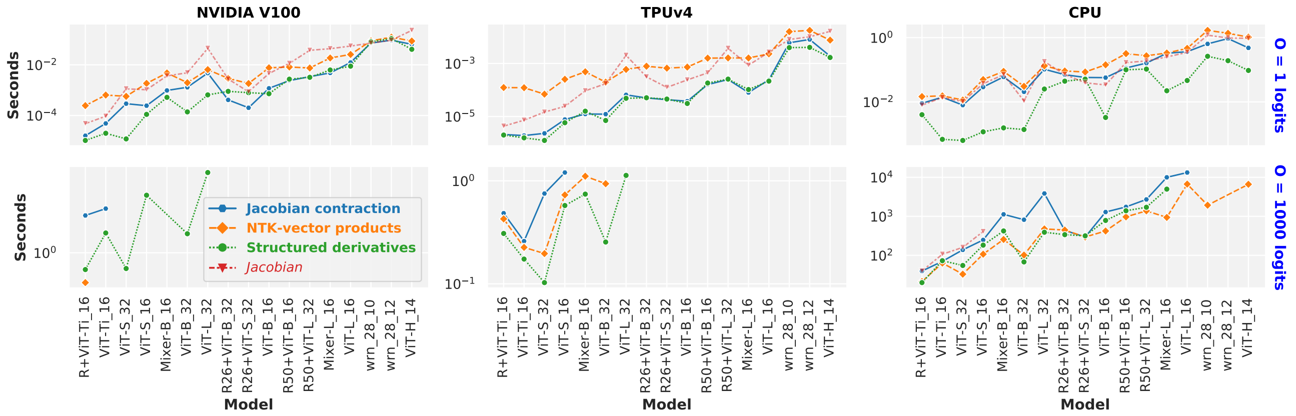

Figure 4. **Wall-clock time per input pair of computing NTK on various ImageNet models** like Vision Tansformers and hybrids (Dosovitskiy et al., 2021; Steiner et al., 2021), WideResNets (Zagoruyko & Komodakis, 2016) and MLP-Mixers (Tolstikhin et al., 2021).

Structured derivatives generally allow fastest computation, but also are able to process more models due to lower memory requirements (lower left; missing points indicate out-of-memory error). For the case of single output logit $\mathbf{O} = 1$ (top row), NTK-vector products are generally detrimental due to a costly forward pass **FP** relative to the size of parameters **P** (i.e. a lot of weight sharing; see Table 1). However, since NTK-vector products scale better than other methods with output size, for $\mathbf{O} = 1000$ (bottom row), they perform comparably or better than other methods.

Finally, we remark that the Jacobian not only runs out of memory faster, but can also take more time to compute. We conjecture that due to a larger memory footprint, XLA can sometimes perform optimizations that trade off speed for memory, and therefore compute the Jacobian in a less optimal way than if it had more memory available. Alternatively, XLA could also be performing simplifications of the NTK expression in these cases, such that those would not be possible in Jacobian computation alone.

See Fig. 2 for ResNets, and §N for details.

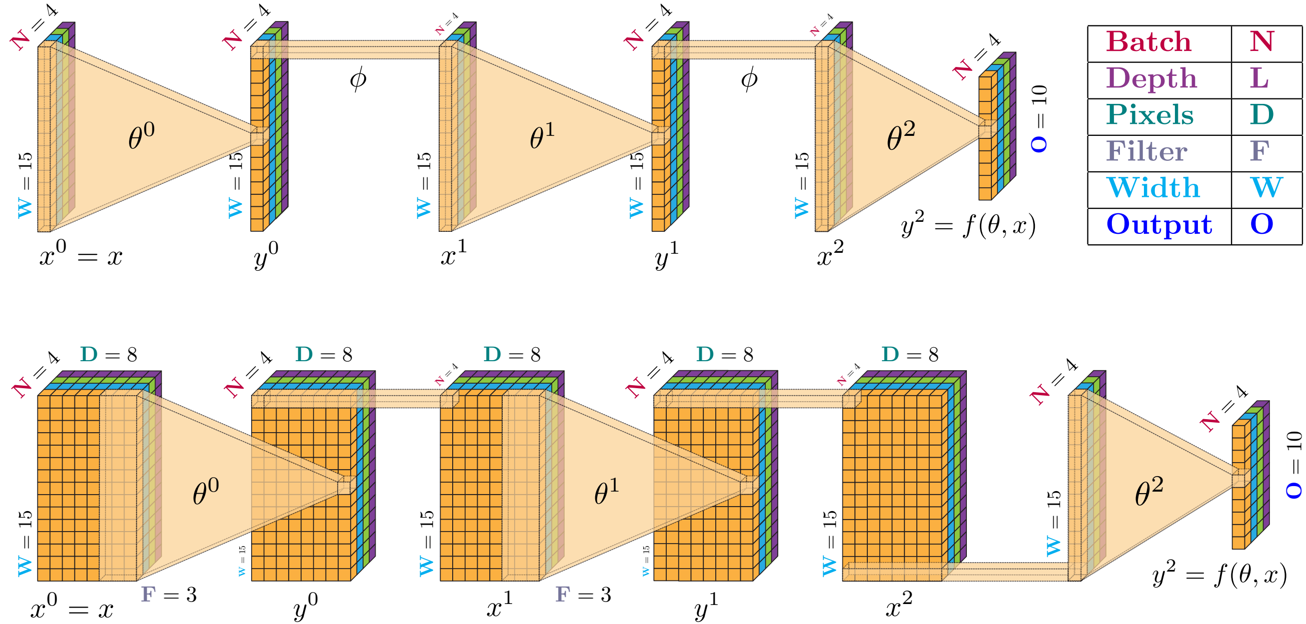

Figure 5. **Notation used in §4.1 (FCN, top) and §F (CNN, bottom).** In both settings $\mathbf{L} = 2$. For FCN, $\mathbf{K} = 5$ (3 matrix multiplication primitives and two nonlinearities ϕ), $\mathbf{D} = \mathbf{F} = 1$. For CNN, there is an extra global average pooling primitive as the penultimate layer, therefore $\mathbf{K} = 6$, and $\mathbf{D} = 8$, $\mathbf{F} = 3$.

B. Glossary

- **N** - batch size of inputs x to the NN $f(\theta, x)$.

  n - batch index ranging from 1 to **N**.

- **O** - output size (e.g. number of logits) of the NN $f(\theta, x)$ for a single (**N** $= 1$) input x.

  o - output index ranging from 1 to **O**.

- The NTK matrix has shape **NO** $\times$ **NO**.

- **W** - width of an FCN, or number of channels of a CNN. Individual inputs x are usually assumed to have the same size / number of channels.

- **L** - number of trainable parameter matrices, that are used in a possibly different number of primitives in the network. Without weight sharing, synonymous with the depth of the network.

  l - depth index ranging from 0 to **L**.

- **K** - number of primitives (nodes in the computation graph) of the network $f(\theta, x)$. Without weight sharing, synonymous with the depth of the network and is proportional to **L**.

  k - primitive index ranging from 1 to **K**.

- **D** - total number of pixels (e.g. 1024 for a 32×32 image; 1 for an FCN) in an input and every intermediate layer of a CNN (`SAME` or `CIRCULAR` padding, unit stride, and no dilation are assumed to ensure the spatial size unchanged from layer to layer).

- **F** - total filter size (e.g. 9 for a 3×3 filter; 1 for an FCN) in a convolutional filter of a CNN.

- **Y** - total output size of a primitive y (e.g. **Y** $=$ **DW** for a layer with **D** pixels and **W** channels; **Y** $=$ **W** for FCN). Depending on the context, can represent size of a single or particular primitive in the network, or the size of all primitives together.

  y - intermediate primitive output or intermediate primitive as a function of parameters $y\left(\theta^l\right)$, depending on the context.

- **C** - in §E, the size of the axis along which a primitive Jacobian $\partial y / \partial \theta$ admits certain structure (**C** can often be equal to **Y** or a significant fraction of it, e.g. **W**).

  c - index along the structured axis, ranging from 1 to **C**.

- **P** - total size of trainable parameters. Depending on the context, can represent the size of a particular weight tensor θ^l in some layer l (e.g. $\mathbf{W}^2$ for width-**W** FCN), or the size of all parameters in the network.

- **FP** - forward pass, cost (time or memory, depending on the context) of evaluating $f(\theta, x)$ on a single (**N** $= 1$) input x.

- **MJJMP** - matrix-Jacobian-Jacobian-matrix product, an AD operation necessary to evaluate the NTK as in Eq. (10), and the respective time and memory cost of the operation. As we describe in §3.4, our efficient implementation of MJJMPs often allows to evaluate the NTK much faster than when using standard AD operations like JVPs and VJPs. Note that unlike other variables, **MJJMP** represents the batched (accounting for **N**) contraction cost, since batched **MJJMP** can have non-trivial dependence on **N**.

C. List of Primitives and their Structures

Transposable primitive in `jax.ad.primitive_transposes`	**Constant block-diagonal**	**Block-diagonal**	**Output block-tiled**
add	✓		✓
add_any	✓		✓
all_gather			
all_to_all			
broadcast_in_dim	✓		✓
call			
closed_call			
complex			
concatenate	✓		
conj			
conv_general_dilated	✓	✓	
convert_element_type	✓		
copy	✓		
cumsum			
custom_lin			
custom_linear_solve			
custom_transpose_call			
device_put	✓		
div	✓	✓	
dot_general	✓	✓	
dynamic_slice			
dynamic_update_slice			
fft			
gather			
imag			
linear_call			
mul	✓	✓	
named_call			
neg	✓		
pad	✓		
pdot			
ppermute			
psum			
real			
reduce_sum	✓		
reduce_window_sum	✓		
remat_call			
reshape	✓		
rev	✓		
scatter			
scatter-add			
scatter-mul			
select_n			
select_and_gather_add			
select_and_scatter_add			
sharding_constraint			
slice			
squeeze	✓		
sub	✓		✓
transpose	✓		
triangular_solve			
while			
xla_call			
xla_pmap			
xmap			
zeros_like	✓		

Table 4. **List of all linear primitives and currently implemented Structured derivatives rules from §E.** In the future, more primitives and more rules can be supported, yet at the time of writing even the small set currently covered enables dramatic speed-up and memory savings in contemporary ImageNet models as in Fig. 2 and Fig. 4.

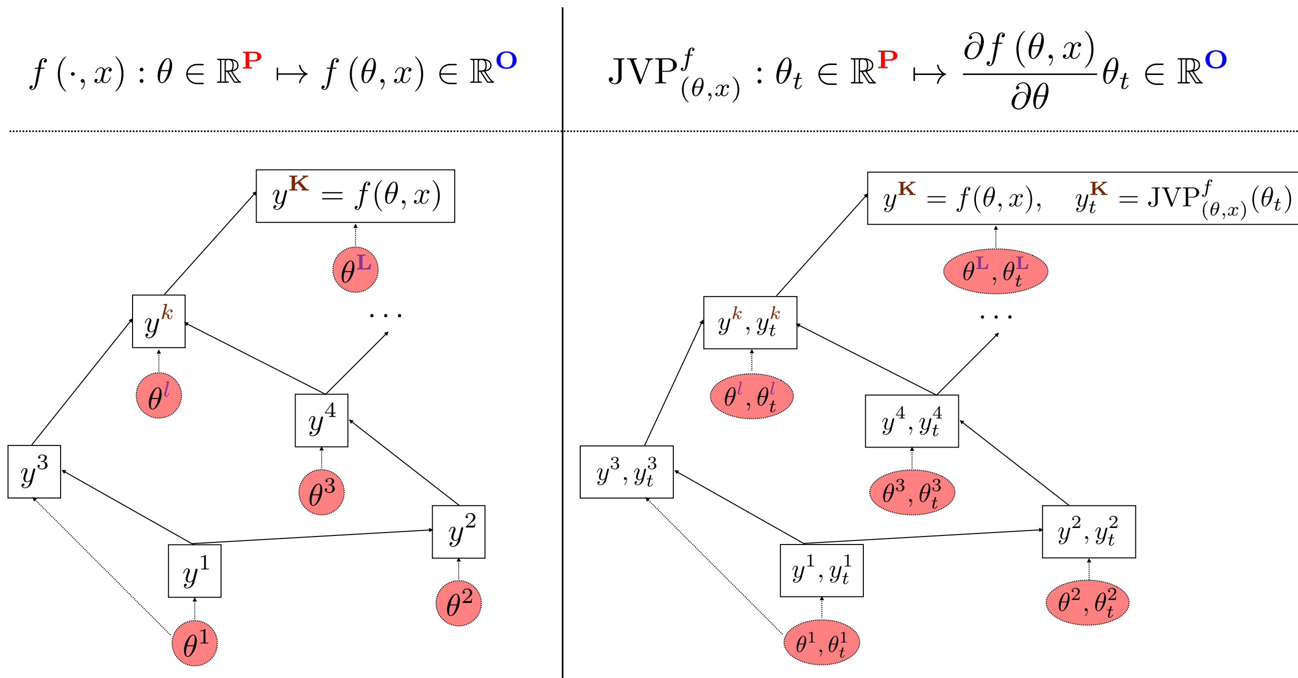

Figure 6. **Visual demonstration for why JVP time and memory costs are asymptotically comparable to the forward pass (FP). Left:** computational graph of the forward pass $f(\theta, x)$. **Right:** computational graph of joint evaluation of the forward pass $f(\theta, x)$ along with $\mathrm{JVP}^{f}_{(\theta,x)}(\theta_t)$. Each node of the JVP graph accepts both primal and tangent inputs, and returns primal and tangent outputs, but the topology of the graph and order of execution remains identical to **FP**. As long as individual nodes of the JVP graph do not differ significantly in time and memory from the **FP** nodes, time and memory of a JVP ends up asymptotically equivalent to **FP** due to identical graph structure. However, in order to create JVP nodes and evaluate them, the time cost does grow by a factor of about 3 compared to **FP**. See §D for discussion.

D. JVP and VJP Costs

Here we provide intuition for why JVP and VJP are asymptotically equivalent in time to the forward pass **FP**, as we mentioned in §3.1.2. See (Griewank & Walther, 2008, Section 3) and (Radul et al., 2022) for a rigorous treatment, and the JAX Autodiff Cookbook for a hands-on introduction.

JVP can be computed by traversing a computational graph of the same topology as **FP**, except for primitive nodes in the graph need to be augmented to compute not only the forward pass of the node, but also the JVP of the node (see Fig. 6). Due to identical topology and order of evaluation, asymptotically time and memory costs remain unchanged. However, constructing the augmented nodes in the JVP graph, and their consequent evaluation results in extra time cost proportional to the size of the graph. Therefore in practice JVP costs about $3 \times$ **FP** time and $2 \times$ **FP** memory.

VJP, as a linear function of cotangents f_c, is precisely the transpose of the linear function JVP. As such, it can be computed by traversing the transpose of the JVP graph (Fig. 6, right), with each JVP node replaced by its transposition as well. This results in identical time and memory costs, as long as node transpositions are implemented efficiently. However, their evaluation requires primal outputs y^k (now inputs to the transpose nodes), which is why VJP necessitates an extra **FP** time cost to compute them (hence costlier than JVP, but still inconsequential asymptotically) and extra memory to store them, which can generally increase asymptotic memory requirements.

E. Types of Structured Derivatives and Associated MJJMP Costs

Here we continue §3.4 and list the types of structures in primitive Jacobians $\partial y / \partial \theta$ that allow linear algebra simplifications of the respective MJJMPs in Eq. (10). Resulting contraction (**MJJMP**) and memory (**J**) costs from the following subsections are summarized in Table 5, and the list of primitives annotated with their types of structure is presented in Table 4.

Structure of $\partial y/\partial\theta \downarrow$	MJJMP (time; minimum of the 3 values)			NJ (memory)
	Outside-in	Left-to-right	Inside-out	
None w/ VJPs & JVPs	$\mathbf{NO}$ [FP] $+\mathbf{N^2O^2P}$	$\mathbf{N^2O}$ [FP]	Not possible	0
None w/ explicit matrices	$\mathbf{NOYP} + \mathbf{N^2O^2P}$	$\mathbf{N^2OYP} + \mathbf{N^2O^2Y}$	$\mathbf{N^2Y^2P} + \mathbf{N^2OY^2} + \mathbf{N^2O^2Y}$	$\mathbf{NYP}$
Block-diagonal	$\mathbf{NOYP/C} + \mathbf{N^2O^2P}$	$\mathbf{N^2OYP/C} + \mathbf{N^2O^2Y}$	$\mathbf{N^2Y^2P/C^2} + \mathbf{N^2OY^2/C} + \mathbf{N^2O^2Y}$	$\mathbf{NYP/C^2}$
Constant block-diagonal	$\mathbf{NOYP/C} + \mathbf{N^2O^2P}$	$\mathbf{N^2OYP/C} + \mathbf{N^2O^2Y}$	$\mathbf{N^2Y^2P/C^3} + \mathbf{N^2OY^2/C} + \mathbf{N^2O^2Y}$	$\mathbf{NYP/C^2}$
Input block-tiled	$\mathbf{NOYP/C} + \mathbf{N^2O^2P}$	$\mathbf{N^2OYP/C} + \mathbf{N^2O^2Y}$	$\mathbf{N^2Y^2P/C} + \mathbf{N^2OY^2} + \mathbf{N^2O^2Y}$	$\mathbf{NYP/C}$
Output block-tiled	$\mathbf{NOYP/C} + \mathbf{N^2O^2P} + \mathbf{NOY}$	$\mathbf{N^2OYP/C} + \mathbf{N^2O^2Y/C} + \mathbf{NOY}$	$\mathbf{N^2Y^2P/C^2} + \mathbf{N^2OY^2/C^2} + \mathbf{N^2O^2Y/C} + \mathbf{NOY}$	$\mathbf{NYP/C}$
Block-tiled	$\mathbf{NOYP/C^2} + \mathbf{N^2O^2P/C} + \mathbf{NOY}$	$\mathbf{N^2OYP/C^2} + \mathbf{N^2O^2Y/C^2} + \mathbf{NOY}$	$\mathbf{N^2Y^2P/C^3} + \mathbf{N^2OY^2/C^2} + \mathbf{N^2O^2Y/C} + \mathbf{NOY}$	$\mathbf{NYP/C^2}$

Table 5. **Asymptotic time (MJJMP) and extra memory (NJ) costs of computing the contractions for NTK summands $\Theta^f_\theta[l, k_1, k_2]$ of shape $\mathbf{NO}\times\mathbf{NO}$ in §3.4 and Table 1.** Time complexity of the MJJMP in Structured derivatives is the minimum (due to using `np.einsum` with optimal contraction order) of the 3 time entries in the row corresponding to the structure present in a pair of primitives $y_1^{k_1}$ and $y_2^{k_2}$. How it compares to Jacobian contraction and NTK-vector products (top row) depends on the specific primitive, notably the cost of evaluating the primitive **FP**. See Table 2 and Table 3 for exact comparison in the case of matrix multiplication and convolution. See §B for legend.

E.1. No Structure

We first consider the default cost of evaluating a single summand in Eq. (10), denoting individual matrix shapes underneath:

$$\Theta^f_\theta[l, k_1, k_2] = \frac{\partial f_1}{\partial y_1^{k_1}}\frac{\partial y_1^{k_1}}{\partial\theta^l}\frac{\partial y_2^{k_2}}{\partial\theta^l}^T\frac{\partial f_2}{\partial y_2^{k_2}}^T =: \overbrace{\underbrace{\frac{\partial f_1}{\partial y_1}}_{\mathbf{O}\times\mathbf{Y}}\underbrace{\frac{\partial y_1}{\partial\theta}}_{\mathbf{Y}\times\mathbf{P}}\underbrace{\frac{\partial y_2}{\partial\theta}^T}_{\mathbf{P}\times\mathbf{Y}}\underbrace{\frac{\partial f_2}{\partial y_2}^T}_{\mathbf{Y}\times\mathbf{O}}}^{\mathbf{O}\times\mathbf{O}} \tag{18}$$

We have dropped indices l, k_1 and k_2 on the right-hand side of Eq. (18) to avoid clutter, and consider $\theta := \theta^l, y_1 := y_1^{k_1}, y_2 := y_2^{k_2}$ until the end of this section. There are 3 ways of contracting Eq. (18) that cost

(a) **Outside-in: $\mathbf{OYP} + \mathbf{O^2P}$**

(b) **Left-to-right and right-to-left: $\mathbf{OYP} + \mathbf{O^2Y}$.**

(c) **Inside-out-left and inside-out-right: $\mathbf{Y^2P} + \mathbf{OY^2} + \mathbf{O^2Y}$.**

In the next sections, we look at how these costs are reduced given certain structure in $\partial y/\partial\theta$.

E.2. Block Diagonal

Assume $\partial y/\partial\theta = \oplus_{c=1}^{\mathbf{C}}\partial y^c/\partial\theta_c$, where $\oplus$ stands for direct sum of matrices, i.e. $\partial y/\partial\theta$ is a block diagonal matrix made of blocks $\{\partial y^c/\partial\theta_c\}_{c=1}^{\mathbf{C}}$, where $\partial y^c/\partial\theta_c$ have shapes $(\mathbf{Y}/\mathbf{C})\times(\mathbf{P}/\mathbf{C})$. Here $\{y^c\}_{c=1}^{\mathbf{C}}$ and $\{\theta_c\}_{c=1}^{\mathbf{C}}$ are partitions of y and θ respectively. In NNs this structure is present in binary bilinear operations (on θ and another argument) such as multiplication, division, batched matrix multiplication, or depthwise convolution. Then Eq. (18) can be re-written as

$$\Theta^f_\theta[l, k_1, k_2] = \frac{\partial f_1}{\partial y_1}\frac{\partial y_1}{\partial\theta}\frac{\partial y_2}{\partial\theta}^T\frac{\partial f_2}{\partial y_2}^T \tag{19}$$

$$= \frac{\partial f_1}{\partial y_1}\left(\oplus_{c=1}^{\mathbf{C}}\frac{\partial y_1^c}{\partial\theta_c}\right)\left(\oplus_{c=1}^{\mathbf{C}}\frac{\partial y_2^c}{\partial\theta_c}\right)^T\frac{\partial f_2}{\partial y_2}^T \tag{20}$$

$$= \frac{\partial f_1}{\partial y_1}\left(\oplus_{c=1}^{\mathbf{C}}\left[\frac{\partial y_1^c}{\partial\theta_c}\frac{\partial y_2^c}{\partial\theta_c}^T\right]\right)\frac{\partial f_2}{\partial y_2}^T \tag{21}$$

$$= \sum_{c=1}^{\mathbf{C}}\frac{\partial f_1}{\partial y_1^c}\left[\frac{\partial y_1^c}{\partial\theta_c}\frac{\partial y_2^c}{\partial\theta_c}^T\right]\frac{\partial f_2}{\partial y_2^c}^T, \tag{22}$$

where we have applied the block matrix identity

$$\left[A^1, \dots, A^{\mathbf{C}}\right]^T\left(\oplus_{c=1}^{\mathbf{C}}B^c\right)\left[D^1, \dots, D^{\mathbf{C}}\right] = \sum_{c=1}^{\mathbf{C}}A^cB^cD^c. \tag{23}$$

We now perform a complexity analysis similar to Eq. (18):

$$\Theta_\theta^f[l, k_1, k_2] = \sum_{c=1}^{\mathbf{C}} \overbrace{\underbrace{\frac{\partial f_1}{\partial y_1^c}}_{\mathbf{O}\times(\mathbf{Y}/\mathbf{C})} \underbrace{\frac{\partial y_1^c}{\partial \theta_c}}_{(\mathbf{Y}/\mathbf{C})\times(\mathbf{P}/\mathbf{C})} \underbrace{\frac{\partial y_2^c}{\partial \theta_c}^T}_{(\mathbf{P}/\mathbf{C})\times(\mathbf{Y}/\mathbf{C})} \underbrace{\frac{\partial f_2}{\partial y_2^c}^T}_{(\mathbf{Y}/\mathbf{C})\times\mathbf{O}}}^{\mathbf{O}\times\mathbf{O}}$$

In this case complexities of the three methods become

1. **Outside-in:** $\mathbf{OYP}/\mathbf{C} + \mathbf{O}^2\mathbf{P}$.
2. **Left-to-right and right-to-left:** $\mathbf{OYP}/\mathbf{C} + \mathbf{O}^2\mathbf{Y}$.
3. **Inside-out-left and inside-out-right:** $\mathbf{Y}^2\mathbf{P}/\mathbf{C}^2 + \mathbf{OY}^2/\mathbf{C} + \mathbf{O}^2\mathbf{Y}$.

E.3. Constant-block Diagonal

Assume $\frac{\partial y}{\partial \theta} = I_{\mathbf{C}} \otimes \frac{\partial y^1}{\partial \theta_1}$, and $\frac{\partial y^1}{\partial \theta_1}$ has shape $(\mathbf{Y}/\mathbf{C}) \times (\mathbf{P}/\mathbf{C})$. In NNs, this is present in fully-connected, convolutional, locally-connected, attention, and many other layers that contain a matrix multiplication along some axis. This is also present in all unary elementwise linear operations on θ like transposition, negation, reshaping and many others. This is a special case of §E.2 with $\frac{\partial y^c}{\partial \theta_c} = \frac{\partial y^1}{\partial \theta_1}$ for any c. Here a similar analysis applies, yielding

$$\Theta_\theta^f[l, k_1, k_2] = \sum_{c=1}^{\mathbf{C}} \overbrace{\underbrace{\frac{\partial f_1}{\partial y_1^c}}_{\mathbf{O}\times(\mathbf{Y}/\mathbf{C})} \underbrace{\frac{\partial y_1^1}{\partial \theta_1}}_{(\mathbf{Y}/\mathbf{C})\times(\mathbf{P}/\mathbf{C})} \underbrace{\frac{\partial y_2^1}{\partial \theta_1}^T}_{(\mathbf{P}/\mathbf{C})\times(\mathbf{Y}/\mathbf{C})} \underbrace{\frac{\partial f_2}{\partial y_2^c}^T}_{(\mathbf{Y}/\mathbf{C})\times\mathbf{O}}}^{\mathbf{O}\times\mathbf{O}}$$

and the same contraction complexities as in §E.2, except for the **Inside-out** order, where the inner contraction term costs only $\mathbf{Y}^2\mathbf{P}/\mathbf{C}^3$, since it is only contracted once instead of $\mathbf{C}$ times as in the Block-diagonal case.

E.4. Input Block-tiled

Assume $\frac{\partial y}{\partial \theta} = \mathbb{1}_{(1,\mathbf{C})} \otimes \frac{\partial y}{\partial \theta_1}$, where $\mathbb{1}_{(1,\mathbf{C})}$ is an all-ones matrix of shape $1 \times \mathbf{C}$, and $\frac{\partial y}{\partial \theta_1}$ has shape $\mathbf{Y} \times (\mathbf{P}/\mathbf{C})$. This occurs, for example, in summation or taking the mean.

$$\Theta_\theta^f[l, k_1, k_2] = \frac{\partial f_1}{\partial y_1}\frac{\partial y_1}{\partial \theta}\frac{\partial y_2}{\partial \theta}^T\frac{\partial f_2}{\partial y_2}^T \tag{24}$$

$$= \frac{\partial f_1}{\partial y_1}\left(\mathbb{1}_{(1,\mathbf{C})} \otimes \frac{\partial y_1}{\partial \theta_1}\right)\left(\mathbb{1}_{(1,\mathbf{C})} \otimes \frac{\partial y_2}{\partial \theta_1}\right)^T\frac{\partial f_2}{\partial y_2}^T \tag{25}$$

$$= \frac{\partial f_1}{\partial y_1}\left(\mathbf{C}\mathbb{1}_{(1,1)} \otimes \left[\frac{\partial y_1}{\partial \theta_1}\frac{\partial y_2}{\partial \theta_1}^T\right]\right)\frac{\partial f_2}{\partial y_2}^T \tag{26}$$

$$= \mathbf{C}\frac{\partial f_1}{\partial y_1}\left[\frac{\partial y_1}{\partial \theta_1}\frac{\partial y_2}{\partial \theta_1}^T\right]\frac{\partial f_2}{\partial y_2}^T. \tag{27}$$

The matrix shapes are

$$\Theta_\theta^f[l, k_1, k_2] = \mathbf{C}\overbrace{\underbrace{\frac{\partial f_1}{\partial y_1}}_{\mathbf{O}\times\mathbf{Y}} \underbrace{\frac{\partial y_1}{\partial \theta_1}}_{\mathbf{Y}\times(\mathbf{P}/\mathbf{C})} \underbrace{\frac{\partial y_2}{\partial \theta_1}^T}_{(\mathbf{P}/\mathbf{C})\times\mathbf{Y}} \underbrace{\frac{\partial f_2}{\partial y_2}^T}_{\mathbf{Y}\times\mathbf{O}}}^{\mathbf{O}\times\mathbf{O}}$$

Which leads to the following resulting complexities:

1. **Outside-in:** $\mathbf{OYP}/\mathbf{C} + \mathbf{O}^2\mathbf{P}$.
2. **Left-to-right and right-to-left:** $\mathbf{OYP}/\mathbf{C} + \mathbf{O}^2\mathbf{Y}$.
3. **Inside-out and inside-out-right:** $\mathbf{Y}^2\mathbf{P}/\mathbf{C} + \mathbf{OY}^2 + \mathbf{O}^2\mathbf{Y}$.

E.5. Output Block-tiled

Assume $\frac{\partial y}{\partial \theta} = \mathbb{1}_{(\mathbf{C},1)} \otimes \frac{\partial y^1}{\partial \theta}$, where $\frac{\partial y^1}{\partial \theta}$ has shape $(\mathbf{Y}/\mathbf{C}) \times \mathbf{P}$. This occurs during broadcasting or broadcasted arithmetic operations. In this case

$$\Theta_\theta^f[l, k_1, k_2] = \frac{\partial f_1}{\partial y_1}\frac{\partial y_1}{\partial \theta}\frac{\partial y_2}{\partial \theta}^T\frac{\partial f_2}{\partial y_2}^T \tag{28}$$

$$= \frac{\partial f_1}{\partial y_1}\left(\mathbb{1}_{(\mathbf{C},1)} \otimes \frac{\partial y_1^1}{\partial \theta}\right)\left(\mathbb{1}_{(\mathbf{C},1)} \otimes \frac{\partial y_2^1}{\partial \theta}\right)^T\frac{\partial f_2}{\partial y_2}^T \tag{29}$$

$$= \frac{\partial f_1}{\partial y_1}\left(\mathbb{1}_{(\mathbf{C},\mathbf{C})} \otimes \left[\frac{\partial y_1^1}{\partial \theta}\frac{\partial y_2^1}{\partial \theta}^T\right]\right)\frac{\partial f_2}{\partial y_2}^T \tag{30}$$

$$= \left(\sum_{c=1}^{\mathbf{C}}\frac{\partial f_1}{\partial y_1^c}\right)\left[\frac{\partial y_1^1}{\partial \theta_1}\frac{\partial y_2^1}{\partial \theta_1}^T\right]\left(\sum_{c=1}^{\mathbf{C}}\frac{\partial f_2}{\partial y_2^c}^T\right), \tag{31}$$

where we have used a block matrix identity

$$\left[A^1, \ldots, A^{\mathbf{C}}\right]^T\left(\mathbb{1}_{(\mathbf{C},\mathbf{C})} \otimes B\right)\left[D^1, \ldots, D^{\mathbf{C}}\right] = \left(\sum_{c=1}^{\mathbf{C}} A^c\right) B \left(\sum_{c=1}^{\mathbf{C}} D^c\right).$$

Finally, denoting the shapes,

$$\Theta_\theta^f[l, k_1, k_2] = \overbrace{\underbrace{\left(\sum_{c=1}^{\mathbf{C}}\frac{\partial f_1}{\partial y_1^c}\right)}_{\mathbf{O}\times(\mathbf{Y}/\mathbf{C})}\underbrace{\frac{\partial y_1^1}{\partial \theta}}_{(\mathbf{Y}/\mathbf{C})\times\mathbf{P}}\underbrace{\frac{\partial y_2^1}{\partial \theta}^T}_{\mathbf{P}\times(\mathbf{Y}/\mathbf{C})}\underbrace{\left(\sum_{c=1}^{\mathbf{C}}\frac{\partial f_2}{\partial y_2^c}^T\right)}_{(\mathbf{Y}/\mathbf{C})\times\mathbf{O}}}^{\mathbf{O}\times\mathbf{O}},$$

complexities of the three methods become (notice we add an $\mathbf{OY}$ term to perform the sums)

1. **Outside-in:** $\mathbf{OYP}/\mathbf{C} + \mathbf{O}^2\mathbf{P} + \mathbf{OY}$.
2. **Left-to-right:** $\mathbf{OYP}/\mathbf{C} + \mathbf{O}^2\mathbf{Y}/\mathbf{C} + \mathbf{OY}$.
3. **Inside-out:** $\mathbf{Y}^2\mathbf{P}/\mathbf{C}^2 + \mathbf{OY}^2/\mathbf{C}^2 + \mathbf{O}^2\mathbf{Y}/\mathbf{C} + \mathbf{OY}$.

E.6. Block-tiled

Assume $\frac{\partial y}{\partial \theta} = \mathbb{1}_{(\mathbf{C},\mathbf{C})} \otimes \frac{\partial y^1}{\partial \theta_1}$, where $\frac{\partial y^1}{\partial \theta_1}$ has shape $(\mathbf{Y}/\mathbf{C}) \times (\mathbf{P}/\mathbf{C})$. This occurs for instance when y is a constant. In this case

$$\Theta_\theta^f[l, k_1, k_2] = \frac{\partial f_1}{\partial y_1}\frac{\partial y_1}{\partial \theta}\frac{\partial y_2}{\partial \theta}^T\frac{\partial f_2}{\partial y_2}^T \tag{32}$$

$$= \frac{\partial f_1}{\partial y_1}\left(\mathbb{1}_{(\mathbf{C},\mathbf{C})} \otimes \frac{\partial y_1^1}{\partial \theta_1}\right)\left(\mathbb{1}_{(\mathbf{C},\mathbf{C})} \otimes \frac{\partial y_2^1}{\partial \theta_1}\right)^T\frac{\partial f_2}{\partial y_2}^T \tag{33}$$

$$= \frac{\partial f_1}{\partial y_1}\left(\mathbf{C}\mathbb{1}_{(\mathbf{C},\mathbf{C})} \otimes \left[\frac{\partial y_1^1}{\partial \theta_1}\frac{\partial y_2^1}{\partial \theta_1}^T\right]\right)\frac{\partial f_2}{\partial y_2}^T \tag{34}$$

$$= \mathbf{C}\left(\sum_{c=1}^{\mathbf{C}}\frac{\partial f_1}{\partial y_1^c}\right)\left[\frac{\partial y_1^1}{\partial \theta_1}\frac{\partial y_2^1}{\partial \theta_1}^T\right]\left(\sum_{c=1}^{\mathbf{C}}\frac{\partial f_2}{\partial y_2^c}^T\right), \tag{35}$$

This results in the following contraction:

$$\Theta_\theta^f[l, k_1, k_2] = \mathbf{C}\overbrace{\underbrace{\left(\sum_{c=1}^{\mathbf{C}}\frac{\partial f_1}{\partial y_1^c}\right)}_{\mathbf{O}\times(\mathbf{Y}/\mathbf{C})}\underbrace{\frac{\partial y_1^1}{\partial \theta_1}}_{(\mathbf{Y}/\mathbf{C})\times(\mathbf{P}/\mathbf{C})}\underbrace{\frac{\partial y_2^1}{\partial \theta_1}^T}_{(\mathbf{P}/\mathbf{C})\times(\mathbf{Y}/\mathbf{C})}\underbrace{\left(\sum_{c=1}^{\mathbf{C}}\frac{\partial f_2}{\partial y_2^c}^T\right)}_{(\mathbf{Y}/\mathbf{C})\times\mathbf{O}}}^{\mathbf{O}\times\mathbf{O}},$$

Structure of $\partial y / \partial \theta \downarrow$	MJJMP (time; minimum of the 3 values)			NJ (memory)
	Outside-in	**Left-to-right**	**Inside-out**	
None w/ JVPs and VJPs	$\mathbf{N}^2\mathbf{O}^2\mathbf{W}^2$	$\mathbf{N}^2\mathbf{O}\mathbf{W}^2$	Not possible	0
Constant block-diagonal	$\mathbf{N}^2\mathbf{O}^2\mathbf{W}^2$	$\mathbf{N}^2\mathbf{O}\mathbf{W}^2 + \mathbf{N}^2\mathbf{O}^2\mathbf{W}$	$\mathbf{N}^2\mathbf{O}^2\mathbf{W}$	$\mathbf{N}\mathbf{W}$

Table 6. **Asymptotic time complexities of computing a single FCN layer NTK contribution**. See §E.9 for discussion, Table 5 for a more general setting, Table 2 for the case of deep networks, and §B for detailed legend.

with final complexities of

1. **Outside-in:** $\mathbf{OYP}/\mathbf{C}^2 + \mathbf{O}^2\mathbf{P} + \mathbf{OY}$.
2. **Left-to-right:** $\mathbf{OYP}/\mathbf{C}^2 + \mathbf{O}^2\mathbf{Y}/\mathbf{C}^2 + \mathbf{OY}$.
3. **Inside-out:** $\mathbf{Y}^2\mathbf{P}/\mathbf{C}^3 + \mathbf{OY}^2/\mathbf{C}^2 + \mathbf{O}^2\mathbf{Y}/\mathbf{C} + \mathbf{OY}$.

E.7. Batched NTK Cost Analysis

For simplicity, above we have considered evaluating the NTK for $\mathbf{N} = 1$. In the more general case of inputs of batch sizes $\mathbf{N}_1$ and $\mathbf{N}_2$, the same argument as in previous section follows, but the cost of contractions involving terms from different batches grow by a multiplicative factor of $\mathbf{N}_1\mathbf{N}_2$, while all other costs grow by a factor of $\mathbf{N}_1$ or $\mathbf{N}_2$. To declutter notation we consider $\mathbf{N}_1 = \mathbf{N}_2 = \mathbf{N}$, and summarize resulting costs in Table 5.

E.8. Complex Structure Cost Analysis

In previous sections we have considered $\partial y_1 / \partial \theta$ and $\partial y_2 / \partial \theta$ admitting the same, and at most one kind of structure. While this is a common case, in general these Jacobians may admit multiple types of structures along multiple axes (for instance, addition is Constant block-diagonal along non-broadcasted axes, and Output block-tiled along the broadcasted axes), and $\partial y_1 / \partial \theta$ and $\partial y_2 / \partial \theta$ may have different types of structures and respective axes, if the same weight θ is used in multiple different primitives of different kind. In such cases, equivalent optimizations are possible (and are implemented in the code) along the largest common subsets of axes for each type of structure that $\partial y_1 / \partial \theta$ and $\partial y_2 / \partial \theta$ have.

For example, let θ be a matrix in $\mathbb{R}^{\mathbf{W}\times\mathbf{W}}$, y_1 be multiplication by a scalar $y_1(\theta) = 2\theta$, and y_2 be matrix-vector multiplication $y_2(\theta) = \theta x$, $x \in \mathbb{R}^{\mathbf{W}}$. In this case $\partial y_1 / \partial \theta = 2I_{\mathbf{W}} \otimes I_{\mathbf{W}}$, i.e. it is Constant block-diagonal along axes 1 and 2. $\partial y_2 / \partial \theta = I_{\mathbf{W}} \otimes x^T$, i.e. it is also Constant block-diagonal, but only along axis 1. Hence, the NTK term containing $\partial y_1 / \partial \theta$ and $\partial y_2 / \partial \theta$ will be computed with Constant block-diagonal simplification along axis 1. There are probably more computationally optimal ways of processing different structure combinations, as well as more types of structures to be leveraged for NTK computation, and we intend to investigate it in future work.

E.9. Example: FCN Layer

In §4.1 we have derived the **J** and **MJJMP** costs for an FCN in an improvised manner. Here we arrive at identical complexities using the framework in §E and Table 5. Precisely, per Table 4, matrix-vector multiplication (`dot_general`) has Constant block-diagonal structure with $\mathbf{C} = \mathbf{Y} = \mathbf{W}$. Substituting it (along with $\mathbf{P} = \mathbf{W}^2$, $\mathbf{FP} = \mathbf{W}^2$), we obtain Table 6, which reproduces our conclusions in §4.1 and Table 2 for a single layer $l < \mathbf{L}$.

E.10. Example: CNN Layer

Per Table 4 convolution (`conv_general_dilated`) also has Constant block-diagonal structure along the channel axis with $\mathbf{C} = \mathbf{W}$. Substituting it (along with $\mathbf{P} = \mathbf{FW}^2$, $\mathbf{FP} = \mathbf{DFW}^2$, $\mathbf{Y} = \mathbf{DW}$), we obtain Table 7, which we next use in §F.

F. CNN NTK Complexity

Here we derive the CNN NTK complexity for §4.2, identically to §4.1. We assume the same setup, except for denoting the total filter size as $\mathbf{F}$, total number of pixels $\mathbf{D}$ (unchanged from layer to layer, assuming `SAME` or `CIRCULAR` padding and

Structure of $\partial y / \partial \theta \downarrow$	**MJJMP** (time; minimum of the 3 values)			**NJ** (memory)
	Outside-in	**Left-to-right**	**Inside-out**	
None w/ JVPs and VJPs	$\mathbf{N}^2\mathbf{O}^2\mathbf{F}\mathbf{W}^2 + \mathbf{NODFW}^2$	$\mathbf{N}^2\mathbf{ODFW}^2$	Not possible	0
Constant block-diagonal	$\mathbf{N}^2\mathbf{O}^2\mathbf{F}\mathbf{W}^2 + \mathbf{NODFW}^2$	$\mathbf{N}^2\mathbf{ODFW}^2 + \mathbf{N}^2\mathbf{O}^2\mathbf{DW}$	$\mathbf{N}^2\mathbf{D}^2\mathbf{FW} + \mathbf{N}^2\mathbf{OD}^2\mathbf{W} + \mathbf{N}^2\mathbf{O}^2\mathbf{DW}$	$\mathbf{NDFW}$

Table 7. **Asymptotic time complexities of computing a single CNN layer NTK contribution**. See §E.10 for discussion, Table 5 for a more general setting, Table 3 for the case of deep networks, and §B for detailed legend.

unit stride, no dilation), and adding an extra global average pooling layer before the final readout layer (Fig. 5, bottom). Note that FCN (§4.1) is a particular case of CNN with $\mathbf{D} = \mathbf{F} = 1$.

In this setting $\mathbf{K} = 2\mathbf{L} + 2$ (one more than in §4.1, due to the extra global average pooling primitive). $\mathbf{P}^l = \mathbf{FW}^2$ for $l < \mathbf{L}$ and $\mathbf{OW}$ for $l = \mathbf{L}$, $\mathbf{P} = \mathbf{LFW}^2 + \mathbf{OW}$, $\mathbf{Y}^k = \mathbf{DW}$ for $k < \mathbf{K} - 1$, $\mathbf{W}$ for $k = \mathbf{K} - 1$, and $\mathbf{O}$ for $k = \mathbf{K}$; the total primitive output size is $\mathbf{Y} \sim \mathbf{LDW} + \mathbf{O}$. Finally, the forward pass costs $\mathbf{FP} \sim \mathbf{LDFW}^2 + \mathbf{OW}$ time and $\mathbf{DW} + \mathbf{FW}^2 + \mathbf{OW}$ memory.

As in §4.1 we plug the above into the cost of the Jacobian contraction in §3.2, and obtain

> CNN Jacobian contraction costs $\mathbf{N}^2\mathbf{O}^2\mathbf{LFW}^2 + \mathbf{N}^2\mathbf{O}^3\mathbf{W} + \mathbf{NOLDFW}^2$ time; $\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOFW}^2 + \mathbf{NLDW} + \mathbf{NO}^2\mathbf{W} + \mathbf{LFW}^2$ memory.

Analogously, the cost of NTK-vector products from §3.3 becomes

> CNN NTK-vector products cost $\mathbf{N}^2\mathbf{OLDFW}^2 + \mathbf{N}^2\mathbf{O}^2\mathbf{W}$ time; $\mathbf{N}^2\mathbf{O}^2 + \mathbf{NOFW}^2 + \mathbf{NLDW} + \mathbf{NO}^2\mathbf{W} + \mathbf{LFW}^2$ memory.

For Structured derivatives (§3.4), we remark that **MJJMP** and $\mathbf{J}$ costs for $l = \mathbf{L}$ are identical to FCN (§4.1, §E.9), hence $\mathbf{N}^2\mathbf{O}^3 + \mathbf{N}^2\mathbf{W}$ time and $\mathbf{N}^2\mathbf{O}^2 + \mathbf{NW}$ memory.

For convolutional layers $l < \mathbf{L}$, we reference §E.10 and specifically Table 7, to obtain $\mathbf{J} = \mathbf{DFW}$ for both memory and time, since the primitive Jacobian subarrays can be computed via $\mathbf{F}$ calls of `np.roll` on the $\mathbf{D} \times \mathbf{W}$ input x. However, note that on actual hardware such operations are known to perform *worse* than highly optimized convolution-based solutions that scale as $\mathbf{DF}^2\mathbf{W}$ (this would be performing $\mathbf{F}$ JVP calls of the convolutional primitive, each costing $\mathbf{DFW}$ time, or, equivalently, invoking `jax.lax.conv_general_dilated_patches`). For this reason we currently use the convolutional approach in our implementation, but note that it would be trivial to extend the codebase to switch to the `np.roll`-based solution in the case of very large filters $\mathbf{F} > \mathbf{OW}$.

Finally, from Table 7, $\mathbf{MJJMP} = \mathbf{N}^2\mathbf{O}^2 \min\left(\mathbf{FW}^2, \mathbf{DW} + \frac{\mathbf{DFW}^2}{\mathbf{O}}, \mathbf{DW} + \frac{\mathbf{D}^2\mathbf{W}}{\mathbf{O}} + \frac{\mathbf{D}^2\mathbf{FW}}{\mathbf{O}^2}\right)$ (we have ignored the $\mathbf{NODFW}^2$ term in **Outside-in** contraction order since this is dominated by the cost $\mathbf{NO}\,[\mathbf{FP}]$ of computing primitive output cotangents).

Adding all the costs up, we obtain

> CNN Structured derivatives cost
> $\mathbf{NOLDFW}^2 + \mathbf{NO}^2\mathbf{W} + \mathbf{N}^2\mathbf{O}^3 + \mathbf{N}^2\mathbf{O}^2\mathbf{L}\min\left(\mathbf{FW}^2, \mathbf{DW} + \frac{\mathbf{DFW}^2}{\mathbf{O}}, \mathbf{DW} + \frac{\mathbf{D}^2\mathbf{W}}{\mathbf{O}} + \frac{\mathbf{D}^2\mathbf{FW}}{\mathbf{O}^2}\right)$ time;
> $\mathbf{N}^2\mathbf{O}^2 + \mathbf{NODW} + \mathbf{NDFW} + \mathbf{NLDW} + \mathbf{LFW}^2 + \mathbf{OW}^2$ memory.

G. Jacobian Rules for Structured Derivatives

Here we discuss computing primitive $\partial y / \partial \theta$ Jacobians as part of our Structured derivatives implementation in §6. We provide 4 options to compute them through arguments `_j_rules` and `_fwd`:

1. **Forward mode**, `_fwd = True`, is equivalent to `jax.jacfwd`, forward mode Jacobian computation, performed by applying the JVP to $\mathbf{P}$ columns of the $I_\mathbf{P}$ identity matrix. Best for $\mathbf{P} < \mathbf{Y}$.
2. **Reverse mode**, `_fwd = False`, is equivalent to `jax.jacrev`, reverse mode Jacobian computation, performed by applying the VJP to $\mathbf{Y}$ columns of the $I_\mathbf{Y}$ identity matrix. Best for $\mathbf{P} > \mathbf{Y}$.

3. **Automatic mode**, `_fwd = None`, selects forward or reverse mode for each y based on parameters and output shapes.

4. **Rule mode**, `_j_rules = True`, queries a dictionary of Jacobian rules (similar to the dictionary of structure rules) with our custom implementations of primitive Jacobians, instead of computing them through VJPs or JVPs. The reason for introducing custom rules follows our discussion in §3.4: while JAX has computationally optimal VJP and JVP rules, the respective Jacobian computations are not guaranteed to be most efficient. In practice, we find our rules to be most often faster, however this effect is not perfectly consistent (can occasionally be slower) and often negligible, requiring further investigation.

The default setting is `_j_rules = True`, `_fwd = None`, i.e. a custom Jacobian implementation is preferred, and, if absent, Jacobian is computed in forward or reverse mode based on parameters and output sizes. Note that in all settings, structure of $\partial y/\partial\theta$ is used to compute only the smallest Jacobian subarray necessary, and therefore most often inputs to VJP/JVP will be smaller identity matrices $I_{\mathbf{P}/\mathbf{C}}$ or $I_{\mathbf{Y}/\mathbf{C}}$ respectively, and all methods will return a smaller Jacobian matrix of size $(\mathbf{Y}/\mathbf{C}) \times (\mathbf{P}/\mathbf{C})$. If for any reason (for example debugging) you want the whole $\partial y/\partial\theta$ Jacobians computed, you can set the `_s_rules=False`, i.e. disable structure rules.

H. Known Issues

We will continue improving our function transformations in various ways after release, and welcome bug reports and feature requests. Below are the missing features / issues at the time of submission:

1. No support for complex differentiation.

2. Not tested on functions with advanced JAX primitives like parallel collectives (`psum`, `pmean`, etc.), gradient checkpointing (`remat`), compiled loops (`scan`; Python loops are supported).

3. Our current implementation of NTK-vector products relies on XLA's common subexpression elimination (CSE) in order to reuse computation across different pairs of inputs x_1 and x_2, and, as shown in Fig. 1 and Fig. 3, can have somewhat unpredictable wall-clock time performance and memory requirements. We believe this could correspond to CSE not always working perfectly, and are looking into a more explicitly efficient implementation.

I. Finite and Infinite Width NTK

In this work we focus on the finite width NTK $\Theta_\theta^f(x_1, x_2)$, defined in Eq. (1), repeated below with a batch size $\mathbf{N}$:

$$\textbf{F-NTK (finite width):} \quad \underbrace{\Theta_\theta^f(x_1, x_2)}_{\mathbf{NO}\times\mathbf{NO}} \coloneqq \underbrace{\left[\partial f(\theta, x_1)/\partial\theta\right]}_{\mathbf{NO}\times\mathbf{P}} \underbrace{\left[\partial f(\theta, x_2)/\partial\theta\right]^T}_{\mathbf{P}\times\mathbf{NO}}. \tag{36}$$

Another important object in deep learning is the *infinite width* NTK $\Theta^f(x_1, x_2)$, introduced by Jacot et al. (2018):

$$\textbf{I-NTK (infinite width):} \quad \underbrace{\Theta^f(x_1, x_2)}_{\mathbf{NO}\times\mathbf{NO}} \coloneqq \lim_{\mathbf{W}\to\infty} \mathbb{E}_{\theta\sim\mathcal{N}(\mathbf{0}, I_{\mathbf{P}})}\left[\underbrace{\Theta_\theta^f(x_1, x_2)}_{\mathbf{NO}\times\mathbf{NO}}\right]. \tag{37}$$

A natural question to ask is what are the similarities and differences of F- and I-NTK, when is one more applicable than the other, and what are their implementation and compute costs.

Applications. At a high level, F-NTK describes the local/linearized behavior of the finite width NN $f(\theta, x)$ (Lee et al., 2019). In contrast, I-NTK is an approximation that is exact only in the infinite width $\mathbf{W}$ limit, and only at initialization ($\theta \sim \mathcal{N}(\mathbf{0}, I_{\mathbf{P}})$). As such, the resulting I-NTK has no notion of width $\mathbf{W}$, parameters θ, and cannot be computed during draining, or in a transfer or meta-learning setting, where the parameters θ are updated. As a consequence, any application to finite width networks (§2) is better served by the F-NTK, and often impossible with the I-NTK.

In contrast, I-NTK describes the behavior of an infinite ensemble of infinitely wide NNs. In certain settings this can be desirable, such as when studying the inductive bias of certain NN architectures (Xiao et al., 2020) or uncertainty (Adlam et al., 2020), marginalizing away the dependence on specific parameters θ. However, care should be taken when applying

I-NTK findings to the finite width realm, since many works have demonstrated substantial finite width effects that cannot be captured by the I-NTK (Novak et al., 2019; Arora et al., 2019b; Lee et al., 2019; Yaida, 2020; Hanin & Nica, 2020; Lee et al., 2020).

Mathematical scope. Another significant difference between F- and I-NTK is the scope of their definitions in Eq. (36) and Eq. (37) and mathematical tractability.

The F-NTK is well-defined for any differentiable (w.r.t. θ) function f, and so are our algorithms. In fact, our implementations support any Tangent Kernels (not necessarily "Neural"), and are not specific to NNs at all.

In contrast, the I-NTK requires the function f to have the concept of width $\mathbf{W}$ (that can be meaningfully taken to infinity) to begin with, and further requires f and θ to satisfy many conditions in order for the I-NTK to be well-defined (Yang, 2019). In order for I-NTK to be well-defined *and computable in closed-form*, f needs to be built out of a relatively small, hand-selected number of primitives that admit certain Gaussian integrals to have closed-form solutions. Examples of ubiquitous primitives that *don't* allow a closed-form solution include attention with standard parameterization (Hron et al., 2020b); max-pooling; sigmoid, (log-)softmax, tanh, and many other nonlinearities; various kinds of normalization (Yang et al., 2019); non-trivial weight sharing (Yang, 2020); and many other settings. Going forward, it is unlikely that the I-NTK will scale to the enormous variety of architectures introduced by the research community each year.

Implementation tractability. Above we have demonstrated that the I-NTK is defined for a very small subset of functions admitting the F-NTK. A closed-form solution exists for an even smaller subset. However, even when the I-NTK admits a closed-form solution, it is important to consider the complexity of implementing it.

Our implementation for computing the F-NTK is applicable to any differentiable function f, and requires no extra effort when switching to a different function g. It is similar to JAX's highly-generic function transformations such as `jax.vmap`.

In contrast, there is no known way to compute the I-NTK for an arbitrary function f, even if the I-NTK exists in closed form. The best existing solution to date is provided by Novak et al. (2020), which allows to *construct* f out of the limited set of building blocks provided by the authors. However, one cannot compute the I-NTK for a function implemented in a different library such as Flax (Heek et al., 2020), or Haiku (Hennigan et al., 2020), or bare-bone JAX. One would have to re-implement it using the primitives provided by Novak et al. (2020). Further, for a generic architecture, the primitive set is unlikely to be sufficient, and the function will need to be adapted to admit a closed-form I-NTK.

Computational tractability. F-NTK and I-NTK have different time and memory complexities, and a fully general comparison is an interesting direction for future work. Here we provide discussion for deep FCNs and CNNs.

Networks having a fully-connected top ($\mathbf{L}$) readout layer have a constant-block diagonal I-NTK, hence its cost *does not* scale with $\mathbf{O}$. The cost of computing the I-NTK for a deep FCN scales as $\mathbf{N}^2\mathbf{L}$ for time and $\mathbf{N}^2$ for memory. A deep CNN without pooling costs $\mathbf{N}^2\mathbf{D}\mathbf{L}$ time and $\mathbf{N}^2\mathbf{D}$ memory (where $\mathbf{D}$ is the total number of pixels in a single input/activation; $\mathbf{D} = 1$ for FCNs). Finally, a deep CNN with pooling, or any other generic architecture that leverages the spatial structure of inputs/activations, costs $\mathbf{N}^2\mathbf{D}^2\mathbf{L}$ time and $\mathbf{N}^2\mathbf{D}^2$ memory. This applies to all models in Fig. 2 and Fig. 4, Graph Neural Networks (Du et al., 2019), and the vast majority of other architectures used in practice.

The quadratic scaling of the I-NTK cost with $\mathbf{D}$ is especially burdensome, since, for example, for ImageNet $\mathbf{D}^2 = 224^4 = 2,517,630,976$. As a result, it would be impossible to evaluate the I-NTK on even a single ($\mathbf{N} = 1$) pair of inputs with a V100 GPU for any model for which we've successfully evaluated the F-NTK in Fig. 2 and Fig. 4.

The F-NTK time and memory only scale linearly with $\mathbf{D}$ (Table 3). However, the F-NTK cost scales with other parameters such as width $\mathbf{W}$ or number of outputs $\mathbf{O}$, and in general the relative F- and I-NTK performance will depend on these parameters. As a rough point of comparison, we consider the cost of evaluating the I-NTK of a 20-layer binary classification ReLU CNN with pooling on a V100 GPU used by Arora et al. (2019b) against the respective F-NTK with $\mathbf{W} = 128$ also used by Arora et al. (2019b, Section B). Arora et al. (2019b) and Novak et al. (2020) report from 0.002 to 0.003 seconds per I-NTK entry on a pair of CIFAR-10 inputs. Using Structured derivatives, we can compute the respective F-NTK entry on same hardware in at most 0.000014 seconds, i.e. at least 100 times faster than the I-NTK. In 0.002 – 0.003 seconds per NTK entry, we can compute the F-NTK on a pair of *ImageNet* inputs (about 50x larger than CIFAR-10) for a *200-layer ResNet* (about 10x deeper than the model above) in Fig. 2 (top left).

Finally, we remark that efficient NTK-vector products without instantiating the entire $\mathbf{NO} \times \mathbf{NO}$ NTK are only possible using the F-NTK (§L).

J. Relationship Between the NTK and the Hessian

Here we briefly touch on the difference between the NTK

$$\textbf{NTK:} \quad \underbrace{\Theta_\theta^f(x_1, x_2)}_{\mathbf{NO}\times\mathbf{NO}} := \underbrace{\left[\partial f(\theta, x_1)/\partial\theta\right]}_{\mathbf{NO}\times\mathbf{P}} \underbrace{\left[\partial f(\theta, x_2)/\partial\theta\right]^T}_{\mathbf{P}\times\mathbf{NO}}, \tag{38}$$

and the Hessian:

$$\textbf{Hessian:} \quad \underbrace{\mathbf{H}_\theta(x)}_{\mathbf{P}\times\mathbf{P}} := \frac{\partial^2 \mathcal{L}\left(f\left(\theta, x\right)\right)}{\partial\theta^2}, \tag{39}$$

defined for some differentiable loss function on the output space $\mathcal{L}\colon \mathbb{R}^{\mathbf{NO}} \to \mathbb{R}$.

Both matrices characterize localized training dynamics of a NN, and the NTK can be used as a more tractable quantity in cases where the Hessian is infeasible to instantiate (for example, **P** amounts to tens of millions in models considered in Fig. 2 and Fig. 4).

The connection between the NTK and the Hessian can be established for when $\mathcal{L}$ is the squared error (SE), i.e. $\mathcal{L}(y) = \|y - \mathcal{Y}\|_2^2/2$, where $\mathcal{Y} \in \mathbb{R}^{\mathbf{NO}}$ are the training targets. In this case, as presented in (Pennington & Bahri, 2017, Section 2) and (Grosse, 2021, Equation 13; page 21):

$$\underbrace{\mathbf{H}_\theta(x)}_{\mathbf{P}\times\mathbf{P}} = \underbrace{\left[\frac{\partial f(\theta, x)}{\partial\theta}^T \frac{\partial f(\theta, x)}{\partial\theta}\right]}_{:=\mathbf{H}_\theta^0(x)} + \underbrace{\sum_{n,o=1}^{\mathbf{N},\mathbf{O}} \left(f\left(\theta, x\right) - \mathcal{Y}\right)^{n,o} \frac{\partial^2 f\left(\theta, x\right)^{n,o}}{\partial^2\theta}}_{:=\mathbf{H}_\theta^1(x)}, \tag{40}$$

where we have decomposed the Hessian $\mathbf{H}_\theta(x)$ into two summands $\mathbf{H}_\theta^0(x)$ and $\mathbf{H}_\theta^1(x)$ following the notation of Pennington & Bahri (2017).

Notice that if $f\left(\theta, x\right) = \mathcal{Y}$, i.e. the SE loss is 0, $\mathbf{H}_\theta^1(x) = 0$, yielding

$$\underbrace{\mathbf{H}_\theta(x)}_{\mathbf{P}\times\mathbf{P}} = \mathbf{H}_\theta^0(x) = \underbrace{\frac{\partial f(\theta, x)}{\partial\theta}^T}_{\mathbf{P}\times\mathbf{NO}} \underbrace{\frac{\partial f(\theta, x)}{\partial\theta}}_{\mathbf{NO}\times\mathbf{P}}, \qquad \underbrace{\Theta_\theta^f(x, x)}_{\mathbf{NO}\times\mathbf{NO}} = \underbrace{\frac{\partial f(\theta, x)}{\partial\theta}}_{\mathbf{NO}\times\mathbf{P}} \underbrace{\frac{\partial f(\theta, x_2)}{\partial\theta}^T}_{\mathbf{P}\times\mathbf{NO}}, \tag{41}$$

and, as a consequence, the Hessian and the NTK have the same eigenvalues (see also Grosse (2021, Page 21)) in this particular case. Moreover, the Hessian (and Hessian-vector products) can be computed very similarly to NTK-vector products, by switching the order of VJP and JVP operations in Eq. (7).

However, except for zero SE loss case above, the NTK and the Hessian have different spectra, and their computations share less similarity. Precisely, Hessian-vector products (and consequently the Hessian) are computed in JAX through a composition of JVPs and VJPs similar to NTK-vector products:

$$\mathbf{H}_\theta(x)v = \left[\frac{\partial^2 \mathcal{L}\left(f\left(\theta, x\right)\right)}{\partial\theta^2}\right] v = \frac{\partial}{\partial\theta}\left[\frac{\partial \mathcal{L}\left(f\left(\theta, x\right)\right)}{\partial\theta}\right] v = \frac{\partial\left[\mathrm{VJP}_{(\theta,x)}^{\mathcal{L}\circ f}(1)\right]}{\partial\theta} v = \mathrm{JVP}_{(\theta,x)}^{\left[\mathrm{VJP}_{(\cdot,\cdot)}^{\mathcal{L}\circ f}(1)\right]}(v)\,. \tag{42}$$

Although Eq. (42) is similar to Eq. (7) in that both are compositions of JVPs and VJPs, in Eq. (7) the result of a VJP is the input tangent to the JVP of f, while in Eq. (42) it is the function to be differentiated by the JVP (instead of f).

K. Relationship Between Structured derivatives and K-FAC

Similarly to K-FAC (Martens & Grosse, 2015), in the simple example of §3.4, we leverage the structure in the matrix-vector product derivative w.r.t. the matrix, and we use the mixed-product property, i.e. $(A \otimes B)(C \otimes D) = (AC) \otimes (BD)$. However, in the general case this is not enough, and Structured derivatives rely on three components: (1) a direct sum linear algebra Eq. (23), (2) symbolic simplification of expressions with an identity matrix (§E.3), and (3) optimal order of contractions in Eq. (10) (e.g. "Inside-out" (Table 6), which is not possible to achieve with standard AD tools). All three components are necessary to achieve our asymptotic complexities, and cannot be achieved by leveraging the mixed-product property alone.

L. Applications with a Limited Compute Budget

While our methods allow to dramatically speed-up the computation of the NTK, all of them still scale as $\mathbf{N}^2\mathbf{O}^2$ for both time and memory, which can be intractable for large datasets and/or large outputs.

Here we present several settings in which our proposed methods still provide substantial time and memory savings, even when instantiating the entire $\mathbf{NO} \times \mathbf{NO}$ NTK is not feasible or not necessary.

- **NTK-vector products**. In many applications one only requires computing the NTK-vector product linear map

$$\Theta_\theta^f : v \in \mathbb{R}^{\mathbf{NO}} \mapsto \Theta_\theta^f v \in \mathbf{NO}, \tag{43}$$

  without computing the entire NTK matrix Θ_θ. A common setting is using the power iteration method (Müntz, 1913) to compute the NTK condition number and hence trainability of the respective NN (Lee et al., 2019; Chen et al., 2021a;b). Another setting is using conjugate gradients to compute $\Theta_\theta^{-1}\mathcal{Y}$ when doing kernel ridge regression with the NTK (Jacot et al., 2018; Lee et al., 2019; Zhou et al., 2021).

  Eq. (43) is the same map as the one we considered in §3.3, and naturally, NTK-vector products can provide a substantial speed-up over Jacobian contraction in this setting. Precisely, a straightforward application of Jacobian contraction yields

$$\underbrace{\Theta_\theta^f v}_{\mathbf{NO}\times 1} = \underbrace{\frac{\partial f(\theta, x_1)}{\partial \theta}}_{\mathbf{NO}\times\mathbf{P}} \underbrace{\frac{\partial f(\theta, x_2)}{\partial \theta}^T}_{\mathbf{P}\times\mathbf{NO}} \underbrace{v}_{\mathbf{NO}\times 1}. \tag{44}$$

  Combined with the cost of computing the weight space cotangents $\partial f/\partial\theta$, such evaluation costs $\mathbf{NO}$ $[\mathbf{FP}]$ time, i.e. the cost of instantiating the entire Jacobian. Alternatively, one could store the entire Jacobians of sizes $\mathbf{NOP}$ in memory, and compute a single NTK-vector product in $\mathbf{NOP}$ time.

  In contrast, NTK-vector products allow to compute an NTK-vector product at a cost asymptotically equivalent to a single VJP call (§3.3), i.e. $\mathbf{N}$ $[\mathbf{FP}]$, $\mathbf{O}$ times faster than Jacobian contraction without caching. With caching, fastest method will vary based on the cost of $[\mathbf{FP}]$ relative to $\mathbf{OP}$, as discussed in §3.3, but NTK-vector products will remain substantially more memory-efficient due to not caching the entire $\mathbf{NOP}$ Jacobians.

- **Batching.** In many applications it suffices to compute the NTK over small batches of the data. For example Dauphin & Schoenholz (2019); Chen et al. (2021a;b) estimate the conditioning by computing an approximation to the NTK on $\mathbf{N}$ equal to 128, 32, and 48 examples respectively. Similarly, Zhou et al. (2021) use a small batch size of $\mathbf{N} = 25$ to meta-learn the network parameters by replacing the inner SGD training loop with NTK regression.

- **Pseudo-NTK.** Many applications (§2) compute a pseudo-NTK of size $\mathbf{N} \times \mathbf{N}$, which is commonly equal to one of its $\mathbf{O}$ diagonal blocks, or to the mean of all $\mathbf{O}$ blocks. The reason for considering such approximation is that in the infinite width limit, off-diagonal entries often converge to zero, and for wide-enough networks this approximation can be justified. Compute-wise, these approximations are equivalent to having $\mathbf{O} = 1$. While an important contribution of our work is to enable computing the full $\mathbf{NO} \times \mathbf{NO}$ NTK quickly, if necessary, Structured derivatives can be combined with the $\mathbf{O} = 1$ approximations, and still provide an asymptotic speed-up and memory savings relative to prior works.

M. Leveraging JAX Design for Efficient NTK Computation

At the time of writing, Tensorflow (Abadi et al., 2016, TF) and PyTorch (Paszke et al., 2019) are more widely used than JAX (Bradbury et al., 2018). However, certain JAX features and design choices made it much more suitable, if not indispensable, for our project:

1. Structured derivatives require manual implementation of structure rules for different primitives in the computational graph of a function $f(\theta, x)$. JAX has a small primitive set of about 136 primitives, while PyTorch and Tensorflow have more than 400 (Frostig et al., 2021, Section 2). Further, by leveraging `jax.linearize`, we reduce our task to implementing structure rules for only *linear* primitives, of which JAX has only 56.[9] To our knowledge neither PyTorch nor Tensorflow have an equivalent transformation, which makes JAX a natural choice due to the very concise set of primitives that we need to handle (Table 4).

[9] See §3.4, as well as (Frostig et al., 2021, Section 1) for how JAX uses the same insight to not implement all 136 VJP rules, but only implement 56 transpose rules for reverse mode AD.

2. NTK-vector products critically rely on forward mode AD (JVP), and Structured derivatives also use it (albeit it's not crucial; see §G). At the time of writing, PyTorch does not implement an efficient forward mode AD.

3. Structured derivatives rely crucially on the ability to traverse the computation graph to rewrite contractions using our substitution rules. JAX provides a highly-convenient graph representation in the form of a Jaxpr, as well as tooling and documentation for writing custom Jaxpr interpreters.

4. All implementations (even Jacobian contraction) rely heavily on `jax.vmap` (and in many cases, it is indispensable). While PyTorch has released a prototype of `vmap` in May 2021, it was not available when we started this project.

M.1. Interfacing with Tensorflow

Since JAX and TF leverage the same underlying compiler XLA, we are able to construct a seamless TF $\to$ JAX $\to$ TF pipeline using Jax2TF and TF2Jax (Babuschkin et al., 2020). We will provide a prototype implementation of this pipeline in the `experimental` folder.

M.2. Interfacing with PyTorch

Recently introduced Functorch (Horace He, 2021) enables implementing Jacobian contraction and NTK-vector products, and our code has been ported to PyTorch in the "Neural Tangent Kernels" tutorial. However, due to JAX features (1) and (3) from §M implementing Structured derivatives in PyTorch remains challenging. If necessary, a pipeline PyTorch $\to$ TF $\to$ JAX $\to$ PyTorch using ONNX (Bai et al., 2019) and DLPack can be constructed. We will include an example implementation in the `experimental` folder.

N. Experimental Details

All experiments were performed in JAX (Bradbury et al., 2018) using 32-bit precision.

Throughout this work we assume the cost of multiplying two matrices of shapes (M, K) and (K, P) to be MKP. While there are faster algorithms for very large matrices, the XLA compiler (used by JAX, among other libraries) does not implement them, so our assumption is accurate in practice.

Hardware. **CPU** experiments were run on Dual 28-core Intel Skylake CPUs with at least 240 GiB of RAM. **NVIDIA V100** and **NVIDIA P100** used a respective GPU with 16 GiB GPU RAM. **TPUv3** and **TPUv4** have 8 and 32 GiB of RAM respectively, and use the default 16/32-bit mixed precision.

Fig. 1 and Fig. 3: a 10-layer, ReLU FCN was constructed with the Neural Tangents (Novak et al., 2020) `nt.stax` API. Default settings (weight variance 1, no bias) were used. Individual inputs x had size 3. Jacobian contraction was evaluated using `nt.empirical_ntk_fn` with `trace_axes=(), diagonal_axes=(), vmap_axes=0`. Jacobian was evaluated using `jax.jacobian` with a `vmap` over inputs x. For time measurements, all functions were `jax.jit` ted, and timing was measured as the average of 100 random samples (compilation time was not included). For FLOPs, the function was not JITted, and FLOPs were measured on CPU using the `utils.get_flops` function that is released together with our code.[10]

Fig. 2 and Fig. 4: for ResNets, implementations from Flax (Heek et al., 2020) were used, specifically `flax.examples.imagenet.models`. For WideResNets, the code sample from Novak et al. (2020) was used.[11] For all other models, we used implementations from `https://github.com/google-research/vision_transformer`. Inputs were random arrays of shapes $224 \times 224 \times 3$. All models were JITted. All reported values are averages over 10 random samples. For each setting, we ran a grid search over the batch size $\mathbf{N}$ in $\left\{2^k\right\}_{k=0}^{9}$, and reported the best time divided by $\mathbf{N}^2$, i.e. best possible throughput in each setting.

Title page ribbon is adapted from Arfian (2022).

[10] The XLA team has let us know that if JITted, the FLOPs are currently correctly computed only on TPU, but are incorrect on other platforms. Therefore we compute FLOPs of non-JITted functions.

[11] We replaced `stax.AvgPool((8, 8)), stax.Flatten()` with `stax.GlobalAvgPool()`.