

Dimensionality Reduced Training by Pruning and Freezing Parts of a Deep Neural Network, a Survey

Paul Wimmer (✉ Paul.Wimmer@de.bosch.com)

Robert Bosch (Germany)

Jens Mehnert

Robert Bosch (Germany)

Alexandru Paul Condurache

Robert Bosch (Germany)

Research Article

Keywords: Pruning, Freezing, Lottery Ticket Hypothesis, Dynamic Sparse Training, Pruning at Initialization

Posted Date: January 11th, 2023

DOI: <https://doi.org/10.21203/rs.3.rs-2458016/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: Competing interest reported. During preparation of this manuscript, all authors were employed by the Robert Bosch GmbH. Also, Alexandru Paul Condurache and Paul Wimmer were part of the Institute for Signal Processing of the University of Luebeck.

Version of Record: A version of this preprint was published at Artificial Intelligence Review on May 1st, 2023. See the published version at <https://doi.org/10.1007/s10462-023-10489-1>.

Dimensionality Reduced Training by Pruning and Freezing Parts of a Deep Neural Network, a Survey

Paul Wimmer^{1,2*}, Jens Mehnert¹ and Alexandru Paul
Condurache^{1,2}

^{1*}Automated Driving Research, Robert Bosch GmbH,
Burgenlandstraße 44, Stuttgart, 70469, Germany.

²Institute for Signal Processing, University of Lübeck,
Ratzeburger Allee 160, Lübeck, 23562, Germany.

*Corresponding author(s). E-mail(s):

Paul.Wimmer@de.bosch.com;

Contributing authors: JensEricMarkus.Mehnert@de.bosch.com;

AlexandruPaul.Cundurache@de.bosch.com;

Abstract

State-of-the-art deep learning models have a parameter count that reaches into the billions. Training, storing and transferring such models is energy and time consuming, thus costly. A big part of these costs is caused by training the network. Model compression lowers storage and transfer costs, and can further make training more efficient by decreasing the number of computations in the forward and/or backward pass. Thus, compressing networks also at training time while maintaining a high performance is an important research topic. This work is a survey on methods which reduce the number of trained weights in deep learning models throughout the training. Most of the introduced methods set network parameters to zero which is called pruning. The presented pruning approaches are categorized into pruning at initialization, lottery tickets and dynamic sparse training. Moreover, we discuss methods that freeze parts of a network at its random initialization. By freezing weights, the number of trainable parameters is shrunken which reduces gradient computations and the dimensionality of the model's optimization space. In this survey we first propose dimensionality reduced training as an underlying mathematical

model that covers pruning and freezing during training. Afterwards, we present and discuss different dimensionality reduced training methods – with a strong focus on unstructured pruning and freezing methods.

Keywords: Pruning, Freezing, Lottery Ticket Hypothesis, Dynamic Sparse Training, Pruning at Initialization

1 Introduction

In recent years, deep neural networks (DNNs) have shown state-of-the-art (SOTA) performance in many artificial intelligence applications, like image classification (Pham et al, 2021), speech recognition (Park et al, 2020) or object detection (Wang et al, 2021a). These applications require optimizing large models with up to billions of parameters. Training and testing such large models has been made possible due to technological advances. As a consequence, SOTA models are trained on specialized hardware for fast tensor computations, like GPUs or TPUs. Usually, not only one GPU/TPU is utilized to train these models but many of them. For example, XLNet (Yang et al, 2019) needs 5.5 days of training on 512 TPU v3 chips. Not only training, but transferring, storing and evaluating big models is costly, too (Schwartz et al, 2020). In order to train SOTA models, huge amount of data is required (Peters et al, 2018; Mahajan et al, 2018; Kolesnikov et al, 2020; Dosovitskiy et al, 2021) which needs resources for collecting, labeling, storing and transferring it. Of course, the large number of parameters and data, along with high computational demands leads to excessive energy consumption for training and evaluating deep learning (DL) models. For example, training a big transformer in conjunction with a previously performed neural architecture search results in emissions of 284t CO₂ (Strubell et al, 2019). This is about 315× the CO₂ emissions of a passenger traveling by air from New York City to San Francisco.

Between 2012 and 2018, computations required for DL research have been increased by estimated 300,000 times which corresponds to doubling the amount of computations every few months (Schwartz et al, 2020), see Figure 1. This rate outruns by far the predicted one by Moore’s Law (Gustafson, 2011). The performance improvements of DL models are to a great extent induced by raising the number of parameters in the networks and/or increasing the number of computations needed to train and infer the network (Strubell et al, 2019, 2020). Orthogonal to the development of new, big SOTA models which need massive amount of data, hardware resources and training time it is important to simultaneously improve parameter, computational, energy and data efficiency of DL models.

Model compression lowers storage and transfer costs, speeds up inference by reducing the number of computations or accelerates the training which uses less energy. It can be achieved by methods such as *quantization*, *weight*




Fig. 1 (From the OpenAI blog post (Amodei et al, 2018 [Online])) Evolution of the number of computations required for training SOTA models in the first era and in the modern era of AI systems.

sharing, tensor decomposition, low rank tensor approximation, pruning or freezing. Quantization reduces the number of bits used to represent the network's weights and/or activation maps, thereby decreasing memory consumption and speeding up inference (Han et al, 2015; Courbariaux et al, 2015; Wu et al, 2016; Zhou et al, 2017; Zhang et al, 2018; Jacob et al, 2018; Li et al, 2019). Memory reduction and speed up can also be achieved by weight sharing (Nowlan and Hinton, 1992; Chen et al, 2015; Ullrich et al, 2017), tensor decomposition (Xue et al, 2013; Lebedev et al, 2015; Novikov et al, 2015) or low rank tensor approximation (Sainath et al, 2013; Denton et al, 2014; Liu et al, 2015) to name only a few. *Network pruning* (Janowsky, 1989; Mozer and Smolensky, 1989; Karnin, 1990; LeCun et al, 1990; Han et al, 2015; Guo et al, 2016; Qian and Klabjan, 2021) sets parts of a DNN's weights to zero. This can help to reduce the model's complexity and memory requirements, speed up inference (Blalock et al, 2020) and even improve the network's generalization ability (LeCun et al, 1990; Arora et al, 2018; Bartoldson et al, 2020; Barsbey et al, 2021).

Pruning DNNs can be divided into *structured* and *unstructured* pruning. Structured pruning removes channels, neurons or even coarser structures of the network (Anwar et al, 2017; Chen et al, 2020a; Huang and Wang, 2018; Zhuang et al, 2018, 2020; Joo et al, 2021; Wang et al, 2020b). This generates leaner architectures, resulting in reduced computational time. A more fine-grained approach is given by unstructured pruning where single weights are set to zero (Mozer and Smolensky, 1989; Janowsky, 1989; Karnin, 1990; LeCun et al, 1990; Frankle and Carbin, 2018; Lee et al, 2019b; Wang et al, 2020a; Tanaka et al, 2020). Unstructured pruning usually leads to a better performance than

4 Dimensionality Reduced Training, a Survey




Fig. 2 Graphical overview of the proposed dimensionality reduced training methods with a low dimensional initialization $\vartheta^{(0)} \in \mathbb{R}^d$. To use the superior trainability of big networks, the low dimensional space is embedded in a bigger space via $\Theta^{(0)} = \Psi^{(0)}(\vartheta^{(0)}) = \psi^{(0)} \cdot \vartheta^{(0)} + \chi^{(0)}$ where $\Theta^{(0)} \in \mathbb{R}^D$, $\psi^{(0)}$ is the sparse embedding and $d \ll D$. Here, $\chi^{(0)}$ is the optional affine part modeling the frozen parameters. The embedding is obtained by either (a) *scoring* the big random initialization $\Theta^{(0)}$, (b) using a pre-training step to generate $\psi^{(0)}$ or (c) mapping $\vartheta^{(0)}$ onto a randomly chosen subset of $\Theta^{(0)}$. During training, only the low dimensional $\vartheta^{(0)}$ is learned. Furthermore, the embedding $\Psi^{(t)}$ can be adjusted to improve results, see Section 6.

structured pruning (Li et al, 2017; Mao et al, 2017) but specialized soft- and hardware which supports sparse tensor computations is needed to actually reduce the runtime (Han et al, 2016; Parashar et al, 2017). For an overview of the current SOTA pruning methods we refer to Blalock et al (2020).

Freezing a DNN means that only parts of the network are trained, whereas the remaining ones are frozen at their initial/pre-trained values (Huang et al, 2004; Hoffer et al, 2018; Rosenfeld and Tsotsos, 2019; Wimmer et al, 2020; Sung et al, 2021). This leads to faster convergence of the networks (Huang et al, 2004) and reduced communication costs for distributed training (Sung et al, 2021). Furthermore, freezing reduces memory requirements if the networks are initialized with pseudorandom numbers (Wimmer et al, 2020).

In recent years, *training* only a sparse part of the network, *i.e.* a network with many weights fixed at zero or their randomly initialized value, became of interest to the DL community (Bellec et al, 2018; Frankle and Carbin, 2018; Lee et al, 2019b; Tanaka et al, 2020; Wang et al, 2020a; Wimmer et al, 2020, 2021, 2022), providing the benefit of reduced memory requirements and the potential of reduced runtime not only for inference but also for training. Diffeenderfer et al (2021) even show that sparsely trained models are more robust against distributional shifts than (i) their densely trained counterpart and (ii) networks pruned with *classical* techniques. With classical techniques we denote pruning during or after training where the sparse network is fine-tuned afterwards. A high level overview of recent approaches to prune a network at initialization is given in Wang et al (2021b). In our work, the underlying mathematical framework is generalized to include a broader class of training

methods reducing dimensionality. Moreover, our work discusses the analyzed methods in more detail.

Scope of this work

In this work, we focus on methods that reduce the number of trainable parameters in a DL model. For this, we reversely define dimensionality reduction by embedding a low dimensional space into a high dimensional one, see Section 2.2. The high dimensional space corresponds to the original DNN, whereas the low dimensional one is the space where the DNN is actually trained. Therefore, training proceeds with reduced dimensionality.

Pruning and *freezing* parts of the network during training are two methods that yield *dimensionality reduced training* (DRT). This survey compares different DRT strategies. Figure 2 shows a graphical overview of the proposed DRT methods. A structural comparison between *freezing* DNNs at initialization and *pruning* them is given in Figure 3. For pruning, we differentiate between *pruning at initialization* (PaI) which trains a *fixed* set of parameters in the network and *dynamic sparse training* (DST) which adapts the set of trainable parameters during training. Closely related to PaI is the so called *Lottery Ticket Hypothesis* (LTH) which uses well trainable sparse subnetworks, so called *Lottery Tickets* (LTs). LTs are obtained by applying train-prune-reset cycles to the network's parameters, starting with the dense network and finally chiseling out the subnetwork *at initialization* with desired sparsity. Since most pruning/freezing methods applied before training are unstructured, this survey will focus on unstructured DRT methods.

Structure of this work

First we propose the problem formulation and mathematical setup of this survey in Section 2. Then, we discuss different possibilities to reduce the network's dimensionality in Section 3. The LTH and PaI are introduced in Sections 4 and 5, respectively. This is followed by a comparison of different DST methods in Section 6. As last DRT method, we present freezing initial parameters in Section 7. Subsequently, the different DRT methods from Sections 4 - 7 are compared at a high level in Section 8. The survey is closed by drawing conclusions in Section 9.

For better readability, the Figures should best be viewed in the colored online version.

2 Problem formulation

We will start Section 2 with a general introduction into DL. This is followed by defining a mathematical framework describing DRT. This framework comprises PaI, LTH, DST and freezing parts of a randomly initialized DNN.

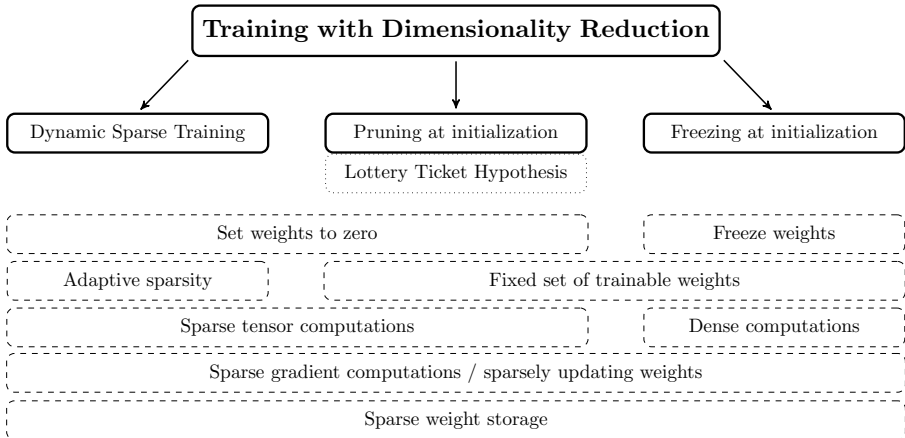


Fig. 3 Structural comparison between the three main categories of DRT. Here, the Lottery Ticket Hypothesis is seen as a sub-category of pruning at initialization.

2.1 General deep learning and setup

Let $f_{\Theta} : \mathbb{R}^m \rightarrow \mathbb{R}^n$ define a DNN with vectorized weights $\Theta \in \mathbb{R}^D \cong \mathbb{R}^{D_1} \times \dots \times \mathbb{R}^{D_L}$. Here, D_l denotes the number of parameters in layer l of a DNN with L layers and $D = \sum_{l=1}^L D_l$ parameters in total. We further assume the global network structure (activation functions, ordering of layers, ordering of weights inside one layer, ...) to be fixed and encoded in f_{Θ} , *i.e.* only the weight values Θ can be changed.

We assume a standard DNN training, starting with *random* initial weights $\Theta^{(0)} \in \mathbb{R}^D$ (He et al, 2015; Glorot and Bengio, 2010; Saxe et al, 2014; Martens, 2010; Sutskever et al, 2013; Hanin and Rolnick, 2018). These initial weights are iteratively updated with gradient based optimization (Robbins and Monro, 1951; Duchi et al, 2011; Kingma and Ba, 2015) to minimize the loss function

$$\mathcal{L} : \mathbb{R}^D \times \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}, (\Theta, X, Y) \mapsto \mathcal{L}(f_{\Theta}, X, Y) \quad (1)$$

over a training dataset $\mathcal{D} = (\mathcal{X}, \mathcal{Y}) \subset \mathbb{R}^n \times \mathbb{R}^m$. After initialization, a model is trained for T iterations forming a sequence of model parameters $\{\Theta^{(0)}, \Theta^{(1)}, \dots, \Theta^{(T)}\}$. Updating the parameters $\Theta^{(t)}$ in iteration t is done by calculating the current gradient $\nabla_{\Theta^{(t)}} \mathcal{L}$ over a batch of training data. The gradient is used, possibly together with former gradients $\nabla_{\Theta^{(0)}} \mathcal{L}, \dots, \nabla_{\Theta^{(t-1)}} \mathcal{L}$ and parameter values $\Theta^{(0)}, \dots, \Theta^{(t-1)}$, to minimize the loss function further. The final model is then given by $f_{\Theta^{(T)}}$.

2.2 Model for dimensionality reduced training

A reduction of dimensionality for $\Theta^{(t)}$ at training step t will be modeled through a transformation, also called *embedding*, $\Psi^{(t)} : \mathbb{R}^d \rightarrow \mathbb{R}^D$, $\Theta^{(t)} =$

$\Psi^{(t)}(\vartheta^{(t)})$ and $d \ll D$.¹ In this work, we restrict $\Psi^{(t)}$ to form an affine linear transformation which includes all standard pruning/freezing approaches. However, $\Psi^{(t)}$ has parameters which need to be stored. Therefore, we not only assume $d \ll D$ but also $\text{size}(\Psi^{(t)}) + d \ll D$, where $\text{size}(\cdot)$ computes the minimal number of parameters needed to express the affine linear transformation $\Psi^{(t)}$. To model DST, $\Psi^{(t)}$ is allowed to change during training.

In our setup, the parameter count in the network is reduced not only after training but also during it. Thus, the trainable parameters of the network are given by $\vartheta^{(t)} \in \mathbb{R}^d$. All methods discussed in this work are restricted to fulfill $\vartheta^{(t)} \in \mathbb{R}^d$ for all training iterations $t \in \{1, \dots, T\}$, *i.e.* reduce dimensionality throughout training. The corresponding model with reduced dimensionality is $f_{\Psi^{(t)}(\vartheta^{(t)})}$.

2.2.1 Pruning and dynamic sparse training

For pruning, $\Psi^{(t)}$ encodes the positions of the non-zero entries whereas $\vartheta^{(t)}$ stores the values of those parameters. A corresponding $\Psi^{(t)}$ can be easily constructed by setting

$$\Psi^{(t)}(\vartheta^{(t)}) = \psi^{(t)} \cdot \vartheta^{(t)}, \quad (2)$$

with the *pruning embedding* $\psi^{(t)} \in \mathbb{R}^{D \times d}$ defined via

$$\psi_{i,j}^{(t)} = \begin{cases} 1, & \text{if } \Theta_i^{(t)} \text{ is the } j^{\text{th}} \text{ un-pruned element.} \\ 0, & \text{else} \end{cases}, \quad (3)$$

where we assume the d un-pruned elements $\Theta_{i_1^{(t)}}^{(t)}, \dots, \Theta_{i_d^{(t)}}^{(t)}$ to have the natural ordering $i_1^{(t)} < \dots < i_d^{(t)}$. Consequently, the pruned version of $\Theta^{(t)}$ is encoded by the low dimensional

$$\vartheta^{(t)} = (\vartheta_j^{(t)})_j = (\Theta_{i_j^{(t)}}^{(t)})_j. \quad (4)$$

Further, we define $p := 1 - d/D$ as the model's *pruning rate*. For pruning at initialization, *i.e.* fixed position $i_1^{(0)}, \dots, i_d^{(0)}$ for the non-zero parameters, $\psi^{(t)} = \psi^{(0)}$ for all training iterations t . Whereas, $\psi^{(t)}$ might adapt for DST.

Since $d \ll D$, the embedding $\Psi^{(t)}(\vartheta^{(t)})$ is automatically sparse in the bigger space \mathbb{R}^D . A possibility to overcome sparse $\Theta^{(t)}$ while still training only a small part $\vartheta^{(t)}$ of a DNN is proposed in [Wimmer et al \(2022\)](#). Here, convolutional $K \times K$ filters are represented via few non-zero coefficients of an adaptive dictionary. The dictionary is shared over one or more layers of the network. This procedure can be described by the adaptive embedding

$$\Psi^{(t)} = \Phi^{(t)} \cdot \psi^{(t)} \quad (5)$$

¹We use the same transformation as introduced in [Mostafa and Wang \(2019\)](#) to embed \mathbb{R}^d into \mathbb{R}^D .

8 Dimensionality Reduced Training, a Survey




Fig. 4 Comparison between a standard DNN and a frozen one. Left: Standard dense model. Middle left: Topology of the frozen weights. Middle right: The forward pass is exactly the same for the dense and the frozen model. Right: The frozen weights are used for the backward pass. However, they are not updated and therefore drawn with dotted arrows.

with the pruning embedding $\psi^{(t)} \in \mathbb{R}^{D \times d}$ as defined in (3) and the trainable dictionary $\Phi^{(t)} \in \mathbb{R}^{D \times D}$.² In this case, the coordinates w.r.t. $\Phi^{(t)}$ are sparse, but the resulting representation in the spatial domain \mathbb{R}^D will be dense (Wimmer et al, 2022).

Orthogonal to that approach, Price and Tanner (2021) combine a sparse pruning embedding (2) with a dense *discrete cosine transformation* (DCT) by summing them up. A DCT is free to store and can be computed with $D \cdot \log D$ FLOPs. By doing so, the network keeps a high information flow while only a sparse part of the network has to be trained. Moreover, the computational cost is almost equal to a fully sparse model.

2.2.2 Freezing parameters

Another approach, which is similar to (Price and Tanner, 2021), proposed by Rosenfeld and Tsotsos (2019); Wimmer et al (2020); Sung et al (2021), uses frozen weights on top of the trainable ones. The resulting transformation is given by

$$\Psi^{(t)}(\vartheta^{(t)}) = \psi^{(t)} \cdot \vartheta^{(t)} + \chi^{(t)} \cdot \Theta^{(0)} \quad (6)$$

with the pruning embedding $\psi^{(t)} \in \mathbb{R}^{D \times d}$ from (3), $\chi^{(t)} = \text{diag}(\chi_1^{(t)}, \dots, \chi_D^{(t)}) \in \mathbb{R}^{D \times D}$ defined as

$$\chi_i^{(t)} = \begin{cases} 0, & \text{if } i \in \{i_1^{(t)}, \dots, i_d^{(t)}\}, \\ 1, & \text{else} \end{cases}, \quad (7)$$

and the random initialization $\Theta^{(0)} \in \mathbb{R}^D$. A graphical overview of freezing parts of a network is given in Figure 4. Similarly to pruning, we define $p := 1 - d/D$ as the model's *freezing rate*, *i.e.* the rate of parameters which are frozen at their random initial value. Note, Rosenfeld and Tsotsos (2019) and Wimmer et al (2020) both use a fixed set of un-frozen weights, *i.e.* $\Psi^{(t)} = \Psi^{(0)}$. Therefore, the network parameters $\Theta^{(t)} = \Psi^{(0)}(\vartheta^{(t)})$ consist of two parts, the dynamic un-frozen weights defined by $\psi^{(0)} \cdot \vartheta^{(t)}$ and the fixed, frozen weights $\chi^{(0)} \cdot \Theta^{(0)}$.

²Here, $\Phi^{(t)}$ corresponds to a block diagonal matrix with shared blocks. By sharing blocks, the total parameter count of $\Phi^{(t)}$ is at most $D/10,000$ (Wimmer et al, 2022). Furthermore, the formulation in Wimmer et al (2022) is not restricted to quadratic $D \times D$ matrices, but also allows *undercomplete* or *overcomplete* systems $\Phi^{(t)}$.




Fig. 5 Structured and unstructured pruning compared. Structured pruning removes whole nodes (corresponding to channels for CNNs) whereas unstructured pruning sets individual weights to zero.

Sung et al (2021) allow the embedding $\Psi^{(t)}$ to adapt during training, however their procedure still leaves a large part of the network completely untouched.

Another construction of $\Theta^{(t)}$ is given by interchanging the pruning embedding $\psi^{(t)}$ in (6) with a fixed, random orthogonal transformation $\psi_o \in \mathbb{R}^{D \times d}$ (Li et al, 2018).

3 High-level overview of dimensionality reducing transformations $\Psi^{(t)}$

In this Section we give a high-level overview of different approaches to determine the dimensionality reducing transformation $\Psi^{(t)}$. First, we discuss the *structure* of the parameters which are pruned or frozen in Section 3.1. Then, we compare *global* and *layer wise* dimensionality reduction in Section 3.2. Section 3.3 covers the *update frequency* for $\Psi^{(t)}$. Afterwards, the most prominent criteria for determining the pruning embedding $\psi^{(t)}$ are proposed in Section 3.4. Finally, we discuss *pre-training* the transformation $\Psi^{(t)}$ in Section 3.5.

Throughout this Section, we restrict the transformation $\Psi^{(t)}$ to use $\psi^{(t)}$ from (3), *i.e.* a sparse linear embedding $\psi^{(t)} \in \{0, 1\}^{D \times d}$ with $\|\psi^{(t)}\|_0 = d$, as linear part. Further we assume the affine part of $\Psi^{(t)}$ to be either zero (pruning) or correspond to the almost free to store pseudorandom initialization (freezing). We use the notion of *trainable* weights (in iteration t) for all weights $\Theta_i^{(t)}$ with $i \in \{i_1^{(t)}, \dots, i_d^{(t)}\}$, *i.e.* elements of $\{i : \exists j \text{ s.t. } \psi_{i,j}^{(t)} \neq 0\}$.

3.1 Structure of trainable weights

Pruning is generally distinguished in *structured* and *unstructured* pruning, see Figure 5. Of course it can be generalized to our setup, including freezing of parameters. Structured freezing/pruning means freezing/pruning whole neurons or channels or even coarser structures of the network. For pruning, this immediately results in reduced computational costs, whereas gradient computations can be skipped for the frozen structure. For structured freezing, usually whole layers are frozen (Huang et al, 2011; Qing et al, 2020; Saxe et al, 2011; Hoffer et al, 2018) with the extreme case of freezing the *entire* network (Giryes et al, 2016). For pruning it is more common to prune on the level of channels/neurons (Liu et al, 2019; Wang et al, 2020c; Verdenius et al, 2020).

Unstructured pruning of single weights usually improves results compared to structured pruning (Li et al, 2017; Mao et al, 2017). Unstructured pruning

has the disadvantage to require software which supports sparse computations to actually speed up the forward and backward propagation in DNNs. Furthermore, such software only accelerates sparse DNNs on CPUs (Park et al, 2017; Liu et al, 2021a) or specialized hardware (Han et al, 2016; Parashar et al, 2017; Elsen et al, 2020; Gale et al, 2020; Wang, 2020). Weights can also be frozen in an unstructured manner (Li et al, 2018; Wimmer et al, 2020; Sung et al, 2021), leading to reduced gradient computations, memory requirements and communication costs for distributed training.

In recent years, a semi-structured pruning strategy developed, the so called $N : M$ sparsity (Hubara et al, 2021; Zhou et al, 2021a; Sun et al, 2021; Pool and Yu, 2021; Holmes et al, 2021; NVIDIA, 2020). A tensor is defined as $N : M$ sparse if each block of size M contains (at least) N zeros. Here, the tensor is covered by non-overlapping, homogeneous blocks of size M .

As shown in Frankle and Carbin (2018), using sophisticated methods to determine a sparse, fixed subset of trainable parameters at random initialization greatly outperforms choosing them randomly. Contrarily, choosing and fixing the trainable channels randomly or via well performing classical techniques leads to similar results for structured pruning at initialization (Liu et al, 2019). Therefore, this survey focuses in the following on unstructured pruning and freezing methods.

3.2 Global or layerwise dimensionality reduction

An important distinction for DRT methods is whether the rate of the trained parameters is chosen separately for each layer or globally. As presented in Section 3.4, the *importance* of a weight $\Theta_i^{(0)}$ for training is measured via a score $s_i \in \mathbb{R}$. A weight is frozen or pruned if the corresponding score is below a threshold τ_i^3 . This threshold can be determined *globally* (Wimmer et al, 2020; Lee et al, 2019b; Frankle and Carbin, 2018; Wang et al, 2020a; Sanh et al, 2020; Bellec et al, 2018; Mostafa and Wang, 2019) or *layerwise*. For thresholding the score layerwise we differentiate between setting a constant rate of trainable parameters in each layer (Jayakumar et al, 2020), using heuristics or optimized hyperparameters to find the best rate of trainable parameters for each layer individually (Mocanu et al, 2018; Evci et al, 2020) or letting the number of trained parameters in each layer *adapt* dynamically during training (Dettmers and Zettlemoyer, 2019; Mostafa and Wang, 2019).

3.3 Update frequency

The frequency of updating the dimensionality reducing transformation $\Psi^{(t)}$ is called *update frequency* F_p , see Figure 6. Given an update frequency $F_p \in \mathbb{N}^+ \cup \{\infty\}$, it holds

$$\Psi^{(t_0)} = \Psi^{(t_0+1)} = \dots = \Psi^{(t_0+F_p-1)} \quad (8)$$

³We index the threshold with the same index as the weight to show that the threshold might depend on the position of $\Theta_i^{(0)}$ but not the weight itself.




Fig. 6 Dynamic sparse training where the initial pruning embedding $\Psi^{(0)}$ is usually determined randomly.

where $t_0 \in \{0, F_p, 2F_p, \dots\}$ is an iteration where the transformation is updated. If $\Psi^{(t)} = \Psi^{(0)}$ is constant, the update frequency equals $F_p = \infty$. We assume the number of trained parameters to be fixed. Consequently, the proposed methods usually have $F_p \gg 1$. This guarantees (i) stability of the training and (ii) a chance for newly trainable weights to grow big enough to be not pruned/frozen immediately in the next update step.

3.4 Criteria to choose trainable parameters

In the following, we present the most common criteria to determine the trainable weights.

Random criterion

Randomly selecting trainable parameters means that a given percentage of them are chosen to be trained by a purely random process, see Figure 2 (c). This can be done by sampling a score for each weight from a standard Gaussian distribution and training those weights with the d biggest samples. Random pruning is often used as first comparison for a newly developed pruning method. For dynamic sparse training, initially trained parameters are often chosen by random (Mocanu et al, 2018; Evci et al, 2020) as the dynamics of the sparse topology will find a well performing architecture anyway. Also for DRT methods where information flow is not a limiting factor, like PaI combined with DCTs (Price and Tanner, 2021) or freezing (Rosenfeld and Tsotsos, 2019; Li et al, 2018), trainable weights are often chosen randomly.

Magnitude criterion

A straight forward way to determine trainable weights is to use those with the highest magnitude. This is done by the LTH (Frankle and Carbin, 2018; Zhou et al, 2019) and also by many DST methods for updating the sparse architecture during training (Mocanu et al, 2018; Mostafa and Wang, 2019; Dettmers and Zettlemoyer, 2019; Evci et al, 2020). Magnitude pruning for Θ is equivalent to the solution of

$$\min_{\bar{\Theta} \in \mathbb{R}^D, \|\bar{\Theta}\|_0 \leq d} \|\Theta - \bar{\Theta}\|_q, \quad (9)$$

i.e. the best d -sparse approximation of Θ w.r.t. $\|\cdot\|_q$. Here, $q \in (0, \infty)$ is arbitrary. This shows the ability of magnitude pruning to approximate dense networks sparsely if the pruned parameters are not too big. In practice, this is guaranteed by pruning only small fractions of the parameters in one step (Frankle and Carbin, 2018; Frankle et al, 2020a; Han et al, 2015; Liu et al, 2019). The magnitude criterion in the viewpoint of dynamical systems is analyzed in Redman et al (2022). They show that magnitude pruning is equivalent to pruning small modes of the Koopman operator (Mezić, 2005) which determines the networks convergence behavior. Consequently, pruning small magnitudes only slightly disturbs the long term behavior of DNNs.

If dimensionality reduction is applied at initialization without any pre-training, the magnitude criterion chooses trainable parameters purely by their initial, randomly drawn weight. Still as shown in Frankle et al (2021a), magnitude pruning at initialization is a non-trivial baseline for other PaI methods and outperforms random PaI.

Gradient based criterion

As discussed above, the magnitude criterion is for the most part random at initialization. Consequently, different criteria to chose the trainable weights are used for SOTA PaI methods. The most common gradient based criterion relies on the first order Taylor expansion of the loss function at the beginning of training (Lee et al, 2019b, 2020; Verdenius et al, 2020; de Jorge et al, 2021; Hayou et al, 2021; Wimmer et al, 2020; Sanh et al, 2020) and is mainly used for PaI in our setup. It measures how disturbing the weights at initialization will affect the loss function. It holds

$$\mathcal{L}(\Theta^{(0)} - \delta) = \mathcal{L}(\Theta^{(0)}) - \nabla_{\Theta^{(0)}} \mathcal{L} \cdot \delta + \mathcal{O}(\|\delta\|^2). \quad (10)$$

In our setup, the disturbance introduced by the dimensionality reduction is given by $\delta = \Theta^{(0)} - \Psi^{(0)}(\vartheta^{(0)})$. Neglecting the higher order terms and the sign, the following optimization problem has to be solved

$$\max_{\Psi, \vartheta} |\nabla_{\Theta^{(0)}} \mathcal{L}| \cdot |\Psi(\vartheta)|, \quad (11)$$

in order to determine $\Psi^{(0)}$ and $\vartheta^{(0)}$. If Ψ is restricted to be a pruning transformation defined by (2) and (3), the optimization problem (11) is solved by training the indices with top- d $|(\nabla_{\Theta^{(0)}} \mathcal{L})_i \cdot \Theta_i^{(0)}|$.

Using the aforementioned gradient based criterion for PaI might lead to a loss of information flow for high pruning rates since some layers are pruned too aggressively (Wang et al, 2020a; Tanaka et al, 2020; Wimmer et al, 2020). Without information flow in the network, the gradient is vanishing and no training is possible. Thus, there are several approaches to overcome a weak gradient flow in networks pruned according to (11). Using a random initialization, fulfilling the layerwise *dynamic isometry* property (Saxe et al, 2014; Poole et al, 2016; Schoenholz et al, 2017; Xiao et al, 2018) will lead to an

improved information flow in the pruned network (Lee et al, 2020). Another approach (Hayou et al, 2021) uses a rescaling of the pruned network to bring it to the *edge of chaos* (Poole et al, 2016; Schoenholz et al, 2017; Xiao et al, 2018) which is beneficial for DNN training. A straight forward way, proposed by Wimmer et al (2020), is given by freezing the un-trained weights instead of pruning them.

Conserving information flow

Sufficient information flow can also be guaranteed directly by the criterion for selecting trainable parameters. Wang et al (2020a) train the sparse network with the highest total gradient ℓ_2 norm *after pruning*. Consequently, the information flow is not the bottleneck for training.

Synaptic saliency of weights is introduced in Tanaka et al (2020). A synaptic saliency is defined as

$$S(\Theta) = \nabla_{\Theta} \mathcal{R} \odot \Theta, \quad (12)$$

where $\mathcal{R} : \mathbb{R}^D \rightarrow \mathbb{R}$ is a function, dependent on the weights Θ . In Tanaka et al (2020), a conservation law for a layer’s total synaptic saliency is shown. By iteratively pruning a small fraction of parameters based on a synaptic saliency score (12), the conservation law guarantees a faithful information flow in the sparse network. Examples for a synaptic flow based pruning criterion are setting $\mathcal{R}(\Theta)$ as the ℓ_1 path norm (Neyshabur et al, 2015a,b) of the DNN f_{Θ} (Tanaka et al, 2020) or the ℓ_2 path norm (Gebhart et al, 2021; Patil and Dovrolis, 2021).

Trainable $\psi^{(0)}$

All aforementioned criteria are based on heuristics, usually limited to work well in some scenarios but fail for other ones, see Frankle et al (2021a) for an empirical comparison for the SOTA PaI methods Lee et al (2019b); Wang et al (2020a); Tanaka et al (2020). A natural way to overcome these limitations is given by *training* the embedding $\psi^{(0)}$. This approach is especially in the interest for pruning randomly generated networks *without* training the weights afterwards, where the determination of $\psi^{(0)}$ is the only way to optimize the network (Zhou et al, 2019; Ramanujan et al, 2020; Diffenderfer and Kaikhura, 2021; Aladago and Torresani, 2021).

Finding $\psi^{(0)} \in \{0, 1\}^{D \times d}$ can not be done by simple backpropagation since $\psi^{(0)}$ is optimized in a *discrete* set $\{\psi \in \{0, 1\}^{D \times d} : \|\psi\|_0 = d\}$. Ramanujan et al (2020); Mallya et al (2018); Diffenderfer and Kaikhura (2021); Aladago and Torresani (2021); Zhang et al (2021b) find the pruned weights with a trainable score $s \in \mathbb{R}^D$ together with a (shifted) sign function. The score s is optimized via backpropagating the error of the sparse network on the training set by using the straight through estimator (Bengio et al, 2013) to bypass the zero gradient of the sign function. Overcoming the vanishing gradient of zeroed scores can also be achieved by training a pruning probability for each weight

Table 1 Comparison of different methods to find lottery tickets. Data size ● means using the full dataset for finding LTs whereas ◐ shrinks the full dataset during the process of finding LTs. Note, all methods can be used in an iterative manner, but also as one-shot methods.

Method	Prune criterion	Late rewinding	Early stopping	Low precision	Data size	Transfer	Additional Trafo Φ
Frankle and Carbin (2018)	$ \Theta_i^{(0)} $	✗	✗	✗	●	✗	✗
Frankle et al (2020a)	$ \Theta_i^{(0)} $	✓	✗	✗	●	✗	✗
Morcos et al (2019)	$ \Theta_i^{(0)} $	✓	✗	✗	●	Datasets/Optimizers	✗
Zhang et al (2021a)	$ \Theta_i^{(0)} $	✓	✓	✗	◐	✗	✗
You et al (2020)	$ \Theta_i^{(0)} $	✗	✓	✓	●	✗	✗
Rosenfeld and Tsotsos (2019)	$ \Theta_i^{(0)} $	✓	✗	✗	●	Predict Performance	✗
Wimmer et al (2022)	$ \Theta_i^{(0)} $	✓	✗	✗	●	✗	Trainable
Lee et al (2021)	$ \Theta_i^{(0)} \left(\sum_{ \Theta_j^{(0)} \geq \Theta_i^{(0)} } \Theta_j^{(0)} \right)^{-1}$	✓	✗	✗	●	✗	✗
Zhang et al (2021b)	$ \Theta_i^{(0)} $	✓	✗	✗	●	✗	Inertial Manifold

and sampling a corresponding $\psi \in \{0, 1\}^{D \times d}$ in each optimization step (Zhou et al, 2019).

3.5 Pre-training the transformation

Similar to a trainable $\psi^{(0)}$, discussed in Section 3.4, also a pre-training step for finding $\psi^{(0)}$ can be applied. The difference is that for pre-training $\psi^{(0)}$ the corresponding dense weights $\Theta^{(0)}$ can be trained as well. Using pre-training for finding the initial transformation $\psi^{(0)}$ is of course costly in terms of time and also in terms of an increased parameter count during the pre-training phase. Note, *after* pre-training $\psi^{(0)}$, the pre-trained weights are *not allowed* to be used, but reset to $\Theta^{(0)}$.

The most prominent example for this is given by the LTH which trains $\Theta^{(0)}$ to convergence, prunes $p_0 \cdot 100\%$ of the non-zero weights and resets the remaining non-zero weights to their initial values. This procedure is continued until the desired pruning rate is reached. Then, the pruning transformation $\psi^{(0)}$ is applied to the initial weights $\Theta^{(0)}$ and fixed for the training of the sparse architecture. Here, the *full* network is used for finding a well performing sparse architecture while for the actual training only the sparse part of the network, starting from the random initialization, is used.

Another approach for pre-training $\psi^{(0)}$ is given in Liu and Zenke (2020) which pre-trains a sparse network to mimic the training dynamics of a randomly initialized dense network, measured by the *neural tangent kernel* (Jacot et al, 2018; Arora et al, 2019; Lee et al, 2019a).⁴

4 Lottery ticket hypothesis

In this Section we discuss the LTH, proposed by Frankle and Carbin (2018), in detail. Frankle and Carbin (2018) show that extremely sparse subnetworks of a randomly initialized network can be found which, after being trained, match or even outperform their densely trained baseline. Furthermore, the sparsely

⁴To be precise, in order to mimic the training dynamics, Liu and Zenke (2020) also pre-train the weights $\vartheta^{(0)}$ by using $\Theta^{(0)}$. Therefore, it is not within the narrow bounds of this work.

trained network converges at least as fast as standard training but usually faster. Table 1 summarizes different methods to find LTs.

The procedure to determine the sparse networks is as follows: First, the dense network is trained to convergence and the pruning embedding ψ is constructed according to the magnitude criterion, see Section 3.4. The un-pruned weights are reset to their values $\Theta_i^{(0)}$. Now, this procedure can be done *one-shot* or *iterative*. In the one-shot case, all weights are pruned after training the dense network at once. Then, the sparse network, the so called *Lottery Ticket* (LT), is trained to convergence. For the iterative procedure, not all coefficients but only $p_0 \cdot 100\%$ (usually 20%) of the un-pruned ones are pruned in one iteration. The remaining non-zero weights are reset to their initial value and trained again until convergence. This procedure is applied iteratively until the desired pruning rate p is achieved. Finally, the LT is optimized in a last, sparse training.

Iterative LTH has shown to find better performing LTs than the one-shot approach (Frankle and Carbin, 2018; Frankle et al, 2020a). But, in order to reach a final pruning rate p , the iterative procedure takes

$$\left\lceil \frac{\log(1-p)}{\log 4 - \log 5} \right\rceil \quad (13)$$

many pre-trainings plus the final, sparse training. For the final pruning rate $p = 0.9$, the iterative LTH approach needs in total 12 trainings of the network. On the other hand, the one-shot approach always requires 2 trainings for arbitrary pruning rates p .

Frankle and Carbin (2018); Gale et al (2019); Frankle et al (2020a) show that resetting un-pruned weights to their *initial* value only generates well trainable sparse architectures if the networks are not too big. For modern network architectures like ResNets (He et al, 2016), resetting the weights to their initial value does not lead to similar results as the densely trained baseline. This problem can be overcome by *late rewinding*, *i.e.* resetting the weights to values reached early in the first, dense training (Frankle et al, 2020a,b; Morcos et al, 2019). Moreover, pruning coefficients w.r.t. a dynamic, adaptive basis for $K \times K$ convolutional filters improves late rewinding even further (Wimmer et al, 2022).

Lee et al (2021) use scaled magnitudes to improve results for LTs at high sparsity levels. Their scale approximates the *best choice of layerwise sparsity ratios*. LTs as *equilibria of dynamical systems* are analyzed in Zhang et al (2021b). Theoretical guarantees for recovering sparse linear networks via LTs are described in Elesedy et al (2021).

Costs for finding LTs via iterative magnitude pruning can be reduced by using early stopping and low precision training for each pre-training iteration (You et al, 2020), sharing LTs for different datasets and optimizers (Morcos et al, 2019) or iteratively reducing the dataset together with the number of non-zero parameters (Zhang et al, 2021a). Also, the error of a LT from a given family of network configurations (*i.e.* ResNet with varying sparsity, width,

Table 2 Comparison of different methods for PaI. Here, $\Theta^{(0)} \in \mathbb{R}^D$ denotes the dense, random initialization of the model, $g^{(0)}$ the gradient and $H^{(0)}$ the Hessian of the loss at beginning of training. The function $\mathcal{R} : \mathbb{R}^D \rightarrow \mathbb{R}$, used for Tanaka et al (2020), can be chosen as any almost everywhere differentiable function.

Method	Criterion	One-shot	Pre-train $\psi^{(0)}$	Train $\vartheta^{(0)}$	Initialization of $\vartheta^{(0)}$
Lee et al (2019b)	$ \Theta^{(0)} \odot g^{(0)} $	✓	✗	✓	standard
Lee et al (2020)	$ \Theta^{(0)} \odot g^{(0)} $	✓	✗	✓	dynamic isometry
Hayou et al (2021)	$\mathbb{E} \Theta^{(0)} \odot g^{(0)} ^2$	✓	✗	✓	edge of chaos
de Jorge et al (2021)	$ \Theta^{(0)} \odot g^{(0)} $	✗	✗	✓	standard
Wang et al (2020a)	$-\Theta^{(0)} \odot H^{(0)}g^{(0)}$	✓	✗	✓	standard
Tanaka et al (2020)	$\Theta^{(0)} \odot \nabla_{\Theta^{(0)}} \mathcal{R}$	✗	✗	✓	standard
Wimmer et al (2021)	$\lambda \Theta^{(0)} \odot g^{(0)} + (1 - \lambda)(\Theta^{(0)} \odot H^{(0)}g^{(0)})$	✓	✗	✓	standard
Su et al (2020)	random (+ prior knowledge)	✓	✗	✓	standard
Patil and Dvornik (2021)	path norm	✗	✗	✓	standard
Lubana and Dick (2021)	$ \Theta^{(0)} \odot g^{(0)} \odot \Theta^{(0)} $	✓	✗	✓	standard
Lubana and Dick (2021)	$ \Theta^{(0)} \odot H^{(0)}g^{(0)} $	✓	✗	✓	standard
Zhang and Stadie (2020)	temporal Jacobian	✓	✗	✓	standard
Alizadeh et al (2022)	meta-gradient	✓	✓	✓	standard
Zhou et al (2019)	trainable prune probability	✗	✓	✗	standard
Ramanujan et al (2020)	trainable score s	✗	✓	✗	standard
Diffenderfer and Kaikhura (2021)	trainable score s	✗	✓	✗	binarization
Koster et al (2022)	trainable score s (incl. sign swap)	✗	✓	✗	binarization
Chen et al (2022)	$\Theta^{(0)} \odot \nabla_{\Theta^{(0)}} \mathcal{R}$ (+ trained sign swap)	✗	✓	✗	standard
Aladago and Torresani (2021)	trainable quality score	✗	✓	✓	$\vartheta_k^{(0)}$ out of fixed $\{w_k^{(1)}, \dots, w_k^{(n)}\}$

depth and number of used training examples) can be well estimated by knowing the performance of only a few trained networks from this network family (Rosenfeld et al, 2021). Despite their first applications on image classification, LTs have also shown to be successful in self-supervised learning (Chen et al, 2021a), natural language processing (Yu et al, 2020; Chen et al, 2020b), reinforcement learning tasks (Yu et al, 2020; Vischer et al, 2022), transfer learning (Soelen and Sheppard, 2019) and object recognition tasks like semantic segmentation or object detection (Girish et al, 2021). Moreover, Chen et al (2021b) propose methods to verify the ownership of LTs and hereby protect the rightful owner against intellectual property infringement.

LTs outperform randomly reinitializing sparse networks in the unstructured pruning case (Frankle and Carbin, 2018). Contrarily, for structured pruning there seems to be no difference between randomly reinitializing the sparse network and resetting the weights to their random initialization (Liu et al, 2019).

Closely related to the LTH with late rewinding, Renda et al (2020); Le and Hua (2021) show that fine-tuning a pruned network with a learning rate schedule rewound to earlier stages in training outperforms classical fine-tuning of sparse networks with small learning rates. Bai et al (2022) show that *arbitrary* randomly chosen pruning masks can lead to successful sparse models *if* the full network is used during training. For this, they extrude information of weights, which will be pruned eventually, into the sparse network during a pre-training step. Thus, Bai et al (2022) *is not* a DRT method.

5 Pruning at initialization

Overcoming the high cost for the train-prune-reset cycle(s) needed to find LTs is one of the main motivation to use pruning at initialization (PaI) (Lee et al, 2019b). With pruning at initialization we mean methods that start with a

randomly initialized network and do not perform any pre-training of the network’s weights to find the pruning transformation $\psi^{(0)}$. Furthermore for PaI, $\psi^{(t)} = \psi^{(0)}$ holds for all training iterations t . Different variants of PaI methods include one-shot pruning (Lee et al, 2019b, 2020; Wang et al, 2020a; Zhang and Stadie, 2020; Hayou et al, 2021; Wimmer et al, 2021, 2022; Alizadeh et al, 2022), iterative pruning (Tanaka et al, 2020; de Jorge et al, 2021; Verdenius et al, 2020; Patil and Dvovolis, 2021) and training the pruning transformation (Zhou et al, 2019; Ramanujan et al, 2020; Diffenderfer and Kailkhura, 2021; Aladago and Torresani, 2021). On top of that, there are methods that train the network after pruning and methods that do *not train* the non-zero parameters at all. Table 2 summarizes the PaI methods proposed in this Section.

5.1 PaI followed by training non-zero weights

First, we start with comparing different methods that train weights after the pruning step. We can group most of them into *gradient* based approaches and *information flow* based approaches. For a detailed comparison of the three popular PaI methods Lee et al (2019b); Tanaka et al (2020); Wang et al (2020a) we refer to Frankle et al (2021a). Moreover, Fischer and Burkholz (2022) compares them on generated tasks, where known and extremely sparse target networks are *planted* in a randomly initialized model. They show that current PaI methods fail to find those sparse models at extreme sparsity. However, the performance of other pruning methods than PaI is not evaluated and it remains an open question if they are able to find these planted subnetworks.

Gradient based approaches

Gradient based approaches (Lee et al, 2019b, 2020; Verdenius et al, 2020; de Jorge et al, 2021; Hayou et al, 2021) try to construct the sparse network which has the best influence in changing the loss function at the beginning of training, as presented in equations (10) and (11). However, it was shown that one-shot gradient based approaches have the problem of a vanishing gradient flow, if too many parameters are pruned (Lee et al, 2020; Wang et al, 2020a; Wimmer et al, 2020). Methods to overcome this are given by using an iterative approach (de Jorge et al, 2021; Verdenius et al, 2020) or use an initialization of the network, adjusted to the sparse network (Lee et al, 2020; Hayou et al, 2021). These adjusted initialization include so called *dynamic isometric* networks (Lee et al, 2020) and networks *at the edge of chaos* (Hayou et al, 2021).

Methods preserving information flow

Other PaI methods primarily focus on generating sparse networks with a sufficient information flow (Wang et al, 2020a; Tanaka et al, 2020; Zhang and Stadie, 2020; Patil and Dvovolis, 2021). For randomly initialized DNNs, the loss is no better than chance. Therefore, Wang et al (2020a) argue that “at the beginning of training, it is more important to preserve the training dynamics than the loss itself.” Consequently, Wang et al (2020a) try to find the sparse network with the highest gradient norm after pruning. They do so, by training

the weights with highest $-\Theta_i^{(0)}(H^{(0)}\nabla_{\Theta^{(0)}}\mathcal{L})_i$, where $H^{(0)}$ defines the Hessian of the loss function at initialization. Another approach is given by [Tanaka et al \(2020\)](#) and [Patil and Dovrolis \(2021\)](#), which try to maximize the path norm in the sparse network, see Section 3.4 *Conserving information flow* for more details. For RNNs or LSTMs, standard PaI methods do not work well ([Zhang and Stadie, 2020](#)). By pruning based on singular values of the temporal Jacobian, [Zhang and Stadie \(2020\)](#) are able to preserve weights that propagate a high amount of information through the network’s temporal depth.

Hybrid and other approaches

As shown in [Wimmer et al \(2021\)](#), only optimizing the sparse network to have the highest information flow possible does not lead to the best sparse architectures – even for high pruning rates. They conclude that information flow is a *necessary* condition for sparse training but not a *sufficient* one. Therefore, they combine gradient based and information flow based methods to get the best out of both worlds. With their approach, they improve gradient based PaI and PaI based on preserving information flow at one blow.

Another method guaranteeing a faithful information flow in the network and at the same time improving performance is given by *pruning in the interspace* ([Wimmer et al, 2022](#)). [Wimmer et al \(2022\)](#) represent $K \times K$ filters of a convolutional network in the *interspace* – a linear space spanned by an underlying filter basis. After pruning the filter’s coefficients, the filter basis is trained jointly with the non-zero coefficients. By adapting the interspace during training, networks tend to recover from low information flow. Furthermore, using interspace representations has shown to improve not only PaI for gradient based and information flow preserving methods, but also LTH, DST, freezing *and* training dense networks.

An orthogonal approach for guaranteeing high information flow while training only a sparse part of the network is proposed by [Price and Tanner \(2021\)](#). Random pruning is combined with a bypassing DCT in each layer. Since DCTs correspond to basis transformations with an orthonormal basis, the information of a layer’s input is maintained even if only a tiny fraction of parameters is trained. Therefore, results are improved tremendously for extreme pruning rates. Note, adding DCT improves performance for high p despite using a simple random selection of trained weights. As shown by [Price and Tanner \(2021\)](#), adding DCTs to each layer also improves DST methods for high p .

[Alizadeh et al \(2022\)](#) improve [Lee et al \(2019b\)](#) by modeling the effect of pruning on the loss function at training iteration $M > 0$. In contrast, [Lee et al \(2019b\)](#) only analyze the effect of pruning on the loss at initialization. To do so, [Alizadeh et al \(2022\)](#) compute a *meta-gradient* which is achieved by pre-training the dense network for M epochs. The meta-gradient is computed w.r.t. pruning mask. After pruning, the non-zero weights are reset to their initialization.




Fig. 7 Graphical explanation how [Malach et al \(2020\)](#) approximate a single target weight through random connections by adding a wide enough layer $l + 1/2$ between layers l and $l + 1$ and afterwards pruning unneeded connections.

[Lubana and Dick \(2021\)](#) theoretically compare *magnitude*, *gradient based* and *information flow preserving* PaI methods. Their analysis shows that magnitude based approaches lead to a rapid decrease in the training loss, thus converge fast. Furthermore, gradient based pruning conserves the loss function, removes the slowest changing parameters and preserves the first order dynamics of a model's evolution. They combine magnitude and gradient based pruning to get the best of both worlds which yields the pruning score $|\nabla_{\Theta^{(0)}} \mathcal{L}_i| \cdot |\Theta_i^{(0)}|^2$. Finally, information flow preserving pruning by using [Wang et al \(2020a\)](#)'s criterion to get the sparse model with maximal gradient norm removes the weights which maximally increase the loss function. However, this does not preserve the second order dynamics of the model. *Preserving* the gradient norm by training the weights with highest $|\Theta_i^{(0)}(H^{(0)} \nabla_{\Theta^{(0)}} \mathcal{L})_i|$ on the other hand also preserves the second order dynamics and shows improved results compared to ([Wang et al, 2020a](#)).

[Su et al \(2020\)](#); [Frankle et al \(2021a\)](#) show that the most popular PaI methods [Lee et al \(2019b\)](#); [Wang et al \(2020a\)](#); [Tanaka et al \(2020\)](#) do not significantly lose performance if positions of non-zero weights are shuffled randomly in each layer if the sparsity is not too extreme. Consequently, these methods do not appear to find the best sparse architecture, but rather well performing layerwise pruning rates for the given network architecture and global pruning rate. Using the knowledge of well performing PaI methods, [Su et al \(2020\)](#) show impressive results by randomly pruning weights. Hereby, they use layerwise pruning rates derived by observing the layerwise pruning rates of other well performing PaI methods. [Liu et al \(2022\)](#) also show that random PaI can reach competitive results to other PaI methods. They experimentally demonstrate that the gap between random PaI and dense training gets smaller if the underlying baseline network becomes bigger. Therefore, random PaI can provide a strong sparse training baseline, especially for large models.

5.2 PaI without training non-zero weights

[Mallya et al \(2018\)](#) show that a pre-trained network can be pruned for a specific task to have good performance even without fine-tuning. Inspired by that, several PaI methods showed that the same holds true for a randomly initialized network. Before we go into detail *how* pruning masks for random weights can be found, we will discuss theoretical works that cover the *universal*

approximator ability of pruning a randomly initialized network, also called *strong lottery ticket hypothesis* (Malach et al, 2020; Pensia et al, 2020).

Theoretical background

An intuition that big, randomly initialized networks contain well performing sparse subnetworks is given in Ramanujan et al (2020). Let $f_{\Theta^*} : \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a target network with $\Theta^* \in \mathbb{R}^{d^*}$. Further, let $f_{\Theta} : \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a network with randomly initialized $\Theta \in \mathbb{R}^D$ and $\tilde{\Theta}$ be a randomly chosen sparse projection of Θ with $\|\tilde{\Theta}\|_0 = d \ll D$ non-zero parameters. Consequently, $f_{\tilde{\Theta}}$ is an arbitrary d -sparse subnetwork of f_{Θ} . Assuming $D \gg d^* \approx d$, Ramanujan et al (2020) argue that the chance of $f_{\tilde{\Theta}}$ being a well approximator of f_{Θ^*} is small, but equal to $\delta > 0$. The big, random network has $\binom{D}{d}$ subnetworks with d non-zero parameters. Consequently, the chance of *all* sparse subnetworks $f_{\tilde{\Theta}}$ not being a good approximator for f_{Θ^*} is

$$(1 - \delta)^{\binom{D}{d}} \xrightarrow{D \rightarrow \infty} 0. \quad (14)$$

Thus, if the randomly initialized network f_{Θ} is chosen big enough, it will contain, with high chance, a well approximator for f_{Θ^*} with d non-zero weights.

This intuition is proven in Malach et al (2020) for MLPs by using a slightly different idea. Here, instead of making each layer in the randomly initialized network arbitrary wide, an intermediate layer $l + 1/2$ is added between each two layers l and $l + 1$ of the randomly initialized network, see Figure 7 right. These intermediate layers are made wide enough so that each weight w^* in the target network can be approximated well enough. It holds

$$w^* \cdot x = \underbrace{(+1)}_{\approx \text{sign}_+} \cdot \underbrace{\text{ReLU}(|w^*| \cdot x)}_{\approx w^*_+} + \underbrace{(-1)}_{\approx \text{sign}_-} \cdot \underbrace{\text{ReLU}(-|w^*| \cdot x)}_{\approx w^*_-}. \quad (15)$$

Therefore, for each weight w^* in the target network two paths are constructed. The first one approximates the positive part $|w^*|$ together with sign $+1$ and the other the negative part $-|w^*|$ together with sign -1 , see middle right of Figure 7. All remaining weights in the random network are pruned as shown in the middle left of Figure 7. As a consequence, Malach et al (2020) show that, with some assumptions on the target network, pruning big randomly initialized networks is a universal approximator. But, for each target weight, a polynomial number of intermediate weights have to be added. Follow up works Orseau et al (2020) and Pensia et al (2020) improve the parameter efficiency by shrinking the big, randomized network. They do so by using more than two paths to approximate a weight in the original network. By allowing the random values to be resampled, the width of the intermediate layer $l + 1/2$ can be reduced to 2 times the size of the original layer l (Chijiwa et al, 2021). Burkholz et al (2022) allow the randomly initialized network to be deeper than 2 times the target network. By approximating layers with combinations of univariate and multivariate linear functions, they are also able to approximate

Table 3 Comparison of different methods for dynamic sparse training.

Method	Sparse initialization	Sparsity		Pruning criterion	Regrow criterion	F_p
		Global	Adaptive			
Bellec et al (2018)	random	✗	✗	sign change	random	triggered by pruning
Mocanu et al (2018)	random	✗	✗	magnitude	random	hyperparameter
Mostafa and Wang (2019)	random	✓	✓	magnitude	random	hyperparameter
Dettmers and Zettlemoyer (2019)	magnitude	✗	✓	magnitude	gradient momentum	1 epoch
Evci et al (2020)	random	✗	✗	magnitude	gradient	hyperparameter

convolutional layers. Independent to the approach of [Burkholz et al \(2022\)](#), [da Cunha et al \(2022\)](#) extend the result of [Pensia et al \(2020\)](#) to CNNs by restricting the network’s input to be non-negative.

Methods to train $\psi^{(0)}$

As already mentioned before, training $\psi^{(0)} \in \{0, 1\}^{D \times d}$ requires optimizing in a discrete space. Consequently, different approaches are used to find $\psi^{(0)}$ for a randomly initialized network. In [Zhou et al \(2019\)](#), weights are pruned with probabilities modeled by Bernoulli samplers with corresponding trainable parameters. While helping during training, the stochasticity of this procedure may limit the performance at testing time ([Ramanujan et al, 2020](#)). The stochasticity is overcome in [Ramanujan et al \(2020\)](#) by the **edge-popup** algorithm. For each weight $\Theta_i^{(0)}$ in the network a corresponding pruning score s_i is trained. The weights with top- d scores are kept at their initial value, the remaining ones are pruned. The score s_i is then optimized with the help of the so called *straight through* estimator ([Bengio et al, 2013](#)) and the pruning transformation is updated in each iteration. [Ramanujan et al \(2020\)](#) show that a random Wide-ResNet50 ([Zagoruyko and Komodakis, 2016](#)) contains a sub-network with smaller size than a ResNet34 ([He et al, 2016](#)) but the same test accuracy on ImageNet ([Deng et al, 2009](#)). [Koster et al \(2022\)](#) and [Chen et al \(2022\)](#) independently show that allowing the randomly initialized weights to switch signs, improves the performance of **edge-popup**. [Chijiwa et al \(2021\)](#) improve **edge-popup** by allowing re-sampling of the un-trained weights. Binarizing the un-pruned, random weights can also be combined with the **edge-popup** algorithm ([Diffenderfer and Kailkhura, 2021](#)). Finally, [Aladago and Torresani \(2021\)](#) sample for each weight $\vartheta_i^{(0)}$ in a target network a set of n possible weights $w_i^{(1)}, \dots, w_i^{(n)}$ in a *hallucinated*, n times bigger network. This hallucinated network is fixed. During training, each possible weight $w_i^{(j)}$ has a corresponding *quality score* which is increased if the weight is a good choice for $\vartheta_i^{(0)}$ and decreased otherwise. In the end, the $w_i^{(j)}$ with the highest quality score is set as $\vartheta_i^{(0)}$ whereas the remaining hallucinated weights are discarded, *i.e.* pruned.

6 Dynamic sparse training

For high pruning rates, PaI is not able to perform equally well as classical pruning methods. One possible explanation for the performance gap is that a

randomly initialized network does not contain enough information to find a suitable sparse subnetwork that trains well. Adapting the pruning transformation during training or using network pre-training to find $\psi^{(0)}$ overcomes this lack of information. As shown in Section 4, finding LTs is expensive. Furthermore, it is not answered if resetting weights to their initial value can match late rewinding for big scale datasets. Dynamic sparse training performs well for high pruning rates while only needing one, fully sparse training. This is achieved by redistributing sparsity over the network in the course of training. By always fixing a given rate of parameters at zero, the number of trained weights is kept constant during training. Different DST methods are collected in Table 3.

6.1 Dynamic sparse training methods

Inspired by the rewiring of synaptic connectivity during the learning process in the human brain (Chambers and Rumpel, 2017), DEEP-R (Bellec et al, 2018) trains sparse DNNs while allowing the non-zero connections to rewire during training. To do so, pruning and rewiring is modeled as stochastic sampling of network configurations from a posterior. However, this procedure is computationally costly and challenging to apply to big networks and datasets (Dettmers and Zettlemoyer, 2019).

In Mocanu et al (2018), a sparse subnetwork with d non-zero parameters is chosen randomly at the beginning of training. Here, an initial pruning rate has to be defined separately for every layer. After each epoch, trained weights with the smallest magnitude are pruned layerwise with rate q and the same number of weights is regrown at random positions in this layer. Note, the pruning rate q used for *updating* $\psi^{(t)}$ is different from the *initial* pruning rate p . This procedure is done until training converges. Thus, in each epoch the same number of parameters is pruned and regrown which makes training hard to converge. Furthermore, the number of non-zero parameters needs to be defined for each layer before training and can not be adapted during training.

Dynamic sparse reparameterization (Mostafa and Wang, 2019) overcomes this problem by using magnitude pruning with an adaptive, global threshold. Furthermore, the number of regrown weights in each layer is adapted proportionally to the number of non-zero weights in that layer.

Regrowing weights not randomly, but based on their gradient's momentum was proposed by Dettmers and Zettlemoyer (2019). Here, not only the regrown weights are determined by their momentum, but also their number in each layer is.

RigL (Evcı et al, 2020) bypasses the need of Dettmers and Zettlemoyer (2019) to compute the dense gradient in each training iteration, and only regrows weights based on their actual gradient in the updating iteration of $\psi^{(t)}$, *i.e.* $t \in \{F_p, 2F_p, \dots\}$. Also, the number of pruned and regrown parameters is reduced by a cosine schedule to accelerate and improve convergence.

Overparameterized, densely trained DNNs in combination with SGD have shown good generalization abilities (Du et al, 2019; Li and Liang, 2018;

Brutzkus et al, 2018). Liu et al (2021b) showed that the good generalization ability of DST models can be explained by the so called *in-time-overparameterization*. Training of sparse DNNs can be viewed in the *space-time* manifold. To overparameterize DST models in this manifold, three properties must be fulfilled:

1. The dense baseline network has to be big enough.
2. Exploration of trained weights has to be guaranteed during training.
3. The training time has to be long enough so that the network can test enough sparse architectures in training.

If in-time-overparameterization is guaranteed for DST methods by these three criteria, Liu et al (2021b) show that sparse networks are able to outperform the dense, standard overparameterized ones. In-time-overparameterization can be achieved by either increasing the number of training epochs, or by reducing the batch size while keeping F_p constant. The latter leads to more updates of $\psi^{(t)}$ while not increasing the number of training epochs.

6.2 Closely related methods

We want to highlight that there exists more methods which are called *dynamic sparse training* in literature which do not fulfill our definition of it. We do not present these methods here in detail since they allow to update *all* weights of the network and only mask them out in the forward pass. Examples for such methods are Guo et al (2016); Ding et al (2019); Kusupati et al (2020); Sanh et al (2020); Liu et al (2020). Other dynamic training methods (Jayakumar et al, 2020; Schwarz et al, 2021; Zhou et al, 2021c,b) sample different subnetworks for each training iteration, update this subnetwork while keeping all un-trained parameters fixed at their previous position. In the next training iteration, a new subnetwork is sampled. Therefore, they can not store the network's parameters $\Theta^{(t)}$ in its reduced form $\vartheta^{(t)}$ for all $t \in \{0, 1, \dots, T\}$. Schwarz et al (2021) additionally exchange the standard weights $\Theta^{(t)}$ through a reparameterization using the power $\Theta^{(t)} = \phi^{(t)}|\phi^{(t)}|^{\alpha-1}$ with $\alpha > 1$. By doing so, weights close to 0 are unlikely to grow. As a consequence, the parameters will form a heavy tailed distribution at convergence. This improves results since the parameters are pruned based on their magnitude and heavy tailed distributions are highly compressible via magnitude pruning (Barsbey et al, 2021). Using this reparameterization might also improve other DRT methods proposed in this work. Peste et al (2021) analyze projections of the dense parameters on a sparse subspace via the *iterative hard thresholding* algorithm. Under some assumptions on the smoothness of the loss function and its optimum, Peste et al (2021) theoretically show convergence guarantees in this setting. Moreover, they propose a training scheme by iterating between sparse and dense training to approximate the iterative hard thresholding algorithm – the so called *alternating compressed/decompressed (AC/DC) training*.

Table 4 Comparison of different methods for freezing parts of a neural network. Here, $\Theta^{(0)}$ denotes the dense, random initialization of the network and $g^{(0)}$ the gradient of the loss function at beginning of training.

Method	frozen part	structured	criterion	additional projection
ELMs (Huang et al, 2004)	all except classifier	✓	—	✗
Hoffer et al (2018)	only classifier	✓	—	✗
Rosenfeld and Tsotsos (2019)	all except batch normalization	✓	—	✗
Li et al (2018)	unstructured	✗	random	orthogonal projection
Zhou et al (2019)	unstructured	✗	iterative magnitude (LTH)	✗
Wimmer et al (2020)	unstructured	✗	$ \Theta^{(0)} \odot g^{(0)} $	✗
Sung et al (2021)	unstructured	✗	$ \nabla_{\Theta^{(0)}} \log J_{\Theta^{(0)}} ^2$	✗
Rosenfeld and Tsotsos (2019)	structured	✓	random	✗

7 Freezing parts of a network

Contrarily to pruning, *freezing* parts of a randomly initialized network has attracted less interest in research in recent years. The main reason for this is that frozen, non-zero weights must be accounted for in the forward propagation. Thus, no (theoretical) speed up for inference can be obtained. In addition, frozen weights must be stored even after training, while zeros only require a small memory footprint. But, by using pseudorandom number generators for initializing the neural network, frozen weights can be recovered with a single 32bit integer on top of the memory cost for a pruned network (Wimmer et al, 2020). The proposed methods to freeze a network are summarized in Table 4.

7.1 Theoretical background for freezing

Saxe et al (2011) show that convolutional-pooling architectures can be inherently frequency selective while using random weights. In Giryes et al (2016), euclidean distances and angles between input data points are analyzed while propagating through a ReLU network with random i.i.d. Gaussian weights. With ReLU activation functions, each layer of the network shrinks euclidean distances between points inversely proportional to their euclidean angle. This means that points with a small initial angle migrate closer towards each other the deeper the network is. Assuming the data *behaves well*, meaning data points in the same class have a small angle and points from different classes have a bigger angle between another, random Gaussian networks can be seen as an *universal system that separates any data*. On the other hand, if the data is not perfect, training might be needed to overcome big intra-class angles or small inter-class angles to achieve good generalization.

Freezing parameters at their initial value was used in Li et al (2018) in order to measure the *intrinsic dimension* of the objective landscape. First, the dense network is optimized to generate the best possible solution. Then, the number of frozen parameters is gradually decreased, starting by freezing all parameters. The network's parameters are computed according to (6), with $\psi^{(t)} = \psi^{(0)}$ equaling a random, orthogonal projection. If a similar performance as the dense solution is reached, the number of non-frozen parameters determines the intrinsic dimension of the objective landscape. While using frozen weights only

as a tool to measure the dimension of a problem, [Li et al \(2018\)](#) showed that freezing parameters can lead to competitive results.

7.2 Freezing methods

We want to start with the so called *extreme learning machines* (ELMs) ([Huang et al, 2004, 2011](#); [Qing et al, 2020](#)). ELMs are MLPs, usually with one hidden layer. The parameters of the hidden node are frozen and usually initialized randomly. However, the classification layer is trained by a closed form solution of the least-square regression ([Huang et al, 2004](#)). ELMs are universal approximators if the number of hidden nodes is chosen big enough ([Huang et al, 2011](#)). In SOTA networks, ELMs can be used to substitute the classification layer ([Qing et al, 2020](#)). Closely related to ELMs are *random vector functional link neural networks* ([Pao and Takefuji, 1992](#); [Pao et al, 1994](#)) which, in addition, allow links between the input and output layer.

A completely orthogonal approach to ELMs is proposed in [Hoffer et al \(2018\)](#), where the classification layer is substituted by a random orthogonal matrix. For the classification layer, only a temperature parameter T is learnt. Experiments show that the random orthogonal layer with optimized T yields comparable results to a trainable classifier for modern architectures on CIFAR-10/100 ([Krizhevsky, 2012](#)) and ImageNet ([Deng et al, 2009](#)).

In [Zhou et al \(2019\)](#), pruning weights and freezing them at their initial values is compared in the LTH framework. They show that freezing usually performs better for a low number of trained parameters whereas pruning has better results if more parameters are trained. Finally, they show that a combination of pruning and freezing – depending whether a weight moved towards zero or away from zero during dense training – reaches the best results.

Freezing at initialization is used in [Wimmer et al \(2020\)](#) to overcome the problem of vanishing gradients for PaI. Similar to [Zhou et al \(2019\)](#) they show that freezing outperforms pruning if only few parameters are trained. For high freezing rates, freezing still guarantees a sufficient information flow in the sparsely trained network. On the other hand, if a higher number of parameters is trained, pruning also performs better than freezing in this setting. But, by using weight decay on the frozen parameters, [Wimmer et al \(2020\)](#) are able to get the best from both worlds, frozen weights at the beginning of training to ensure faithful information flow and sparse networks in the end of it, while improving both of them at the same time.

Inspired by transfer learning, [Sung et al \(2021\)](#) freeze large parts of the model to reduce the size of the newly learned part of the network. Frozen weights are chosen according to their importance for changing the network’s output, measured by the Fisher information matrix. They show that freezing parameters helps to reduce communication costs in distributed training as well as memory requirements for checkpointing networks during training. Unsurprisingly, they reach the best results if the freezing mask is allowed to change during training compared to keeping the freezing mask fixed.

Table 5 Performance of different methods for a ResNet50 trained on ImageNet. Results, except for the LTH experiment and Wimmer et al (2022) are derived from Evci et al (2020), Figure 2. Wimmer et al (2022) uses the pruning method from (Evci et al, 2020) combined with a trainable basis. The result for Frankle et al (2020a) is reported from Evci et al (2022) which is LTH with late rewinding.

Method	Category	Top-1-Acc	Training FLOPs	Testing FLOPs	Top-1-Acc	Training FLOPs	Testing FLOPs
Dense	—	76.8 ± 0.1	1×	1×	76.8 ± 0.1	1×	1×
		pruning rate $p = 0.8$		pruning rate $p = 0.9$			
Random	PaI	70.6 ± 0.1	0.23×	0.23×	65.8 ± 0.0	0.10×	0.10×
Lee et al (2019b)	PaI	72.0 ± 0.1	0.23×	0.23×	67.2 ± 0.1	0.10×	0.10×
Frankle et al (2020a)	LTH	75.8 ± 0.1					
Evci et al (2020)	DST	74.6 ± 0.1	0.23×	0.23×	72.0 ± 0.1	0.10×	0.10×
+ Wimmer et al (2022)	DST	76.0 ± 0.1			74.3 ± 0.1		
Evci et al (2020) ×5	DST	76.6 ± 0.1	1.14×	0.23×	75.7 ± 0.1	0.52×	0.10×
Mocanu et al (2018)	DST	72.9 ± 0.4	0.23×	0.23×	69.6 ± 0.2	0.10×	0.10×
Dettmers and Zettlemoyer (2019)	DST	75.2 ± 0.1	0.61×	0.42×	72.9 ± 0.1	0.50×	0.24×

Rosenfeld and Tsotsos (2019) freeze parameters in a structured way. They mainly focus on training only a fraction of filters (*i.e.* output channels) and determine the frozen parts randomly. In this setting, freezing outperforms pruning for almost all numbers of trained weights.

Furthermore, it is shown in Rosenfeld and Tsotsos (2019); Frankle et al (2021b) that freezing *all* weights except the trainable batch normalization (Ioffe and Szegedy, 2015) parameters leads to non-trivial performance.

8 Comparing and discussing different dimensionality reduced training methods

In this Section we discuss the different DRT approaches introduced in Sections 4, 5, 6 and 7. Figure 3 shows a high-level comparison between LTH/PaI, DST and freezing. Leaning on this structural comparison, we will now discuss different aspects of the methods.

8.1 Performance.

We start with the most important criterion, the *performance* of the methods. With performance we mean, test accuracy (or different metrics for tasks other than image classification) after training. This survey covers methods that reduce training cost. Therefore, performance needs to be brought in the context with the cost for the method. As baseline for comparing different methods, we will use our main cost measure – the number of trained (or equivalently stored) parameters.

For the performance evaluation, we use LTH with late rewinding. Note, this is *not* a real DRT method since training starts with a subset of $\Theta^{(t)}$ for a small $t > 0$ here. However, literature only reports results for modern networks and big scale tasks like ImageNet for LTH with late rewinding. As already

mentioned in Section 4, resetting weights to their initialization shows worse results than late rewinding which should be kept in mind.

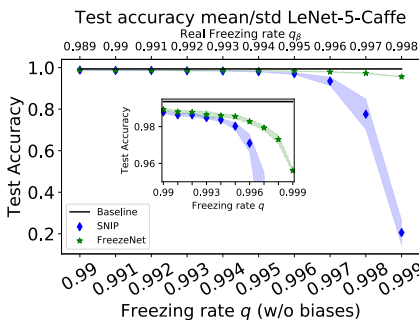


Fig. 8 (Figure 1 in Wimmer et al (2020)) Comparing the PaI method SNIP (Lee et al, 2019b) and the model with the same trained parameters but the un-trained parameters frozen at their initial value, FreezeNet (Wimmer et al, 2020) for a LeNet-5-Caffe (LeCun et al, 1998) on MNIST (LeCun et al, 1998).

twice the training time than a standard DST method. On the other hand, the reported LTH result from Evci et al (2022) does *not* use the standard way to find LTs (Frankle et al, 2020a), but *gradual magnitude pruning* (Gale et al, 2019) to find the sparse subnetwork. Thus, by using the approach to iteratively train-prune-rewind the weights (Frankle et al, 2020a), LTs performance could be improved further. Moreover, Table 5 shows that pruning coefficients w.r.t. adaptive bases (Wimmer et al, 2022) for $K \times K$ filters improves performance of SOTA methods further.

PaI’s inferior performance is not a surprise since PaI methods are clearly the less sophisticated ones. As we will see in the following discussions, this simplicity will reduce costs at other levels, like *training time* or *hyperparameter tuning*. Our first observation therefore is, that pre-training the network to find a well trainable sparse architecture (LT) or adapt the sparse architecture during training (DST) improves results compared to finding a sparse architecture *ad-hoc* at the beginning of training.

Freezing the parameters is compared with pruning them in the PaI setting (Wimmer et al, 2020) and the LTH setup (Zhou et al, 2019). Usually, freezing leads to slightly worse results than pruning for a higher number of trained parameters and better results for fewer trained parameters. Figure 8 shows a comparison between pruning and freezing weights for a low number of trained weights in the PaI setting. Price and Tanner (2021) show that the same holds true if the random dense layer is replaced by a dense DCT. Using DCTs has the advantage that they are cheap to compute.

It was shown and discussed in many works that LTs have better results than PaI for high pruning rates (Wang et al, 2020a; Frankle et al, 2021a). Also, DST improves results compared to PaI (Dettmers and Zettlemoyer, 2019; Evci et al, 2020). As an example, Table 5 compares PaI, LTH and DST for a ResNet50 (He et al, 2015) on ImageNet (Deng et al, 2009). Table 5 shows that LTs approximately reach the same performance than DST. Both of them outperform PaI. But, if the training time spent on DST methods is doubled, DST exceeds LT (Liu et al, 2021b), see also discussion in Section 6. The doubled training time for DST is a fair comparison, since LTs need at least

8.2 Storage cost

Techniques like the compressed sparse row format (Tinney and Walker, 1967) are used to store sparse/frozen networks after training. By using an initialization derived from a pseudorandom number generator,

$$\text{memory}(\text{freeze}) = \text{memory}(\text{prune}) + 32\text{bit} \approx \text{memory}(\text{prune}) \quad (16)$$

holds. Consequently, pruning and freezing have the same memory cost in this case. But, if no pseudorandom number generator can be used, pruning has much lower memory requirements than freezing.

Also, Sections 4 - 7 propose methods that do not use a simple pruning/freezing transformation but also an additional linear transformation to embed the small, trainable parameters $\vartheta \in \mathbb{R}^d$ into the bigger space \mathbb{R}^D . Examples for this is pruning coupled with an additional basis transformation with shared diagonal blocks (Wimmer et al, 2022) or a random, orthogonal projection (Li et al, 2018). The cost for storing these embeddings is neglectable for Wimmer et al (2022). The same holds for Li et al (2018) if the random transformation is created with a pseudorandom number generator.

8.3 Training cost

In the following, we will compare the *training costs* for different methods which mainly are measured by the *training time* for one sparse model and the number of tunable *hyperparameters* for training which implicitly determines the number of training runs needed overall. On top of that, some methods induce additional gradient computations before or during training which will be discussed in the end.

8.3.1 Training time

We assume that all methods, DST, LTs, PaI and freezing train the final sparse model for T iterations. Here, T is the number of training iterations used for the dense network to (i) converge and (ii) have good generalization ability. Consequently, DST, PaI and freezing have approximately the same training time for one model. LTs on the other hand need, as discussed in Section 4, at least twice the number of training iterations than a standard training. By using iterative train-prune-reset cycles, the number of training iterations can easily reach 10 – 20× the standard number. Thus, LTs need a massive amount of pre-training for the pruning transformation $\psi^{(0)}$.

As shown in Liu et al (2021b) and discussed in Section 6, DST massively profit from increasing the training time.

8.3.2 Hyperparameters

All proposed methods induce at least one hyperparameter more than training the corresponding dense model. This hyperparameter is given by the *number of trainable parameters*.⁵

Determining the number of trained parameters

PaI, freezing and LT determine *one* global pruning/freezing rate. DST methods on the other hand often need to specify the rate of trained parameters for each layer separately. If the layerwise pruning rates can be adapted dynamically during training (Dettmers and Zettlemoyer, 2019), their initial choice has shown to be not too important and can be set constant for all layers. Other works like Mocanu et al (2018); Evcı et al (2020); Liu et al (2021b) heuristically determine layerwise pruning rates by an *Erdős-Rényi model* (Erdős and Rényi, 1959).

Additional hyperparameters for PaI

Besides the pruning rate, PaI methods introduce hyperparameters for computing the pruning transformation $\psi^{(0)}$. For Lee et al (2019b); Wang et al (2020a); Verdenius et al (2020); de Jorge et al (2021) this is the number of data batches, needed to compute the *pruning score* of each parameter. The iterative approaches Tanaka et al (2020); Verdenius et al (2020); de Jorge et al (2021) need to determine the number of conducted pruning iterations. However, by using a high number of data batches and many pruning iterations, these parameters are not required to be fine-tuned further (Wang et al, 2020a; Tanaka et al, 2020). A special case in this setting is Wimmer et al (2021), interpolating between the two pruning methods Lee et al (2019b) and Wang et al (2020a). Consequently, they need to tune one hyperparameter to balance between these two methods.

Additional hyperparameters for LTs

As shown in Frankle et al (2020a), resetting weights not to their initial value but a value early on in training leads to better results. Consequently, a well performing *rewinding* iteration has to be found. A proper experimental analysis for various datasets and models is given in Frankle et al (2020a). Furthermore, if iterative pruning is used, the number of pruned weights in each iteration has to be determined. Usually, 20% of the non-zero weights are removed in each iteration (Frankle and Carbin, 2018).

Additional hyperparameters for freezing

Freezing is closely related to either PaI (Wimmer et al, 2020), LTs (Zhou et al, 2019) or random pruning (Rosenfeld and Tsotsos, 2019). Since freezing

⁵In pruning literature, many methods determine the number of trainable parameters not directly but indirect by choosing an associated hyperparameter as for example weighting the ℓ_1 regularization. But, all proposed methods in this work choose the number of trained parameters directly.

does not introduce additional hyperparameters, the same number of additional hyperparameters as for the corresponding pruning methods are needed. Note, to achieve good results with randomly frozen weights, layerwise freezing rates have to be determined which need to be fine-tuned accordingly.

Additional hyperparameters for DST and costs for updating $\psi^{(t)}$

Compared to PaI, DST has the advantage that the pruning transformation is not fixed at $\psi^{(0)}$. This leads to better performance after training while using the same number of trained parameters and training iterations. But this also results in costs for determining $\psi^{(t)}$. Besides computing $\psi^{(t)}$, the *update frequency* F_p and the *update pruning rate* $q_l^{(t)}$ have to be found. The update pruning rate $q_l^{(t)}$ is defined as the ratio of formerly non-zero parameters which are newly pruned in layer l if $\psi^{(t)}$ is updated in iteration t . Note, the update pruning rate might vary between different update iterations of the transformation and also between layers. Also, the *regrowing* rate needs to be determined since it can be different from the update pruning rate (Dettmers and Zettlemoyer, 2019).

Finally, there are the actual costs for updating $\psi^{(t)}$ itself. In all considered methods in Section 6, weights are dried according to having the smallest magnitude – magnitude based pruning. For this, the d un-pruned coefficients have to be sorted. After setting some weights to zero, the same number of weights has to be flagged as trainable again. The costs for regrowing weights can almost be for free by using random regrowing of weights (Mocanu et al, 2018; Mostafa and Wang, 2019). On the other hand, Evci et al (2020) require the gradient of the *whole* network at the update iterations $t \in \{F_p, 2F_p, \dots\}$ to determine $\psi^{(t)}$. Dettmers and Zettlemoyer (2019) even need these gradients in *each* training iteration in order to update a momentum parameter which also requires an additional weighting hyperparameter. A temperature parameter is needed in Bellec et al (2018) to model a *random walk* in the parameter space to explore new weights.

8.3.3 Gradient computations and backward pass.

All proposed methods update only sparse parts of the weights via backpropagation. But, some of them additionally require the computation of the dense gradient for pre-training, determining $\psi^{(0)}$ or to update $\psi^{(t)}$.

First of all, LTs need to *train* the dense network in order to find $\psi^{(0)}$. This can limit the size of the biggest underlying dense network that can be used. For iterative LTH, networks with sparsity $p_k = 1 - 0.8^k$ have to be trained additionally for all k with $p_k > p$.

PaI methods need the computation of a dense gradient before training (Lee et al, 2019b; Tanaka et al, 2020; Verdenius et al, 2020; de Jorge et al, 2021). Even a Hessian-vector product has to be calculated for the method Wang et al (2020a). But after having found $\psi^{(0)}$, PaI methods are trained with a fixed, sparse architecture without the need to compute a dense gradient anymore.

DST methods might not need a dense gradient at all, if regrown weights are found by a random selection (Mocanu et al, 2018; Bellec et al, 2018; Mostafa

and Wang, 2019). As mentioned above, Evci et al (2020) compute the dense gradient at each update step of the pruning transformation $t \in \{F_p, 2F_p, \dots\}$. The most extreme case is given for Dettmers and Zettlemoyer (2019) which need to compute the dense gradient in each iteration.

Finally there is a difference between gradient computations for pruned and frozen models. Pruned networks compute gradients of activation maps with sparse weight tensors. Frozen networks on the other hand compute gradients of activation maps with the help of the frozen weights, *i.e.* require dense computations, see Figure 4 right side. However, frozen weights do not need to be updated. Therefore, it is enough to compute only a sparse part of the weights' gradients. In summary, freezing also reduces computations in the backward pass, but not as much as pruning does.

8.4 Forward Pass.

In Section 8.3.3, gradient computations and the backward pass for the different methods are discussed. Gradients only have to be computed during training whereas the forward pass is needed for both, training and inference. Therefore, we discuss the forward pass in a standalone Section. Still, this discussion should also be seen as a part of the training costs, Section 8.3.

During and after the sparse training, LTH, DST and PaI all have the same cost for inference. This statement is not completely true, since different methods might create varying sparsity distributions for the same global pruning rate. This can result in different FLOP costs for inferring the sparse networks, see for example Table 5. However, analyzing sparsity distributions obtained by different pruning methods is beyond the scope of this work.

If parts of a network are frozen during training, *all* parameters participate in the forward pass. Of course, this also holds true for inference time. Consequently, the computational cost for inference can not be reduced and is equal to the densely trained network.

8.5 Summary

In summary, we see that freezing parameters instead of pruning them leads to the same memory requirements and better results for training a low number of parameters. However, if more parameters are trained, pruning leads to equal or better performance. For pruning, the sparsity of the parameters can reduce the number of required computations for evaluating the network if the used soft- and hardware supports sparse computations. Thus, freezing seems to be a good option if the model size is a bottleneck, *i.e.* only a small part of the network is trained, whereas pruning is preferably used otherwise. Furthermore, the computational overhead induced by freezing layers can be reduced by using dense DCTs instead of dense frozen layers.

For pruning, there is a trade-off between simplicity of the method and performance after training. As shown, DST and LTs have approximately the same generalization ability and are superior to PaI for the same number of non-zero

parameters. However, PaI guarantees a fixed sparse architecture throughout training, determined by only a few dense gradient computations. Furthermore, PaI does not require extensive and costly hyperparameter tuning. Finding LTs requires expensive pre-training of the dense network. Also, LTs at initialization often show unsatisfactory results for modern network architectures and big scale datasets. Resetting weights to an early phase in training is required to achieve good performance. As shown, DST methods need more hyperparameters than PaI and LTH. Further, DST methods might need to regularly compute the full gradient of the network and increase the training time in order to reach their best performance.

As mentioned, presented results for LTs use a small pre-training step for the sparse initialization. Therefore, we conclude that DST methods achieve the best overall performance for completely sparse training. Moreover, using adaptive bases for the $K \times K$ filters of convolutional layers further improves the proposed pruning and freezing approaches.

9 Conclusions

In this work we have introduced a general framework to describe the training of DNNs with reduced dimensionality (DRT). The proposed methods are dominated by pruning neural networks, but we also discussed freezing parts of the network at its random initialization. The methods are categorized in pruning at initialization (PaI), lottery tickets (LTs), dynamic sparse training (DST) and freezing. SOTA methods for each criterion are presented. Furthermore, the different approaches are compared afterwards.

We first discussed that pruning leads to better results than freezing if more parameters are optimized whereas for a low number of trained weights, freezing performs better than pruning. Furthermore, LTs and DST perform better than PaI, while PaI contains the easiest and least expensive methods. For LTs, many trainings, including training the dense model, are needed to find the sparse architecture. Also, LTs achieve their best results if weights are reset to an early phase in training which does not yield a complete sparse training. DST methods usually need more hyperparameters to be tuned than LTs and PaI. All proposed DRT methods can be further enhanced by representing convolutional filters with an adaptive representation instead of the standard, spatial one.

Altogether, finding the best DRT method for a specific setup is a trade-off between available training time and performance of the final network. The training time is mostly influenced by the number of trainings required to find the sparse architecture, ranging from 0 (PaI, DST and freezing) to more than $20\times$ (LTs), and the number of individual trainings required to tune hyperparameters. Also, the need to compute dense gradients (for some DST and PaI methods) or to pre-train the dense network (LTs) has to be considered if a DRT method is chosen. As discussed, frozen models have similar computational costs for inference as densely trained networks. Consequently, if the sparsity together with the used soft- and hardware can actually reduce

computations, pruning should be preferred to freezing, or frozen parameters should be replaced by information preserving transformations which are cheap to compute, as for example the DCT.

We want to end the survey by stating relevant future research topics concerning DRT. For PaI it is important to find methods that reduce, or at best close, the gap between current PaI methods and the better performing DST/LTs methods. On the other hand, DST methods which adapt important hyperparameters automatically during one (or only a few) training runs, instead of using many of them for manual hyperparameter tuning, would help to reduce the costs for DST. Finding LTs is expensive due to the need of at least one, but for good results much more, train-prune-reset cycle(s). As discussed, there is already research done to make LTs less expensive by using low precision training (You et al, 2020; Diffenderfer and Kailkhura, 2021) or transferring the LTs between datasets (Morcos et al, 2019). But, performance of such derived LTs often is worse than using a LT generated for this task. Improving the transferability of LTs would significantly reduce the cost for sparse training. Finally, freezing networks performs worse than pruning if many parameters are trained. Analyzing the performance gap between freezing and pruning could help to better understand the role of random initializations.

Statements and Declarations

Conflicts of Interest

During preparation of this manuscript, all authors were employed by the Robert Bosch GmbH. Also, Alexandru Paul Condurache and Paul Wimmer were part of the Institute for Signal Processing of the University of Lübeck.

References

- Aladago MM, Torresani L (2021) Slot machines: Discovering winning combinations of random weights in neural networks. In: Proceedings of the 38th International Conference on Machine Learning
- Alizadeh M, Tailor SA, Zintgraf LM, et al (2022) Prospect pruning: Finding trainable weights at initialization using meta-gradients. In: 10th International Conference on Learning Representations
- Amodei D, Hernandez D, Sastry G, et al (2018 [Online]) Ai and compute. OpenAI Blog, URL <https://openai.com/blog/ai-and-compute/>, last accessed: 11/16/2022
- Anwar S, Hwang K, Sung W (2017) Structured pruning of deep convolutional neural networks. *ACM Journal on Emerging Technologies in Computing Systems* 13(3):1–18

- Arora S, Ge R, Neyshabur B, et al (2018) Stronger generalization bounds for deep nets via a compression approach. In: Proceedings of the 35th International Conference on Machine Learning
- Arora S, Du SS, Hu W, et al (2019) On exact computation with an infinitely wide neural net. In: Advances in Neural Information Processing Systems 32
- Bai Y, Wang H, TAO Z, et al (2022) Dual lottery ticket hypothesis. In: 10th International Conference on Learning Representations
- Barsbey M, Sefidgaran M, Erdogdu MA, et al (2021) Heavy tails in SGD and compressibility of overparametrized neural networks. In: Advances in Neural Information Processing Systems 34
- Bartoldson B, Morcos A, Barbu A, et al (2020) The generalization-stability tradeoff in neural network pruning. In: Advances in Neural Information Processing Systems 33
- Bellec G, Kappel D, Maass W, et al (2018) Deep rewiring: Training very sparse deep networks. In: 6th International Conference on Learning Representations
- Bengio Y, Léonard N, Courville AC (2013) Estimating or propagating gradients through stochastic neurons for conditional computation. CoRR abs/1308.3432. URL <http://arxiv.org/abs/1308.3432>, last accessed: 10/31/2022
- Blalock DW, Ortiz JJG, Frankle J, et al (2020) What is the state of neural network pruning? In: Proceedings of Machine Learning and Systems 2
- Brutzkus A, Globerson A, Malach E, et al (2018) SGD learns overparameterized networks that provably generalize on linearly separable data. In: 6th International Conference on Learning Representations
- Burkholz R, Laha N, Mukherjee R, et al (2022) On the existence of universal lottery tickets. In: 10th International Conference on Learning Representations
- Chambers AR, Rumpel S (2017) A stable brain from unstable components: Emerging concepts and implications for neural computation. *Neuroscience* 357:172–184
- Chen J, Chen S, Pan SJ (2020a) Storage efficient and dynamic flexible runtime channel pruning via deep reinforcement learning. In: Advances in Neural Information Processing Systems 33

- Chen T, Frankle J, Chang S, et al (2020b) The lottery ticket hypothesis for pre-trained bert networks. In: *Advances in Neural Information Processing Systems* 33
- Chen T, Frankle J, Chang S, et al (2021a) The lottery tickets hypothesis for supervised and self-supervised pre-training in computer vision models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*
- Chen W, Wilson J, Tyree S, et al (2015) Compressing neural networks with the hashing trick. In: *Proceedings of the 32nd International Conference on Machine Learning*
- Chen X, Chen T, Zhang Z, et al (2021b) You are caught stealing my winning lottery ticket! making a lottery ticket claim its ownership. In: *Advances in Neural Information Processing Systems* 34
- Chen X, Zhang J, Wang Z (2022) Peek-a-boo: What (more) is disguised in a randomly weighted neural network, and how to find it efficiently. In: *10th International Conference on Learning Representations*
- Chijiwa D, Yamaguchi S, Ida Y, et al (2021) Pruning randomly initialized neural networks with iterative randomization. In: *Advances in Neural Information Processing Systems* 34
- Courbariaux M, Bengio Y, David JP (2015) Binaryconnect: Training deep neural networks with binary weights during propagations. In: *Advances in Neural Information Processing Systems* 28
- da Cunha A, Natale E, Viennot L (2022) Proving the lottery ticket hypothesis for convolutional neural networks. In: *10th International Conference on Learning Representations*
- Deng J, Dong W, Socher R, et al (2009) Imagenet: A large-scale hierarchical image database. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*
- Denton E, Zaremba W, Bruna J, et al (2014) Exploiting linear structure within convolutional networks for efficient evaluation. In: *Advances in Neural Information Processing Systems* 27
- Dettmers T, Zettlemoyer L (2019) Sparse networks from scratch: Faster training without losing performance. CoRR abs/1907.04840v2. URL <http://arxiv.org/abs/1907.04840v2>, last accessed: 10/02/2022
- Diffenderfer J, Kailkhura B (2021) Multi-prize lottery ticket hypothesis: Finding accurate binary neural networks by pruning a randomly weighted

- network. In: 9th International Conference on Learning Representations
- Diffenderfer J, Bartoldson BR, Chaganti S, et al (2021) A winning hand: Compressing deep networks can improve out-of-distribution robustness. In: *Advances in Neural Information Processing Systems* 34
- Ding X, Ding G, Zhou X, et al (2019) Global sparse momentum sgd for pruning very deep neural networks. In: *Advances in Neural Information Processing Systems* 32
- Dosovitskiy A, Beyer L, Kolesnikov A, et al (2021) An image is worth 16x16 words: Transformers for image recognition at scale. In: 9th International Conference on Learning Representations
- Du SS, Zhai X, Póczos B, et al (2019) Gradient descent provably optimizes over-parameterized neural networks. In: 7th International Conference on Learning Representations
- Duchi J, Hazan E, Singer Y (2011) Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* 12(61):2121–2159
- Elesedy B, Kanade V, Teh YW (2021) Lottery tickets in linear models: An analysis of iterative magnitude pruning. In: *Sparsity in Neural Networks Workshop*
- Elsen E, Dukhan M, Gale T, et al (2020) Fast sparse convnets. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*
- Erdős P, Rényi A (1959) On random graphs I. *Publicationes Mathematicae Debrecen* 6:290–297
- Evcı U, Gale T, Menick J, et al (2020) Rigging the lottery: Making all tickets winners. In: *Proceedings of the 37th International Conference on Machine Learning*
- Evcı U, Dauphin Y, Ioannou Y, et al (2022) Gradient flow in sparse neural networks and how lottery tickets win. In: *Proceedings of the AAAI Conference on Artificial Intelligence*
- Fischer J, Burkholz R (2022) Plant 'n' seek: Can you find the winning ticket? In: 10th International Conference on Learning Representations
- Frankle J, Carbin M (2018) The lottery ticket hypothesis: Finding sparse, trainable neural networks. In: 6th International Conference on Learning Representations

- Frankle J, Dziugaite GK, Roy D, et al (2020a) Linear mode connectivity and the lottery ticket hypothesis. In: Proceedings of the 37th International Conference on Machine Learning
- Frankle J, Schwab DJ, Morcos AS (2020b) The early phase of neural network training. In: 8th International Conference on Learning Representations
- Frankle J, Dziugaite GK, Roy D, et al (2021a) Pruning neural networks at initialization: Why are we missing the mark? In: 9th International Conference on Learning Representations
- Frankle J, Schwab DJ, Morcos AS (2021b) Training batchnorm and only batchnorm: On the expressive power of random features in CNNs. In: 9th International Conference on Learning Representations
- Gale T, Elsen E, Hooker S (2019) The state of sparsity in deep neural networks. In: 36th International Conference on Machine Learning Joint Workshop on On-Device Machine Learning & Compact Deep Neural Network Representations (ODML-CDNNR)
- Gale T, Zaharia M, Young C, et al (2020) Sparse gpu kernels for deep learning. In: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis
- Gebhart T, Saxena U, Schrater P (2021) A unified paths perspective for pruning at initialization. CoRR abs/2101.10552. URL <https://arxiv.org/abs/2101.10552>, last accessed: 11/19/2022
- Girish S, Maiya SR, Gupta K, et al (2021) The lottery ticket hypothesis for object recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition
- Giryes R, Sapiro G, Bronstein AM (2016) Deep neural networks with random gaussian weights: A universal classification strategy? IEEE Trans Signal Process 64(13):3444–3457
- Glorot X, Bengio Y (2010) Understanding the difficulty of training deep feedforward neural networks. In: Proceedings of the 13th International Conference on Artificial Intelligence and Statistics
- Guo Y, Yao A, Chen Y (2016) Dynamic network surgery for efficient dnns. In: Advances in Neural Information Processing Systems 29
- Gustafson JL (2011) Moore’s law. In: Encyclopedia of Parallel Computing, pp 1177–1184

- Han S, Pool J, Tran J, et al (2015) Learning both weights and connections for efficient neural network. In: *Advances in Neural Information Processing Systems* 28
- Han S, Liu X, Mao H, et al (2016) Eie: Efficient inference engine on compressed deep neural network. *ACM SIGARCH Computer Architecture News* 44(3):243–254
- Hanin B, Rolnick D (2018) How to start training: The effect of initialization and architecture. In: *Advances in Neural Information Processing Systems* 31
- Hayou S, Ton JF, Doucet A, et al (2021) Robust pruning at initialization. In: *9th International Conference on Learning Representations*
- He K, Zhang X, Ren S, et al (2015) Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *IEEE International Conference on Computer Vision*
- He K, Zhang X, Ren S, et al (2016) Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*
- Hoffer E, Hubara I, Soudry D (2018) Fix your classifier: the marginal value of training the last weight layer. In: *6th International Conference on Learning Representations*
- Holmes C, Zhang M, He Y, et al (2021) NxMTransformer: Semi-structured sparsification for natural language understanding via ADMM. In: *Advances in Neural Information Processing Systems* 34
- Huang GB, Zhu QY, Siew CK (2004) Extreme learning machine: a new learning scheme of feedforward neural networks. In: *IEEE International Joint Conference on Neural Networks* 2
- Huang GB, Wang DH, Lan Y (2011) Extreme learning machines: a survey. *International journal of machine learning and cybernetics* 2(2):107–122
- Huang Z, Wang N (2018) Data-driven sparse structure selection for deep neural networks. In: *Proceedings of the European Conference on Computer Vision*
- Hubara I, Chmiel B, Island M, et al (2021) Accelerated sparse neural training: A provable and efficient method to find n:m transposable masks. In: *Advances in Neural Information Processing Systems* 34
- Ioffe S, Szegedy C (2015) Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *Proceedings of the 32nd International Conference on Machine Learning*

- Jacob B, Kligys S, Chen B, et al (2018) Quantization and training of neural networks for efficient integer-arithmetic-only inference. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition
- Jacot A, Hongler C, Gabriel F (2018) Neural tangent kernel: Convergence and generalization in neural networks. In: Advances in Neural Information Processing Systems 31
- Janowsky SA (1989) Pruning versus clipping in neural networks. *Physical Review A* 39:6600–6603
- Jayakumar S, Pascanu R, Rae J, et al (2020) Top-kast: Top-k always sparse training. In: Advances in Neural Information Processing Systems 33
- Joo D, Yi E, Baek S, et al (2021) Linearly replaceable filters for deep network channel pruning. In: Proceedings of the AAAI Conference on Artificial Intelligence
- de Jorge P, Sanyal A, Behl H, et al (2021) Progressive skeletonization: Trimming more fat from a network at initialization. In: 9th International Conference on Learning Representations
- Karnin ED (1990) A simple procedure for pruning back-propagation trained neural networks. *IEEE Transactions on Neural Networks* 1(2):239–242
- Kingma DP, Ba J (2015) Adam: A method for stochastic optimization. In: 3rd International Conference on Learning Representations
- Kolesnikov A, Beyer L, Zhai X, et al (2020) Big transfer (bit): General visual representation learning. In: Proceedings of the European Conference on Computer Vision
- Koster N, Grothe O, Rettinger A (2022) Signing the supermask: Keep, hide, invert. In: 10th International Conference on Learning Representations
- Krizhevsky A (2012) Learning multiple layers of features from tiny images. University of Toronto URL: <http://www.cs.toronto.edu/~kriz/cifar.html>, last accessed: 05/13/2022
- Kusupati A, Ramanujan V, Somani R, et al (2020) Soft threshold weight reparameterization for learnable sparsity. In: Proceedings of the 37th International Conference on Machine Learning
- Le DH, Hua BS (2021) Network pruning that matters: A case study on retraining variants. In: 9th International Conference on Learning Representations
- Lebedev V, Ganin Y, Rakhuba M, et al (2015) Speeding-up convolutional neural networks using fine-tuned cp-decomposition. In: 3rd International

Conference on Learning Representations

- LeCun Y, Denker JS, Solla SA (1990) Optimal brain damage. In: *Advances in Neural Information Processing Systems 2*
- LeCun Y, Bottou L, Bengio Y, et al (1998) Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2324
- Lee J, Xiao L, Schoenholz S, et al (2019a) Wide neural networks of any depth evolve as linear models under gradient descent. In: *Advances in Neural Information Processing Systems 32*
- Lee J, Park S, Mo S, et al (2021) Layer-adaptive sparsity for the magnitude-based pruning. In: *9th International Conference on Learning Representations*
- Lee N, Ajanthan T, Torr PH (2019b) SNIP: Single-shot network pruning based on connection sensitivity. In: *7th International Conference on Learning Representations*
- Lee N, Ajanthan T, Gould S, et al (2020) A signal propagation perspective for pruning neural networks at initialization. In: *8th International Conference on Learning Representations*
- Li C, Farkhoor H, Liu R, et al (2018) Measuring the intrinsic dimension of objective landscapes. In: *6th International Conference on Learning Representations*
- Li H, Kadav A, Durdanovic I, et al (2017) Pruning filters for efficient convnets. In: *5th International Conference on Learning Representations*
- Li R, Wang Y, Liang F, et al (2019) Fully quantized network for object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*
- Li Y, Liang Y (2018) Learning overparameterized neural networks via stochastic gradient descent on structured data. In: *Advances in Neural Information Processing Systems 31*
- Liu B, Wang M, Foroosh H, et al (2015) Sparse convolutional neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*
- Liu J, Xu Z, Shi R, et al (2020) Dynamic sparse training: Find efficient sparse network from scratch with trainable masked layers. In: *8th International Conference on Learning Representations*

- Liu S, Mocanu DC, Matavalam ARR, et al (2021a) Sparse evolutionary deep learning with over one million artificial neurons on commodity hardware. *Neural Comput Appl* 33(7):2589–2604
- Liu S, Yin L, Mocanu DC, et al (2021b) Do we actually need dense over-parameterization? In-time over-parameterization in sparse training. In: *Proceedings of the 38th International Conference on Machine Learning*
- Liu S, Chen T, Chen X, et al (2022) The unreasonable effectiveness of random pruning: Return of the most naive baseline for sparse training. In: *10th International Conference on Learning Representations*
- Liu T, Zenke F (2020) Finding trainable sparse networks through neural tangent transfer. In: *Proceedings of the 37th International Conference on Machine Learning*
- Liu Z, Sun M, Zhou T, et al (2019) Rethinking the value of network pruning. In: *7th International Conference on Learning Representations*
- Lubana ES, Dick R (2021) A gradient flow framework for analyzing network pruning. In: *9th International Conference on Learning Representations*
- Mahajan D, Girshick R, Ramanathan V, et al (2018) Exploring the limits of weakly supervised pretraining. In: *Proceedings of the European Conference on Computer Vision*
- Malach E, Yehudai G, Shalev-Schwartz S, et al (2020) Proving the lottery ticket hypothesis: Pruning is all you need. In: *Proceedings of the 37th International Conference on Machine Learning*
- Mallya A, Davis D, Lazebnik S (2018) Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In: *Proceedings of the European Conference on Computer Vision*
- Mao H, Han S, Pool J, et al (2017) Exploring the granularity of sparsity in convolutional neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*
- Martens J (2010) Deep learning via hessian-free optimization. In: *Proceedings of the 27th International Conference on Machine Learning*
- Mezić I (2005) Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics* 41(1):309–325
- Mocanu D, Mocanu E, Stone P, et al (2018) Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature Communications* 9(1):2383

- Morcos A, Yu H, Paganini M, et al (2019) One ticket to win them all: generalizing lottery ticket initializations across datasets and optimizers. In: *Advances in Neural Information Processing Systems* 32
- Mostafa H, Wang X (2019) Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. In: *Proceedings of the 36th International Conference on Machine Learning*
- Mozer MC, Smolensky P (1989) Skeletonization: A technique for trimming the fat from a network via relevance assessment. In: *Advances in Neural Information Processing Systems* 1
- Neyshabur B, Salakhutdinov R, Srebro N (2015a) Path-sgd: Path-normalized optimization in deep neural networks. In: *Advances in Neural Information Processing Systems* 28
- Neyshabur B, Tomioka R, Srebro N (2015b) Norm-based capacity control in neural networks. In: *Proceedings of The 28th Conference on Learning Theory*
- Novikov A, Podoprikin D, Osokin A, et al (2015) Tensorizing neural networks. In: *Advances in Neural Information Processing Systems* 28
- Nowlan SJ, Hinton GE (1992) Simplifying neural networks by soft weight-sharing. *Neural Computation* 4(4):473–493
- NVIDIA (2020) Nvidia a100 tensor core gpu architecture. URL: <https://images.nvidia.com/aem-dam/en-zz/Solutions/data-center/nvidia-ampere-architecture-whitepaper.pdf>, last accessed: 10/31/2022
- Orseau L, Hutter M, Rivasplata O (2020) Logarithmic pruning is all you need. In: *Advances in Neural Information Processing Systems* 33
- Pao YH, Takefuji Y (1992) Functional-link net computing: theory, system architecture, and functionalities. *Computer* 25(5):76–79
- Pao YH, Park GH, Sobajic DJ (1994) Learning and generalization characteristics of the random vector functional-link net. *Neurocomputing* 6(2):163–180
- Parashar A, Rhu M, Mukkara A, et al (2017) Scnn. In: *Proceedings of the 44th Annual International Symposium on Computer Architecture*. ACM
- Park DS, Zhang Y, Chiu C, et al (2020) Specaugment on large scale datasets. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*
- Park J, Li SR, Wen W, et al (2017) Faster cnns with direct sparse convolutions and guided pruning. In: *5th International Conference on Learning Representations*

- Patil SM, Dovrolis C (2021) PHEW: Constructing sparse networks that learn fast and generalize well without training data. In: Proceedings of the 38th International Conference on Machine Learning
- Pensia A, Rajput S, Nagle A, et al (2020) Optimal lottery tickets via subset sum: Logarithmic over-parameterization is sufficient. In: Advances in Neural Information Processing Systems 33
- Peste A, Iofinova E, Vladu A, et al (2021) AC/DC: Alternating compressed/decompressed training of deep neural networks. In: Advances in Neural Information Processing Systems 34
- Peters ME, Neumann M, Iyyer M, et al (2018) Deep contextualized word representations. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies
- Pham H, Dai Z, Xie Q, et al (2021) Meta pseudo labels. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition
- Pool J, Yu C (2021) Channel permutations for n:m sparsity. In: Advances in Neural Information Processing Systems 34
- Poole B, Lahiri S, Raghu M, et al (2016) Exponential expressivity in deep neural networks through transient chaos. In: Advances in Neural Information Processing Systems 29
- Price I, Tanner J (2021) Dense for the price of sparse: Improved performance of sparsely initialized networks via a subspace offset. In: Proceedings of the 38th International Conference on Machine Learning
- Qian X, Klabjan D (2021) A probabilistic approach to neural network pruning. In: Proceedings of the 38th International Conference on Machine Learning
- Qing Y, Zeng Y, Li Y, et al (2020) Deep and wide feature based extreme learning machine for image classification. *Neurocomputing* 412:426–436
- Ramanujan V, Wortsman M, Kembhavi A, et al (2020) What’s hidden in a randomly weighted neural network? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition
- Redman WT, FONOBEROVA M, Mohr R, et al (2022) An operator theoretic view on pruning deep neural networks. In: 10th International Conference on Learning Representations
- Renda A, Frankle J, Carbin M (2020) Comparing rewinding and fine-tuning in neural network pruning. In: 8th International Conference on Learning

Representations

- Robbins H, Monro S (1951) A stochastic approximation method. *The Annals of Mathematical Statistics* 22(3):400–407
- Rosenfeld A, Tsotsos JK (2019) Intriguing properties of randomly weighted networks: Generalizing while learning next to nothing. In: *Conference on Computer and Robot Vision*
- Rosenfeld JS, Frankle J, Carbin M, et al (2021) On the predictability of pruning across scales. In: *Proceedings of the 38th International Conference on Machine Learning*
- Sainath T, Kingsbury B, Sindhvani V, et al (2013) Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*
- Sanh V, Wolf T, Rush AM (2020) Movement pruning: Adaptive sparsity by fine-tuning. In: *Advances in Neural Information Processing Systems* 33
- Saxe A, Koh PW, Chen Z, et al (2011) On random weights and unsupervised feature learning. In: *Proceedings of the 28th International Conference on Machine Learning*
- Saxe AM, McClelland JL, Ganguli S (2014) Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In: *2nd International Conference on Learning Representations*
- Schoenholz SS, Gilmer J, Ganguli S, et al (2017) Deep information propagation. In: *5th International Conference on Learning Representations*
- Schwartz R, Dodge J, Smith NA, et al (2020) Green AI. *Communications of the ACM* 63(12):54–63
- Schwarz J, Jayakumar S, Pascanu R, et al (2021) Powerpropagation: A sparsity inducing weight reparameterisation. In: *Advances in Neural Information Processing Systems* 34
- Soelen RV, Sheppard JW (2019) Using winning lottery tickets in transfer learning for convolutional neural networks. In: *International Joint Conference on Neural Networks*
- Strubell E, Ganesh A, McCallum A (2019) Energy and policy considerations for deep learning in NLP. In: *Proceedings of the 57th Conference of the Association for Computational Linguistics*

- Strubell E, Ganesh A, McCallum A (2020) Energy and policy considerations for modern deep learning research. In: Proceedings of the AAAI Conference on Artificial Intelligence
- Su J, Chen Y, Cai T, et al (2020) Sanity-checking pruning methods: Random tickets can win the jackpot. In: Advances in Neural Information Processing Systems 33
- Sun W, Zhou A, Stuijk S, et al (2021) Dominosearch: Find layer-wise fine-grained n:m sparse schemes from dense neural networks. In: Advances in Neural Information Processing Systems 34
- Sung YL, Nair V, Raffel C (2021) Training neural networks with fixed sparse masks. In: Advances in Neural Information Processing Systems 34
- Sutskever I, Martens J, Dahl G, et al (2013) On the importance of initialization and momentum in deep learning. In: Proceedings of the 30th International Conference on Machine Learning
- Tanaka H, Kunin D, Yamins DL, et al (2020) Pruning neural networks without any data by iteratively conserving synaptic flow. In: Advances in Neural Information Processing Systems 33
- Tinney W, Walker J (1967) Direct solutions of sparse network equations by optimally ordered triangular factorization. Proceedings of the IEEE 55(11):1801–1809
- Ullrich K, Meeds E, Welling M (2017) Soft weight-sharing for neural network compression. In: 5th International Conference on Learning Representations
- Verdenius S, Stol M, Forré P (2020) Pruning via iterative ranking of sensitivity statistics. CoRR abs/2006.00896v2. URL <https://arxiv.org/abs/2006.00896v2>, last accessed: 09/19/2022
- Vischer M, Lange RT, Sprekeler H (2022) On lottery tickets and minimal task representations in deep reinforcement learning. In: 10th International Conference on Learning Representations
- Wang C, Zhang G, Grosse R (2020a) Picking winning tickets before training by preserving gradient flow. In: 8th International Conference on Learning Representations
- Wang CY, Bochkovskiy A, Liao HYM (2021a) Scaled-yolov4: Scaling cross stage partial network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition

- Wang H, Qin C, Zhang Y, et al (2021b) Emerging paradigms of neural network pruning. CoRR abs/2103.06460v2. URL <https://arxiv.org/abs/2103.06460v2>, last accessed: 10/05/2022
- Wang Y, Zhang X, Hu X, et al (2020b) Dynamic network pruning with interpretable layerwise channel selection. In: Proceedings of the AAAI Conference on Artificial Intelligence
- Wang Y, Zhang X, Xie L, et al (2020c) Pruning from scratch. In: Proceedings of the AAAI Conference on Artificial Intelligence
- Wang Z (2020) Sparsert: Accelerating unstructured sparsity on gpus for deep learning inference. In: Proceedings of the ACM International Conference on Parallel Architectures and Compilation Techniques
- Wimmer P, Mehnert J, Condurache AP (2020) FreezeNet: Full performance by reduced storage costs. In: Proceedings of the Asian Conference on Computer Vision
- Wimmer P, Mehnert J, Condurache AP (2021) COPS: Controlled pruning before training starts. In: International Joint Conference on Neural Networks
- Wimmer P, Mehnert J, Condurache AP (2022) Interspace pruning: Using adaptive filter representations to improve training of sparse cnns. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition
- Wu J, Leng C, Wang Y, et al (2016) Quantized convolutional neural networks for mobile devices. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition
- Xiao L, Bahri Y, Sohl-Dickstein J, et al (2018) Dynamical isometry and a mean field theory of cnns: How to train 10, 000-layer vanilla convolutional neural networks. In: Proceedings of the 35th International Conference on Machine Learning
- Xue J, Li J, Gong Y (2013) Restructuring of deep neural network acoustic models with singular value decomposition. In: Interspeech
- Yang Z, Dai Z, Yang Y, et al (2019) Xlnet: Generalized autoregressive pre-training for language understanding. In: Advances in Neural Information Processing Systems 32
- You H, Li C, Xu P, et al (2020) Drawing early-bird tickets: Toward more efficient training of deep networks. In: 8th International Conference on Learning Representations

- Yu H, Edunov S, Tian Y, et al (2020) Playing the lottery with rewards and multiple languages: lottery tickets in rl and nlp. In: 8th International Conference on Learning Representations
- Zagoruyko S, Komodakis N (2016) Wide residual networks. In: Proceedings of the British Machine Vision Conference
- Zhang D, Yang J, Ye D, et al (2018) Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In: Proceedings of the European Conference on Computer Vision
- Zhang S, Stadie BC (2020) One-shot pruning of recurrent neural networks by jacobian spectrum evaluation. In: 8th International Conference on Learning Representations
- Zhang Z, Chen X, Chen T, et al (2021a) Efficient lottery ticket finding: Less data is more. In: Proceedings of the 38th International Conference on Machine Learning
- Zhang Z, Jin J, Zhang Z, et al (2021b) Validating the lottery ticket hypothesis with inertial manifold theory. In: Advances in Neural Information Processing Systems 34
- Zhou A, Yao A, Guo Y, et al (2017) Incremental network quantization: Towards lossless cnns with low-precision weights. In: 5th International Conference on Learning Representations
- Zhou A, Ma Y, Zhu J, et al (2021a) Learning n:m fine-grained structured sparse neural networks from scratch. In: 9th International Conference on Learning Representations
- Zhou H, Lan J, Liu R, et al (2019) Deconstructing lottery tickets: Zeros, signs, and the supermask. In: Advances in Neural Information Processing Systems 32
- Zhou X, Zhang W, Chen Z, et al (2021b) Efficient neural network training via forward and backward propagation sparsification. In: Advances in Neural Information Processing Systems 34
- Zhou X, Zhang W, Xu H, et al (2021c) Effective sparsification of neural networks with global sparsity constraint. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition
- Zhuang T, Zhang Z, Huang Y, et al (2020) Neuron-level structured pruning using polarization regularizer. In: Advances in Neural Information Processing Systems 33

Zhuang Z, Tan M, Zhuang B, et al (2018) Discrimination-aware channel pruning for deep neural networks. In: Advances in Neural Information Processing Systems 31