
Finding trainable sparse networks through Neural Tangent Transfer

Tianlin Liu^{1,2} Friedemann Zenke¹

Abstract

Deep neural networks have dramatically transformed machine learning, but their memory and energy demands are substantial. The requirements of real biological neural networks are rather modest in comparison, and one feature that might underlie this austerity is their sparse connectivity. In deep learning, trainable sparse networks that perform well on a specific task are usually constructed using label-dependent pruning criteria. In this article, we introduce Neural Tangent Transfer, a method that instead finds trainable sparse networks in a label-free manner. Specifically, we find sparse networks whose training dynamics, as characterized by the neural tangent kernel, mimic those of dense networks in function space. Finally, we evaluate our label-agnostic approach on several standard classification tasks and show that the resulting sparse networks achieve higher classification performance while converging faster.

1. Introduction

Deep neural networks achieve human-level performance in a variety of domains (LeCun et al., 2015; Schmidhuber, 2015; Silver et al., 2017). Unlike biological neural networks, however, deep learning systems require extensive computational resources and energy. This demand poses a major impediment for future applications in which deep networks are embedded in always-on edge devices and in the Internet of Things (IoT) (Neftci, 2018; Roy et al., 2019; LeCun, 2019).

One approach to make artificial neural networks more energy-efficient is to exploit sparseness at both the activity and connectivity level (Gong et al., 2014; Lebedev et al., 2015; Jaderberg et al., 2014). This approach echoes the sparseness principle of biological neural systems, which underlies their superior energy efficiency (Sterling & Laughlin,

2017). To this date, most of the studies on sparse artificial neural networks have focused on *post hoc* pruning, whereby sparse networks are derived from *trained* dense networks. This approach, however, is computationally costly because it does not allow us to benefit from sparseness during training.

To tackle this issue, recent work (Lee et al., 2019b; Wang et al., 2020) proposed identifying trainable sparse networks by pruning dense networks *before* training them. Such *foresight pruning* results in sparse networks that can learn quickly and generalize well in subsequent supervised learning tasks despite having only a few adjustable, nonzero parameters. Although they perform well in many scenarios, existing foresight pruning methods can be improved in two ways. First, current methods solely rely on labeled data to determine which parameters to prune. But since labeled data are often scarce (Xu et al., 2019; Rigollet & Weed, 2019), label-free pruning methods that leverage abundant unlabeled data could offer a decisive advantage. Second, existing foresight pruning approaches have focused on global pruning criteria, which are known to remove weights primarily from the fully connected layers of convolutional neural networks (CNNs) while preserving most parameters in the convolutional layers. Yet, because convolutional layers are responsible for most of the computational burden (Yang et al., 2017), global pruning results in comparatively small savings. Here, pruning methods that effectively sparsify convolutional filters in a layerwise manner could offer substantial performance improvements.

In this article, we introduce neural tangent transfer (NTT), a foresight pruning method that works well in the layerwise setting and without labeled data. To the best of our knowledge, this is the first work to find trainable sparse networks in a *label-free* manner. Importantly, we show that our method reliably finds trainable CNNs with sparse convolutional filters, a situation in which existing foresight pruning methods struggle.

2. Prior work

The problem of instantiating sparse neural networks has been considered in several previous studies. They can be broadly divided into three categories. First, *post hoc pruning* approaches rely on removing redundant parameters from trained, dense networks. Second, *on-the-fly pruning* meth-

¹Friedrich Miescher Institute for Biomedical Research, Basel, Switzerland. ²University of Basel, Basel, Switzerland. Correspondence to: Friedemann Zenke <friedemann.zenke@fmi.ch>.

ods enforce sparsity constraints during supervised training. Finally, *foresight pruning* refers to directly finding a sparse network from scratch that can be trained later to high accuracy.

Post hoc pruning approaches operate on trained neural networks and attempt to remove network parameters or units that contribute marginally to task performance using different pruning criteria. Established pruning criteria include parameter magnitude (Hanson & Pratt, 1989; Ström, 1997) and the Hessian of the loss function with respect to the weights (LeCun et al., 1990; Hassibi et al., 1994). Finally, it is common practice to fine-tune or iteratively re-train pruned models, to improve their final performance (Han et al., 2015; Guo et al., 2016; Zhu & Gupta, 2018).

On-the-fly pruning is an alternative approach in which network sparsity is enforced during supervised training. This can be achieved in several ways. First, by introducing a sparsity-inducing term in the penalty function of supervised training (Chauvin, 1989; Collins & Kohli, 2014; Molchanov et al., 2017; Carreira-Perpinan & Idelbayev, 2018; Louizos et al., 2018). Such penalty terms encourage parameters to be close to zero. Second, by use of dynamic network rewiring rules over training time that keeps network sparsity below a given threshold (Mocanu et al., 2018; Bellec et al., 2018; Yan et al., 2019; Evci et al., 2020).

Foresight pruning refers an approach in which one first prunes a network at its initialization and then trains the pruned network to convergence. In part, foresight pruning is motivated by evidence that specific sparse networks can be trained to yield comparable performance to the corresponding dense model (Liu et al., 2019; Frankle & Carbin, 2019; Morcos et al., 2019). To that end, Lee et al. (2019b; 2020) attempted to identify trainable sparse network structure at the network initialization stage based on the connection sensitivity criterion. Wang et al. (2020) proposed to find sparse networks that preserve the error gradients after pruning.

3. Neural Tangent Transfer framework

All previous approaches have relied on labeled data to find trainable sparse neural networks. In this article, we develop a label-free approach to find sparse networks whose training evolution in the function space are comparable to those of their corresponding dense counterparts. The resulting sparse networks exhibit higher performance when trained on subsequent supervised tasks. We approach this problem by considering the evolution of a neural network’s output from a given starting point, determined by its random initialization, to an endpoint, characterized by the trained network parameters. To yield good performance, the endpoint of the sparse network’s output evolution needs to lie close to the output of a well-performing network. However, in su-

pervised learning, the endpoint itself depends on labeled data. The central idea of our approach is to find sparse networks that share the *same* starting point in function space as a corresponding dense network and whose network output evolution during supervised training is characterized by similar dynamics. If both criteria are satisfied, also the endpoints will be closely matched, thus ensuring the trainability of the sparse network (Fig. 1). To influence the training dynamics of a sparse network *before* training it, we leverage recent theoretical insights pertaining to the neural tangent kernel (NTK) (Jacot et al., 2018; Arora et al., 2019; Lee et al., 2019a; Chizat et al., 2019), which allow anticipating the training dynamics in a label-free manner, and devise an efficient algorithm that exploits this knowledge to constrain the sparse network’s output evolution to our advantage.

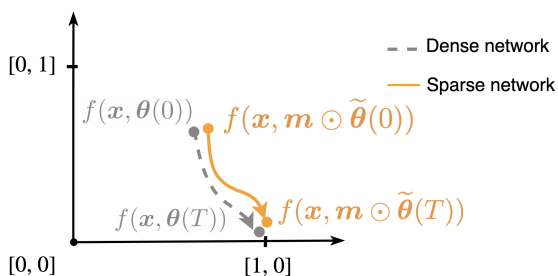


Figure 1. Schematic illustration of neural networks’ output evolution during supervised training from time $t = 0$ (starting point) to $t = T$ (endpoint). Here, $[1, 0]$ and $[0, 1]$ are one-hot targets for a binary classification task; the input data \mathbf{x} comes from the class whose correct label is $[1, 0]$. The dashed grey curve shows to the output evolution of a dense network, which, during training, moves toward the correct target $[1, 0]$. To ensure trainability, we aim to find a sparse network whose starting point and output evolution (orange, solid curve) during subsequent training closely follows the dense one’s, such that both output evolution terminate at similar endpoints.

More formally, consider a densely parameterized feed-forward neural network $f(\mathbf{x}, \boldsymbol{\theta})$ with input data \mathbf{x} and trainable parameters $\boldsymbol{\theta}$. During supervised learning, its output evolution is given by $\{f(\mathbf{x}, \boldsymbol{\theta}(t))\}_{t \geq 0}$ with t being the training time (e.g., iteration number) and $\boldsymbol{\theta}(t)$ the parameters at that time. Similarly, consider training a sparse network $f(\mathbf{x}, \mathbf{m} \odot \tilde{\boldsymbol{\theta}})$, where $\tilde{\boldsymbol{\theta}}$ are the trainable parameters, \mathbf{m} is a fixed sparsity-inducing mask, and \odot is the Hadamard product. The output evolution of the sparse network during training is described by $\{f(\mathbf{x}, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(t))\}_{t \geq 0}$. For the sparse network to yield good performance, we aim to select a good combination of \mathbf{m} and $\tilde{\boldsymbol{\theta}}(0)$ to ensure that its output evolution approximately follows the dense network’s evolution throughout training, i.e.,

$$f(\mathbf{x}, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(t)) \approx f(\mathbf{x}, \boldsymbol{\theta}(t)), \quad (1)$$

for all training inputs \mathbf{x} and training time $t \geq 0$ (cf. Fig. 1).

With this goal in mind, we proceed in the following steps. First, we characterize both the dense and sparse networks' output evolution (Section 3.1) during supervised learning using the tool of NTKs. Second, we motivate an objective function aimed at minimizing the difference between the two respective evolution (Section 3.2). Finally, we introduce an algorithm that minimizes this objective function and, by tuning the sparse network to have similar training dynamics to a corresponding dense network, finds a trainable sparse network in a label-free manner (Section 3.3).

3.1. Label-free characterization of output evolution for dense and sparse networks

In this subsection, we aim to characterize the output evolution of dense and sparse neural networks during training using analytical insights and tools developed in the context of linearized networks and NTK (Jacot et al., 2018; Arora et al., 2019; Lee et al., 2019a; Chizat et al., 2019).

We first focus on a dense neural network $f(\cdot, \theta) : \mathbb{R}^d \rightarrow \mathbb{R}$ for d being the dimension of input and θ the vector of all parameters. Examples of such dense networks include multi-layer perceptrons (MLPs) with fully connected layers and CNNs with dense filters. For ease of notation, we assume that the network has a scalar-valued output, but the framework can easily be extended to the vector-valued case. Given a training dataset $\{\mathbf{x}_i, y_i\}_{i=1}^n \subset \mathbb{R}^d \times \mathbb{R}$ of n input-target pairs, we consider the empirical risk minimization problem with the quadratic loss function

$$L(\theta) = \frac{1}{2} \sum_{i=1}^n (f(\mathbf{x}_i, \theta) - y_i)^2 \quad (2)$$

$$= \frac{1}{2} \|f(\mathcal{X}, \theta) - \mathbf{y}\|_2^2. \quad (3)$$

where for ease of notation we have written $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^n$ as the collection of training inputs, $f(\mathcal{X}, \theta) = (f(\mathbf{x}_i, \theta))_{i \in [n]} \in \mathbb{R}^n$ as the concatenation of network outputs of input data \mathbf{x}_i , and $\mathbf{y} = (y_i)_{i \in [n]} \in \mathbb{R}^n$ as the corresponding targets.

We now consider training the network with continuous-time gradient descent as characterized by the gradient flow

$$\frac{d\theta(t)}{dt} = -\nabla_{\theta} L(\theta(t)), \quad (4)$$

for training time $t \geq 0$.

From Eqns. (3) and (4), and by applying the chain rule follows the network's output evolution during training

$$\frac{df(\mathcal{X}, \theta(t))}{dt} = -\mathbf{H}_{\theta(t)} [f(\mathcal{X}, \theta(t)) - \mathbf{y}], \quad (5)$$

where $\mathbf{H}_{\theta(t)}$ is an $n \times n$ positive semidefinite matrix whose (i, j) -th entry is the value of the inner product

$\langle \nabla_{\theta} f(\mathbf{x}_i, \theta(t)), \nabla_{\theta} f(\mathbf{x}_j, \theta(t)) \rangle$. At initialization $t = 0$, the inner product function

$$K_{\theta(0)}(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}, \\ (\mathbf{x}, \mathbf{x}') \mapsto \langle \nabla_{\theta} f(\mathbf{x}, \theta(0)), \nabla_{\theta} f(\mathbf{x}', \theta(0)) \rangle \quad (6)$$

is called the empirical neural tangent kernel of the network (Jacot et al., 2018).

The dynamics of Eqn. (5) are difficult to analyze because of the time-varying matrix $\mathbf{H}_{\theta(t)}$. To simplify the analysis, previous work has resorted to studying the neural network's linearized approximation around initialization (Lee et al., 2019a; Chizat et al., 2019). Formally, for a neural network model $f(\mathbf{x}, \theta)$ with input \mathbf{x} , parameters θ , and initialization $\theta(0)$, its linearization $f^{\text{lin}}(\mathbf{x}, \theta)$ around initial parameters is defined as

$$f^{\text{lin}}(\mathbf{x}, \theta) := f(\mathbf{x}, \theta(0)) + \langle \theta - \theta(0), \nabla_{\theta} f(\mathbf{x}, \theta(0)) \rangle. \quad (7)$$

Let $f^{\text{lin}}(\mathcal{X}, \theta(t)) = (f^{\text{lin}}(\mathbf{x}_i, \theta(t)))_{i \in [n]} \in \mathbb{R}^n$ be the concatenation of linearized network outputs of all training samples \mathbf{x}_i at time t . The evolution of the linearized network's output is described by the following first order ordinary differential equation (ODE)

$$\frac{df^{\text{lin}}(\mathcal{X}, \theta(t))}{dt} = -\mathbf{H}_{\theta(0)} [f^{\text{lin}}(\mathcal{X}, \theta(t)) - \mathbf{y}]. \quad (8)$$

Note that Eqn. (8) represents a substantial simplification of Eqn. (5) because it only depends on $\mathbf{H}_{\theta(0)}$, which is constant during training and fully characterized by the network's initialization. The solution of the corresponding system of linear ODEs is known analytically

$$f^{\text{lin}}(\mathcal{X}, \theta(t)) = e^{-t\mathbf{H}_{\theta(0)}} f^{\text{lin}}(\mathcal{X}, \theta(0)) + [\mathbf{I} - e^{-t\mathbf{H}_{\theta(0)}}] \mathbf{y}, \quad (9)$$

for $t \geq 0$. While this linear approximation of a network's output evolution becomes exact as the width of the neural network goes towards infinity (Jacot et al., 2018), empirically it is known that the linear approximation is quite accurate even for finite-width networks (Lee et al., 2019a).

So far we have considered the output evolution of dense networks and their linearized approximations. We now turn to the analysis of neural networks with a fixed sparsity-inducing mask. In accordance with Eqn. (7), we define the linearized sparse network with input \mathbf{x} , mask \mathbf{m} , trainable parameters $\tilde{\theta}$, and initialization $\tilde{\theta}(0)$ as

$$f^{\text{lin}}(\mathbf{x}, \mathbf{m} \odot \tilde{\theta}) \\ := f(\mathbf{x}, \mathbf{m} \odot \tilde{\theta}(0)) + \langle \tilde{\theta} - \tilde{\theta}(0), \nabla_{\tilde{\theta}} f(\mathbf{x}, \mathbf{m} \odot \tilde{\theta}(0)) \rangle. \quad (10)$$

In line with Eqn. (9), the output evolution of linearized sparse neural network in Eqn. (10) admits the analytic solution

$$\begin{aligned} f^{\text{lin}}(\mathcal{X}, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(t)) \\ = e^{-t\mathbf{H}_{\mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)}} f^{\text{lin}}(\mathcal{X}, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)) + [\mathbf{I} - e^{-t\mathbf{H}_{\mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)}}] \mathbf{y}, \end{aligned} \quad (11)$$

where $\mathbf{H}_{\mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)}$ is an $n \times n$ positive semidefinite matrix whose (i, j) -th entry is $\langle \nabla_{\tilde{\boldsymbol{\theta}}} f(\mathbf{x}_i, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)), \nabla_{\tilde{\boldsymbol{\theta}}} f(\mathbf{x}_j, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)) \rangle$.

3.2. The neural tangent transfer (NTT) objective function

After having characterized the training dynamics of dense and sparse networks, we proceed to use this knowledge in a label-free manner to find sparse neural networks whose training evolution will remain close to the corresponding dense network’s evolution during subsequent training. In other words, for a given instance of dense network at initialization $f(\mathcal{X}, \boldsymbol{\theta}(0))$, we wish to select \mathbf{m} and $\tilde{\boldsymbol{\theta}}(0)$ that minimize the distance

$$\sum_{t=t_0}^{t_T} \|f(\mathcal{X}, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(t)) - f(\mathcal{X}, \boldsymbol{\theta}(t))\|_2^2 \quad (12)$$

during training in discrete time $0 \leq t_0 < \dots < t_T$. Unfortunately, it is impossible to evaluate Eqn. (12) without labels because the trained parameters $\{\tilde{\boldsymbol{\theta}}(t)\}_{t>0}$ are the result of the label-dependent supervised learning procedure. We therefore proceed in two steps. First, we make use of the linear approximation motivated in the preceding subsection. Second, we derive an auxiliary objective that is sufficient to minimize our objective derived in the first step, but in a label-free manner.

We start by replacing the corresponding output evolution in Eqn. (12) by its linear approximations

$$\sum_{t=t_0}^{t_T} \|f^{\text{lin}}(\mathcal{X}, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(t)) - f^{\text{lin}}(\mathcal{X}, \boldsymbol{\theta}(t))\|_2^2. \quad (13)$$

Note that Eqn. (13) is still label-dependent due to its dependence on the trained parameters $\{\tilde{\boldsymbol{\theta}}(t)\}_{t>0}$. Yet, by comparing the linearized network’s outputs given by Eqn. (9) and Eqn. (11), we see that a sufficient condition for minimizing the quantity in Eqn. (13) is to minimize the following label-free objective:

$$\begin{aligned} J_{\boldsymbol{\theta}(0)}(\mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)) &= \frac{1}{n} \|f(\mathcal{X}, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)) - f(\mathcal{X}, \boldsymbol{\theta}(0))\|_2^2 \\ &\quad + \frac{\gamma^2}{n^2} \|\mathbf{H}_{\mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)} - \mathbf{H}_{\boldsymbol{\theta}(0)}\|_F^2, \end{aligned} \quad (14)$$

where $\gamma^2 > 0$ is a strength parameter and $\|\cdot\|_F$ denotes the Frobenius norm of a matrix. We note that the objective function $J_{\boldsymbol{\theta}(0)}$ contains two terms. The first term $\frac{1}{n} \|f(\mathcal{X}, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)) - f(\mathcal{X}, \boldsymbol{\theta}(0))\|_2^2$ measures the distance between the dense and the sparse network’s starting point (cf. Fig. 1). The second term $\frac{\gamma^2}{n^2} \|\mathbf{H}_{\mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0)} - \mathbf{H}_{\boldsymbol{\theta}(0)}\|_F^2$ measures the discrepancy between the empirical NTKs of the dense and sparse network. When minimized, the latter term encourages the output evolution of a sparse network to remain close to its dense counterpart during training. We call the optimization process that minimizes $J_{\boldsymbol{\theta}(0)}(\mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0))$ *neural tangent transfer* (NTT). As the NTT objective function depends on the dense network’s initialization $\boldsymbol{\theta}(0)$, for convenience, we sometimes refer to the dense network $f(\cdot, \boldsymbol{\theta}(0))$ as the NTT teacher and the sparse network $f(\cdot, \mathbf{m} \odot \tilde{\boldsymbol{\theta}}(0))$ as the NTT student.

Since the NTT objective function is based on a linear approximation of a neural network model, we are prompted to ask what we can expect if the approximation is exact. Proposition 1 states that when the NTT teacher and student models are linear, the sparse student model that minimizes the NTT objective function possesses the same trainability as the dense teacher. That means, upon supervised training, the sparse student’s output evolution matches its dense teacher’s evolution exactly.

Proposition 1. *Consider a linear NTT teacher model $f(\mathbf{x}, \mathbf{a}(0)) = \mathbf{a}(0)^\top \mathbf{x}$, where \mathbf{x} and $\mathbf{a}(0) \in \mathbb{R}^d$ are model input and initial parameters. Suppose that we are given a linear NTT student model $g_{\mathbf{m}}(\mathbf{x}, \tilde{\mathbf{a}}(0)) = (\mathbf{m} \odot \tilde{\mathbf{a}}(0))^\top \mathbf{x}$ whose initial parameters $\tilde{\mathbf{a}}(0)$ are NTT-optimal in the sense that $J_{\mathbf{a}(0)}(\mathbf{m} \odot \tilde{\mathbf{a}}(0)) = 0$. Then upon continuous-time and quadratic-loss based gradient descent training, the dense and sparse models’ outputs evolve in the same way:*

$$f(\mathbf{x}, \mathbf{a}(t)) = g_{\mathbf{m}}(\mathbf{x}, \tilde{\mathbf{a}}(t)),$$

for all training inputs \mathbf{x} and time steps $t \geq 0$.

The proof to the Proposition 1 uses simple calculus and can be found in Appendix A.

3.3. Algorithmic implementation of NTT

The NTT objective function (Eqn. (14)) can readily be optimized using stochastic gradient descent (SGD) techniques. Specifically, we take following procedural steps:

1. **Choose a network architecture.** This includes fixing the number of layers, number of units per layer, etc.
2. **Instantiate the NTT teacher network.** Randomly initialize a dense network $f(\cdot, \boldsymbol{\theta})$ according to an initialization scheme (e.g., Glorot initialization (Glorot & Bengio, 2010)).

3. **Initialize the NTT student network.** For a desired sparsity level of the network $0 < p < 1$, remove $1 - p$ fraction of the weight parameters of θ having the smallest magnitudes to create a binary mask m . The NTT student is initialized to be $f(\cdot, m \odot \tilde{\theta})$, where $\tilde{\theta} = \theta$.
4. **Set the hyperparameters for SGD.** Choose SGD batch size, total number of iterations, optimizer, learning rate, the strength parameter γ^2 in the NTT objective function of Eqn. (14), and a *mask update frequency* (see the next step below).
5. **Adjust NTT student network parameters and masks via SGD.** To update the parameters, we fix a mask m and repeatedly draw minibatches of label-free data samples to perform SGD with weight decay:

$$\tilde{\theta} \leftarrow \tilde{\theta} - \eta \cdot \nabla_{\tilde{\theta}} J_{\theta}(m \odot \tilde{\theta}) - \beta \cdot m \odot \tilde{\theta}, \quad (15)$$

where $\eta > 0$ is the learning rate and β is a small weight-decay constant that helps the learned parameters generalize to the unseen data. Here, the back-propagated gradients flow through the fixed mask m and thus the masked-out parameters are not updated. Every a few steps, as specified by the *mask update frequency*, we update the mask based on the current weight magnitudes.

6. **Return the optimized NTT student** Upon completion of SGD, we are left with an optimized NTT student network $f(\cdot, m \odot \tilde{\theta}(0))$.

For all experiments in this article, we used the procedure outlined above. The specific experimental setup and all hyper-parameter choices are provided in the corresponding experiment sections and Appendix B. Below we offer some general remarks pertaining to the implementation. Neural networks for classification tasks are usually trained using a softmax activation defined on the logits at the output. However, the softmax activation contains no trainable parameters. Therefore, we used the logits output directly for NTT. On subsequent supervised learning tasks, in which labels are present, we equipped the last layer with a softmax activation unless mentioned otherwise. With slight modifications, we applied the outlined steps to produce either layerwise or globally sparse networks (see Sec. 4 for details). In either scheme, we only pruned weights and no bias terms.

4. Experiments

We numerically evaluated the trainability of sparse networks found by the NTT algorithm for both MLPs and CNNs on standard datasets, including MNIST (LeCun et al., 1998), Fashion MNIST (Xiao et al., 2017), CIFAR-10 (Krizhevsky

& Hinton, 2009), and SVHN (Netzer et al., 2011). All experiments were performed in JAX (Bradbury et al., 2018) together with the neural-tangent library (Novak et al., 2020).

General experiment procedure. We adopted the following procedure to evaluate the trainability of the sparse network yielded by NTT. First, we carry out NTT only using label-free data to learn the sparse network initialization (*NTT initialization*). To evaluate the learned results, we presented the labels in subsequent supervised learning tasks and used the NTT initialization to solve the task.

Baselines. We compared NTT against the following baselines:

- Randomly sampled sparse networks from densely initialized neural networks.
- Variance-scaled random sparse networks derived as a modification of Kaiming initialization (He et al., 2015): For w being a random variable taking values as weights in a layer, we set $\text{Var}(w) = 2/(n^{\text{in}} \cdot p)$ prior to pruning, where Var is the variance, n^{in} is the number of input units (fan-in) to that layer and p is fraction of remaining weights.
- Sparse networks pruned with SNIP (Lee et al., 2019b), a label-dependent foresight pruning method. Note that SNIP was originally formulated for global pruning; straightforwardly, we extend SNIP to layerwise pruning and refer to this variant as Layerwise-SNIP (see Appendix B.4).
- Sparse networks pruned with Logit-SNIP, a label-free variant of SNIP obtained by replacing the supervised loss function that occurs in SNIP’s pruning objective with the sum of squares of the network’s logits (for details, see Appendix B.4).

Pruning schemes. We considered both *layerwise* and *global* pruning schemes for the proposed and the baseline methods. In the layerwise scheme, we removed a fixed fraction of the parameters in each layer; in the global scheme, we pruned weights in all layers collectively without enforcing the weight-sparsity level of any specific layer. In the main text of this article, we primarily focused on layerwise pruning (Sec. 4.1 and 4.2). However, we devote Section 4.3 to a case study of global pruning methods and defer a more detailed study of global pruning to the Appendix.

4.1. Visualizing output dynamics of NTT initialization

We first sought to confirm our intuitions underlying the formulation of NTT (cf. Fig. 1). To that end, we devised a simple toy problem in which we are able to visually compare output evolution of sparse and dense networks. Recall

that the goal of the NTT is to learn a well-initialized sparse network (NTT student) so that its output dynamics under supervised training mimics the dynamics of its dense counterpart (NTT teacher), as motivated by Eqn. (1) and Fig. 1. To see whether this goal is empirically fulfilled, we trained both dense and sparse neural network models on a binary classification task using a subset of the MNIST dataset of handwritten digits (0 and 1). Since this is a binary classification task, the network outputs at each training time are two-dimensional vectors and therefore can be visualized.

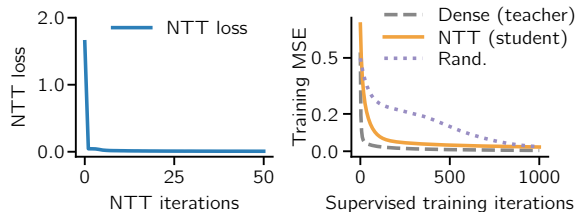


Figure 2. NTT loss and supervised training loss. Using the NTT algorithm (Sec. 3.3), we randomly initialized a dense network (NTT teacher) and used it to derive a sparse network (NTT student). The NTT loss during this process is shown in the left panel. We then used the NTT teacher (dense), the NTT student (sparse), and a random sparse network as initialization to solve the binary digit classification task (right panel). The dense NTT teacher (grey, dashed curve) solved the task with ease. The NTT student (orange curve) solved the task slower than the NTT teacher, but much faster than the random sparse network (purple, dotted curve).

Concretely, we used an MLP with two hidden layers containing 300 and 100 units with rectified linear units (ReLUs) followed with 2 linear output units. We randomly initialized the dense network using the Glorot initialization scheme (Glorot & Bengio, 2010) and designate it as the NTT teacher. We then perform NTT using the algorithm outlined in subsection 3.3 to train a 90% sparse NTT student network (see Appendix C.1 for the NTT parameters). During training, the loss of the NTT objective function dropped sharply (Fig. 2 left). Up to this point, no data label was used.

Next, to assess the trainability of the NTT student, we used a supervised label-dependent loss to train the NTT teacher (dense), NTT student (sparse), and a randomly initialized sparse network with 5000 iterations of gradient descent. During training, we tracked both the loss and the output evolution of all networks. All networks solved this simple problem with low error (Fig. 2 right). Not surprisingly, the dense NTT teacher network solved this simple task easily, reducing the training error with only a few iterations (dashed, grey curve). Importantly, the sparse NTT student network (orange curve) minimized the training error much faster than the randomly initialized sparse network (purple, dotted curve).

Finally, we visualized the networks’ output evolution. Because predictions are two-dimension approximating

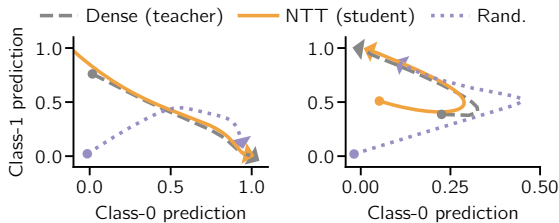


Figure 3. The NTT student’s output dynamics closely mimic the NTT teacher. Left: output evolution of different networks during supervised training obtained as the averaged prediction of digit-0 images. The arrows indicate the direction of the output evolution. Right: Same as left, but for digit-1 images. As expected, we found that the output evolution of the NTT teacher network (dashed grey curve) and the NTT student network (orange curve) closely resemble each other whereas the output evolution of the random sparse control network (dotted purple curve) was notably different.

one-hot encoded binary labels of images, we averaged predictions across images from the same ground-truth class (digit 0 or 1) at each iteration. As anticipated, we found that the output evolution of the NTT teacher and student were similar whereas the random sparse initialization was notably different (Fig. 3).

While the above analysis was conducted using an MLP, we have performed a similar experiment with a CNN and found qualitatively similar dynamics (Appendix C.1). These results confirmed our intuition that NTT initialized networks do indeed behave qualitatively similar to their dense counterparts during training.

4.2. Evaluating layerwise sparse NTT initialization

Having confirmed our intuition on the toy example, we now proceed to examine our method on larger datasets including MNIST, Fashion MNIST, CIFAR-10, and SVHN.

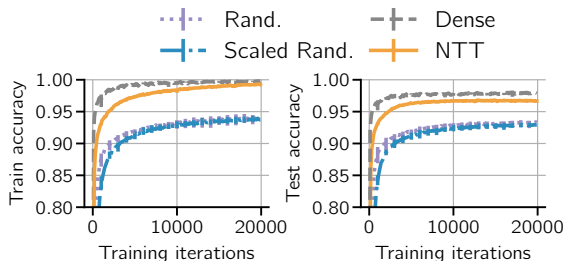


Figure 4. Learning curves of NTT and baseline networks trained on MNIST. The curves illustrate training and testing accuracy of Lenet-300-100 networks. All networks except the dense ones (grey, dashed curves) are sparse networks with 3% nonzero weights. Each curve is the average of five independent trials. Error bars denote the minimum and maximum of any trial.

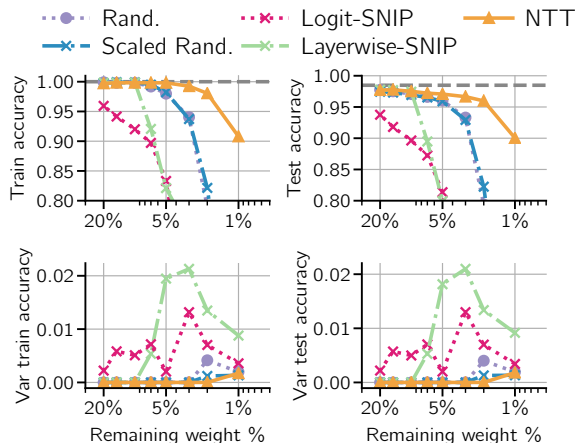


Figure 5. Supervised performance on MNIST using layerwise pruned Lenet-300-100. Top row: Classification accuracy on training (left) and testing datasets (right) for different levels of sparseness in networks with NTT initialization and baseline methods. Bottom row: Corresponding endpoint variance in accuracy of the data points in the row above computed over 5 repetitions.

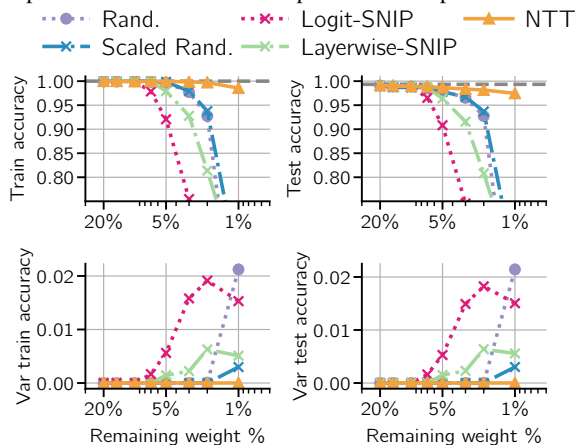


Figure 6. Supervised performance on MNIST using layerwise pruned Lenet-5-Caffe.

MNIST. In the MNIST experiments, we first applied NTT to MLPs of the Lenet-300-100 architecture (LeCun et al., 1998) at different sparsity levels. We noted that NTT initialized networks of a given sparsity could be trained faster and to higher accuracy than corresponding controls (Fig. 4).

When comparing the classification accuracy at different sparsity levels, NTT initialized networks showed generally higher accuracy (Fig. 5 top row). Moreover, we found that NTT initialization resulted in a lower variance of their final classification accuracy between different repetitions of the same experiment (Fig. 5, bottom row). These findings illustrate that NTT, as intended, finds sparse networks that can be trained faster, with higher accuracy, and reduced variability between trials.

To test whether these findings would generalize to CNNs,

we repeated the above experiments using a LeNet-5-Caffe CNN architecture, which is a modified variant of the LeNet-5 model proposed by LeCun et al. (1998). In good agreement with the MLP results, we found that NTT initialized sparse networks showed significantly better performance over baseline methods (Fig. 6 top) and a notable reduction of the variance of their final accuracy (Fig. 6 bottom).

Fashion MNIST. We repeated our experiments on Fashion MNIST. Similar to the MNIST task, we used Lenet-300-100 and Lenet-5. Both MLP (Lenet-300-100) and CNN (Lenet-5-Caffe) trained on fashion MNIST recapitulated our major findings on the MNIST experiments from above: Compared to baseline methods, NTT initialized networks converged faster, exhibited higher endpoint accuracy, and showed less variability over different repeats (Figs. 7 and 8).

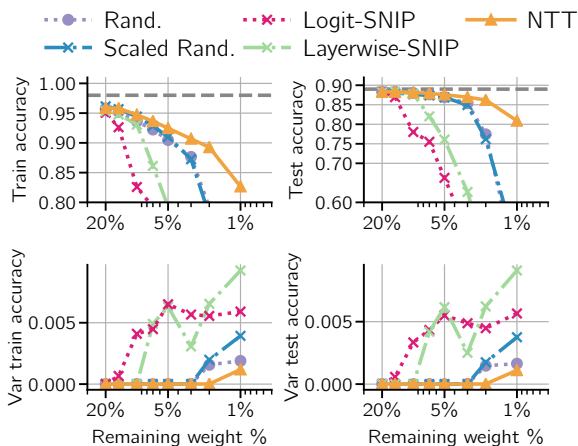


Figure 7. Supervised performance on Fashion MNIST using layerwise pruned Lenet-300-100.

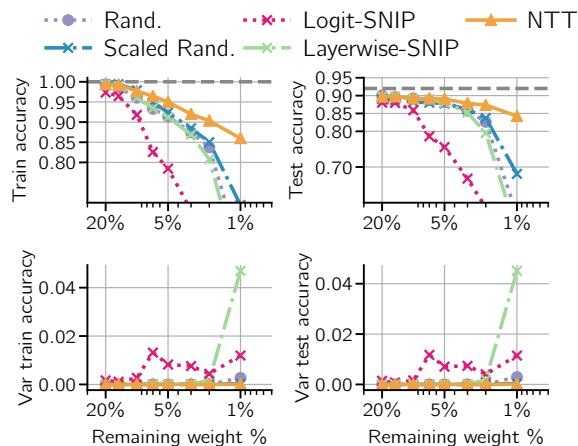


Figure 8. Supervised performance on Fashion MNIST using layerwise pruned Lenet-5-Caffe.

CIFAR-10. For the CIFAR-10 classification task, we trained a CNN model that we called Conv-4. This model

has 4 convolution layers followed by 2 feedforward layers with dropout (see Appendix B.1 for details), corresponding to a slightly scaled-up version of the default CNN example from Keras CIFAR-10 classification tutorial¹.

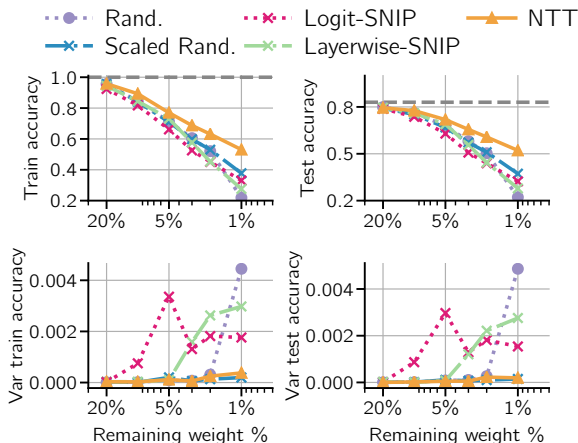


Figure 9. Supervised performance on CIFAR-10 using layerwise pruned Conv-4.

We found qualitatively similar behavior as in the above experiments: NTT achieved higher endpoint accuracy and reduced variance (Fig. 9) over baseline methods.

SVHN. Using Conv-4, we evaluated the performance of NTT on the larger SVHN dataset (Table 1). Consistent with our previous findings, we found NTT initialized networks outperformed other approaches at all sparsity levels.

4.3. Evaluating globally sparse NTT initialization

In the experiments reported above, we have focused on layerwise pruning. Compared to the layerwise pruning setting, global pruning typically results in higher accuracy, but with a lower gain in theoretical speedup in CNNs (Blalock et al., 2020). To investigate the performance of NTT in both settings, we applied NTT to produce globally sparse Conv-4 on CIFAR-10. The results are summarized in Table 2. These results confirmed our expectations that all methods achieved higher performance in the global setting than in the corresponding layerwise setting. Also, we observed that NTT initialized networks offer a small but consistent improvement over other methods in the global setting. Admittedly, the improvement is lower than in the layerwise case.

We next assessed the accuracy-efficiency tradeoff imposed by layerwise and global pruning. First, we numerically confirmed that global pruning tends to preserve orders of magnitude more parameters in the convolutional layers than layerwise pruning at a fixed sparsity level (top panel, Fig. 10). Since the computational cost of CNNs is dominated by these

convolutional layers (Li et al., 2017), global pruning leads to a much lower speedup than layerwise pruning (bottom panel, Fig. 10). Thus NTT can be used to exploit this tradeoff between accuracy and efficiency as it excels at both layerwise and global pruning.

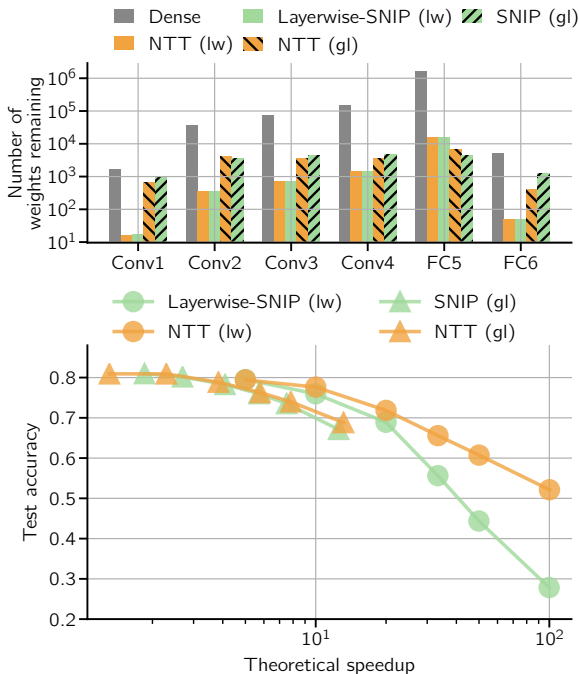


Figure 10. Tradeoff between model quality and efficiency. Layerwise and global pruning methods are respectively indicated by lw and gl in figure legends. Top panel: The number of remaining weights in each layer of 1% sparse Conv-4 networks on CIFAR-10. Bottom panel: Accuracy for Conv-4 on CIFAR-10 for several sparsity levels and their corresponding theoretical speedups. The theoretical speedup is defined as the number of multiply-adds used by the dense model divided by the number used by the sparse model (Blalock et al., 2020).

5. Discussion

We have introduced NTT, a foresight pruning method that finds trainable sparse neural networks capable of learning rapidly and generalizing well in subsequent supervised learning tasks. To that end, our method draws on theoretical insights from linearized neural networks and NTKs to derive sparse neural networks whose training evolution in network output space remains close to the path of a corresponding dense neural network. The two key advantages of our method are that (i) it only requires label-free data, and (ii) it can be used to find trainable layerwise sparse networks, e.g., CNNs with sparse convolutional filters, which are desirable for energy-efficient inference.

Our work was inspired by the Lottery Ticket Hypothesis (LTH) (Frankle & Carbin, 2019) in that we aim to identify “winning tickets,” i.e., sparse networks that can be trained to high accuracy. However, our approach is notably dif-

¹https://keras.io/examples/cifar10_cnn/

Table 1. Supervised performance on CIFAR-10 and SVHN using layerwise pruned Conv-4.

Dataset	CIFAR-10			SVHN		
	Remaining weights %	5%	3%	1%	5%	3%
Rand.	67.94 \pm 0.92	58.09 \pm 1.02	22.17 \pm 6.98	77.87 \pm 0.28	72.09 \pm 0.66	18.70 \pm 1.47
Scaled Rand.	66.86 \pm 1.05	57.17 \pm 0.84	37.24 \pm 1.22	79.55 \pm 0.26	75.22 \pm 0.27	48.47 \pm 3.10
Logit-SNIP	62.93 \pm 5.45	50.77 \pm 3.59	33.30 \pm 3.92	74.89 \pm 4.40	66.21 \pm 4.68	41.92 \pm 5.90
Layerwise-SNIP	68.90 \pm 0.87	55.64 \pm 3.47	27.84 \pm 5.25	75.62 \pm 3.28	64.08 \pm 4.77	43.58 \pm 3.30
NTT	71.84 \pm 0.72	65.58 \pm 0.76	52.14 \pm 1.39	81.49 \pm 0.45	79.04 \pm 0.32	71.63 \pm 0.80

Table 2. Supervised performance on CIFAR-10 using globally pruned Conv-4.

Remaining weights %	5%	3%	1%
Logit-SNIP	77.27 \pm 0.36	75.13 \pm 0.66	65.90 \pm 0.48
SNIP	78.28 \pm 0.33	76.00 \pm 0.48	67.12 \pm 0.46
NTT	78.85 \pm 0.36	76.28 \pm 0.32	68.81 \pm 0.28

ferent because it does not require supervised training of a dense network that is pruned subsequently. Instead, our approach transfers the *training dynamics* of an *untrained* dense network onto a sparse network at the time of its initialization. This idea is partially inspired by “reservoir transfer” (He et al., 2019) from a neuromorphic computing context. Our work is related to other foresight pruning methods that seek to ensure unimpeded error gradients (Lee et al., 2019b; Wang et al., 2020). Notably, both (Wang et al., 2020) and our work involve the NTK, but in different contexts. While NTT relies on the NTK to constrain the student networks’ output evolution under subsequent supervised learning, Wang et al. (2020) invoked the NTK to study the networks’ error gradients during supervised learning. As a result, their pruning criterion depends on a supervised loss function and thus on labeled data. NTT, on the other hand, is independent of the choice of the loss function and thus inherently label-free.

In biology, the sparseness of both connectivity and activity is thought to be a key design principle for energy-efficient information processing (Sterling & Laughlin, 2017). For similar reasons, sparse neural networks are increasingly moving into the focus in machine learning (Gong et al., 2014; Lebedev et al., 2015; Jaderberg et al., 2014) and neuromorphic engineers (Neftci et al., 2019; Roy et al., 2019). In all likelihood, neurobiology has evolved efficient algorithms implemented in its organic wetware to initialize its neural networks in an intelligent and energy-efficient way which may encode inductive biases and promote learning (Zador, 2019; Richards et al., 2019).

One of the limitations of the present work is that the calculation of NTK matrices is itself computationally expensive. Therefore we are particularly excited about the possibility of

exploring novel, more efficient ways to compute the NTK in the future (e.g., Li et al. (2019)). Another limitation is that we have only considered vision tasks and feed-forward networks in this article. It remains future work to explore NTT in other domains and architectures, e.g., temporal signal processing tasks.

Studies on efficient initialization techniques for sparse neural networks are still at their dawn, and profound theoretical insights are missing. In the foreseeable future, when the development of dedicated hardware will translate the hypothetical gains of sparse networks into significant performance gains, effective initialization schemes will become indispensable. With the introduction of NTT, we have made one step in this direction. This work opens up exciting new vistas for future work on elucidating the operating principles of computationally efficient sparse neural networks

Acknowledgements

This work was supported by the Novartis Research Foundation. We thank the anonymous reviewers for their valuable input. Further thanks go to David Belius, Ivan Dokmanić, and Yue (Kris) Wu for helpful discussions. Finally, we would like to thank the JAX and neural-tangents teams for developing libraries that make NTK computation fast and straightforward.

References

- 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 Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8139–8148. Curran Associates, Inc., 2019.
- Bellec, G., Kappel, D., Maass, W., and Legenstein, R. Deep rewiring: Training very sparse deep networks. In *International Conference on Learning Representations*, 2018.
- Blalock, D., Ortiz, J. J. G., Frankle, J., and Guttag, J. What is the state of neural network pruning?, 2020.

- 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>.
- Carreira-Perpinan, M. A. and Idelbayev, Y. "learning-compression" algorithms for neural net pruning. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8532–8541, June 2018.
- Chauvin, Y. A back-propagation algorithm with optimal use of hidden units. In Touretzky, D. S. (ed.), *Advances in Neural Information Processing Systems 1*, pp. 519–526. Morgan-Kaufmann, 1989.
- Chizat, L., Oyallon, E., and Bach, F. On lazy training in differentiable programming. 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. 2933–2943. Curran Associates, Inc., 2019.
- Collins, M. D. and Kohli, P. Memory bounded deep convolutional networks. *ArXiv*, abs/1412.1442, 2014.
- Evcı, U., Gale, T., Menick, J., Castro, P. S., and Elsen, E. Rigging the lottery: Making all tickets winners. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 471–481. 2020.
- Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019.
- Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In Teh, Y. W. and Titterton, M. (eds.), *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pp. 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010.
- Gong, Y., Liu, L., Yang, M., and Bourdev, L. Compressing Deep Convolutional Networks using Vector Quantization. *arXiv:1412.6115 [cs]*, December 2014. arXiv: 1412.6115.
- Guo, Y., Yao, A., and Chen, Y. Dynamic network surgery for efficient dnns. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 1379–1387. Curran Associates, Inc., 2016.
- Han, S., Pool, J., Tran, J., and Dally, W. Learning both weights and connections for efficient neural network. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 28*, pp. 1135–1143. Curran Associates, Inc., 2015.
- Hanson, S. J. and Pratt, L. Y. Comparing biases for minimal network construction with back-propagation. In Touretzky, D. S. (ed.), *Advances in Neural Information Processing Systems 1*, pp. 177–185. Morgan-Kaufmann, 1989.
- Hassibi, B., Stork, D. G., and Wolff, G. Optimal brain surgeon: Extensions and performance comparisons. In Cowan, J. D., Tesauro, G., and Alspector, J. (eds.), *Advances in Neural Information Processing Systems 6*, pp. 263–270. Morgan-Kaufmann, 1994.
- He, K., Zhang, X., Ren, S., and Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *IEEE International Conference on Computer Vision (ICCV 2015)*, 1502, 02 2015.
- He, X., Liu, T., Hadaeghi, F., and Jaeger, H. Reservoir transfer on analog neuromorphic hardware. In *9th International IEEE/EMBS Conference on Neural Engineering (NER)*, pp. 1234–1238, March 2019.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 8571–8580. Curran Associates, Inc., 2018.
- Jaderberg, M., Vedaldi, A., and Zisserman, A. Speeding up Convolutional Neural Networks with Low Rank Expansions. *arXiv:1405.3866 [cs]*, May 2014. arXiv: 1405.3866.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. *Master's thesis, Department of Computer Science, University of Toronto*, 2009.
- Lebedev, V., Ganin, Y., Rakhuba, M., Oseledets, I., and Lempitsky, V. Speeding-up Convolutional Neural Networks Using Fine-tuned CP-Decomposition. *arXiv:1412.6553 [cs]*, April 2015. arXiv: 1412.6553.
- LeCun, Y. Deep Learning Hardware: Past, Present, and Future. In *2019 IEEE International Solid-State Circuits Conference*, February 2019.
- LeCun, Y., Denker, J. S., and Solla, S. A. Optimal brain damage. In Touretzky, D. S. (ed.), *Advances in Neural Information Processing Systems 2*, pp. 598–605. Morgan-Kaufmann, 1990.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov 1998.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *Nature*, 521:436–444, May 2015.

- Lee, J., Xiao, L., Schoenholz, 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 Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8570–8581. Curran Associates, Inc., 2019a.
- Lee, N., Ajanthan, T., and Torr, P. H. S. SNIP: Single-shot network pruning based on connection sensitivity. In *ICLR*, 2019b.
- Lee, N., Ajanthan, T., Gould, S., and Torr, P. H. S. A signal propagation perspective for pruning neural networks at initialization. In *International Conference on Learning Representations*, 2020.
- Li, H., Kadav, A., Durdanovic, I., Samet, H., and Graf, H. P. Pruning filters for efficient convnets. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- Li, Z., Wang, R., Yu, D., Du, S. S., Hu, W., Salakhutdinov, R., and Arora, S. Enhanced convolutional neural tangent kernels, 2019.
- Liu, Z., Sun, M., Zhou, T., Huang, G., and Darrell, T. Re-thinking the value of network pruning. In *International Conference on Learning Representations*, 2019.
- Louizos, C., Welling, M., and Kingma, D. P. Learning sparse neural networks through l0 regularization. In *International Conference on Learning Representations*, 2018.
- Mocanu, D. C., Mocanu, E., Stone, P., Nguyen, P. H., Gibescu, M., and Liotta, A. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature Communications*, 9 (1):2383, 2018.
- Molchanov, D., Ashukha, A., and Vetrov, D. Variational dropout sparsifies deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.
- Morcos, A., Yu, H., Paganini, M., and Tian, Y. One ticket to win them all: generalizing lottery ticket initializations across datasets and optimizers. 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. 4933–4943. Curran Associates, Inc., 2019.
- Neftci, E. O. Data and Power Efficient Intelligence with Neuromorphic Learning Machines. *iScience*, 5:52–68, July 2018.
- Neftci, E. O., Mostafa, H., and Zenke, F. Surrogate Gradient Learning in Spiking Neural Networks: Bringing the Power of Gradient-based optimization to spiking neural networks, November 2019.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011.
- 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>.
- Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R. P., de Berker, A., Ganguli, S., Gillon, C. J., Hafner, D., Kepecs, A., Kriegeskorte, N., Latham, P., Lindsay, G. W., Miller, K. D., Naud, R., Pack, C. C., Poirazi, P., Roelfsema, P., Sacramento, J., Saxe, A., Scellier, B., Schapiro, A. C., Senn, W., Wayne, G., Yamins, D., Zenke, F., Zylberberg, J., Therien, D., and Kording, K. P. A deep learning framework for neuroscience. *Nat Neurosci*, 22 (11):1761–1770, November 2019.
- Rigollet, P. and Weed, J. Uncoupled isotonic regression via minimum Wasserstein deconvolution. *Information and Inference: A Journal of the IMA*, 8(4):691–717, 04 2019.
- Roy, K., Jaiswal, A., and Panda, P. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575(7784):607–617, November 2019.
- Schmidhuber, J. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., Driessche, G. v. d., Graepel, T., and Hassabis, D. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, October 2017.
- Sterling, P. and Laughlin, S. *Principles of Neural Design*. The MIT Press, reprint edition edition, June 2017. ISBN 978-0-262-53468-0.
- Ström, N. Sparse connection and pruning in large dynamic artificial neural networks, 1997.
- Wang, C., Zhang, G., and Grosse, R. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations*, 2020.

- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.
- Xu, L., Honda, J., Niu, G., and Sugiyama, M. Uncoupled regression from pairwise comparison data. 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. 3994–4004. Curran Associates, Inc., 2019.
- Yan, Y., Kappel, D., Neumärker, F., Partzsch, J., Vogginger, B., Höppner, S., Furber, S., Maass, W., Legenstein, R., and Mayr, C. Efficient reward-based structural plasticity on a spinnaker 2 prototype. *IEEE Transactions on Biomedical Circuits and Systems*, 13(3):579–591, June 2019.
- Yang, T., Chen, Y., and Sze, V. Designing energy-efficient convolutional neural networks using energy-aware pruning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pp. 6071–6079. IEEE Computer Society, 2017.
- Zador, A. M. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*, 10(1):3770, 2019.
- Zhu, M. and Gupta, S. To prune, or not to prune: Exploring the efficacy of pruning for model compression. In *6th International Conference on Learning Representations, ICLR 2018*, 2018.